|CURIEUX|

ACADEMIC REVIEW

November Issue

Issue 13 Part 2

Editing Staff

Chief-In-Editor

George Hasnah

Lead Editor

Amanuel Geremew

Ashley Du

Assisting Editors

Jing Yan Huang

Nidhi Potineni

Sophia Kong

Tony Yuan

Ridit Aggarwal

Manu Sharma

Rishi Gadikota

Eric Wang

Table Of Contents

Page 6: Primordial Black Holes as Dark Matter for Beginners: Evidence, Constraints, Future Steps By Christian Lee

Page 22: Artificial Intelligence in Gastritis: Current Applications, Challenges, Clinical Readiness, and Future Directions By Fuad Sudheer Kannu

Page 43: The Influence of Algorithmically Curated Social Media Content on Consumer Purchasing Behaviour Among Generation Z and Millennials: A Literature Review Paper By Yashita Jain

Page 53: AI-Based Predictive Maintenance With Generative AI Integration for Multi-Sensor Diagnosis and Prediction of Aircraft Gearbox Failures By Vikrant Raj Sathyamurthy

Page 73: The Architecture of Exception: What Legal Exclusions in Singapore Reveal About Global Migration Governance By Ava Rose Attari

Page 90: Protection Without Borders: Glocalization and the Distribution of Safeguarding in the NBA's Markets By Armaan Batra

Page 100: Balancing Potency, Stability, and Solubility In MMAE-Based Antibody-Drug Conjugates: Dual-Payload and Advanced Linker Strategies By Yen Kai Chan

Page 113: Mental Health and Remembering Dreams: A Review on Applying Dream Recall Frequency Analysis to the Transdiagnostic Method By Emily Elizabeth Yam

Page 120: Differences between Early-Onset Alzheimer's Disease and Late-Onset Alzheimer's Disease in Attention Networks of the Brain By William Mao

Page 131: Investigating the Performance of Computational Methods in Predicting Off-target Effects of CRISPR-Cas9: A Review By Ashwath Kannan

Page 148: SiO2 Cloud-phase Ambiguity Through Nucleation and Annealing Kinetics on HD 189733b By Avni Jain

Page 162: Semiconductors and Strategic Alignment: How U.S.-China Semiconductor Competition Has Influenced Taiwan's Foreign Policy Since 2018 By Nidhi Yellamraju

Page 176: The Deviation of the Triplets: a Comparative Study on Estonia, Kazakhstan, Uzbekistan's STEM Education By Wo Xuan Li

Page 192: Beyond the Music: The Engineering Secrets Behind Concert Hall Acoustics By Jaydzia E. Khan

Page 204: The Psychological Effects of First-Year Engineering Course Work By Mark Kassih

Page 220: A Discussion of the Societal Impact of AI Bias in Football Scouting By Anthony Joung Lim

Page 240: The Global Superconducting Quantum Computer Race: A Review of the Achievements and Remaining Challenges that Superconducting Quantum Processors Still Face by Comparing Competitors By Max Chen

Page 254: Can a Vegan Diet Meet the Nutritional Needs of Children to Promote Health? A Narrative Review of Current Evidence By Carlos Orfanidis Cordoba

Page 275: Exploring ADHD Neurochemistry, Pharmacological, and Non-Pharmacological Interventions: A Literature Review By Gaven Z Li

Page 293: From Viral Trends to Business Growth: The Impact of TikTok and Instagram Marketing Strategies on The Entrepreneurial Success of Emerging Dance Businesses in the United States By Annie Li

Page 309: The Effects of Dance Movement Therapy on Executive Function in individuals with Autism Spectrum Disorder By Christopher Tsai

Page 320: The Case for Edible Coatings in Bread Waste Prevention By Ryan Wu

Page 330: Modeling 1D Plasma Transport & Magnetic Shear Dynamics Using Burgers' Equation as an Analytical Proxy By Ethan Meline

Page 342: Systematic Algorithmic Optimization: A case study on Cookie Clicker By Namish Balakerthy Punyakoti

Page 360: What Shapes Entrepreneurial Intent? Active Learning and Multicultural Experiences among Japanese High School Students By Toma Ihara

Page 373: Education-Labor Mismatch And Brain Drain: A Political Economy Analysis Of Youth Exclusion In Jordan And Lebanon By Aarohi Bhattacherjee

Page 381: Evaluating ASPM as a Genetic Modifier of Congenital Zika-Associated Microcephaly By Aanya Sood

Page 392: Feasibility of Nuclear Electric Propulsion for Interplanetary Transport By Hibiki Narita

Page 408: Understanding Inflation in Uzbekistan, 2016–2025 By Akbarxon Nasirov

Page 417: Consumer Behavior Towards Sustainable Cosmetic Products: A Literature Review By Ava Tan Li

Page 430: Evaluating The Longitudinal Structural Decay Of Composite Hockey Sticks By Ethan Zhu

Page 451: The Effects of Feynman Technique, Blurting, and SQ3R on Immediate Recall of Biological Information Among Middle School Students By Maria Astezheva

Page 464: Geometric Predictive Control of Wheeled Mobile Robots with an Actuated Directional Sensor on $SE(2) \times SO(2)$ By Winson Zhang

Page 476: Exploring CRISPR/Cas9 as a Potential Gene Therapy for Tay-Sachs Disease By Lydia Xia

Page 486: Progress and limitations of iPSC-derived Brain Organoids in Neurological Disease Research By Sydney Kroker McCormack

Page 501: Hormonal Fluctuations and Genetic Variability: A Review of the Origin of Perimenopausal Depression By Shanaya Ailsinghani

Page 509: A Comparative Mathematical Analysis of Female Influencers' Impact on E-Commerce Growth in the UAE and South Africa By Zia Jooma

Page 516: Strengths, Limitations, and Modern Evaluation of Mendelian Genetics Research By Aarav Tripathi

Page 528: The Role of Social Connectedness in Veteran Reintegration in the United States By Seira Watanabe

Page 533: Agroforestry and Nature-Based Solutions as Tsunami Protection strategies: a Case Study of Sendai By Rio Ito

Page 553: Why Mid-Sized Firms Lag: Barriers to AI-Based Facility Management Adoption in Indian Commercial Real Estate — Maharashtra as a Case Study By Aadith Palresha

Page 565: Air Quality Perceptions and Public Transport Attitudes: A Literature Review of Delhi By Ira Chinmay Dhamne

Page 577:Which Basic Emotion Increases Brand Recall More: Happiness, Surprise or Fear? By Zeynep Nur Ayan

Page 589: Why Some Words Stick: Contextual Diversity and Free-Recall Memorability Across 576 Words By Samaira Roy

Page 600: The Chemistry Underlying Chronic Stress-Induced Gut Dysbiosis By Rajit Reddi

Page 611: A Narrative Study on Legal Consciousness of Chinese Domestic Workers – From the Perspective of Workers' Lived Experiences By Wenxin Yu

Page 621: Game Theory In Football Market Transfers By Atharva Kumar

Primordial Black Holes as Dark Matter for Beginners: Evidence, Constraints, Future Steps
By Christian Lee

Abstract

Primordial Black Holes (PBHs) have re-emerged as a strong non-particle candidate for dark matter content in the universe. PBHs are formed from high-density fluctuations in the very early universe, spanning a vast range of potential mass windows, from asteroid-sized to multiple times the mass of our Sun. This review combines models of PBH formation using observational evidence, constraints and signatures from microlensing events, and Cosmic Microwave Background (CMB) distortions. We pay close attention to gravitational waves using the Laser Interferometer Gravitational-Wave Observatory (LIGO) as a tool for discovering and constraining PBHs. Additionally, we look at some papers to see how upcoming observatories, such as the Laser Interferometer Space Antenna (LISA), will probe new mass windows. Finally, we address new developments, including the implications of small PBHs, Hawking radiation, and the memory burden effect on black hole lifetimes.

Keywords: Primordial black holes, dark matter, cosmic microwave background, microlensing

Introduction

Dark matter remains one of the greatest mysteries of modern physics, completely invisible to us yet accounting for ~85% of all matter content in our universe (Aghanim et al.). For decades, researchers have searched for new elementary particles and other possibilities to explain dark matter, yet there is no definitive answer. The possibility that dark matter is made of primordial black holes (PBHs) has arisen.

PBHs are an especially attractive candidate because they fall under the category of non-baryonic MAssive Compact Halo Objects (MACHOs). Compared to another leading category of candidates, Weakly Interacting Massive Particles (WIMPs), MACHOs are much more convenient, since they do not require a new particle to be found (Green and Kavanagh 043001).

Unlike conventional black holes, PBHs could span an enormous range of masses. This is because according to the most prominent theory of PBH formation, the mass of a PBH depends primarily on the cosmological horizon mass at the time of its collapse, the collapse occurring when fluctuations stretched outside the horizon during inflation re-enter it after inflation ends (Carr and Kühnel 355–94; Carr 1-19). If they exist in a proper abundance, PBHs could account for some or even all of the dark matter in the universe (Chapline 251). However, decades of observations have constrained possible mass ranges for PBHs, whether through CMB or microlensing observations, making PBHs seem less and less likely as potential candidates (Carr et al., "Observational Evidence" 1–68), discussed more in detail in section 3. However, some mass ranges still exist today, and the new development of gravitational wave astronomy has introduced a new method of detecting PBHs through their gravitational wave signatures (Bird et

al. 201301; Sasaki et al. 061101; Clesse and García-Bellido 142-47). Gravitational waves are ripples in the form of gravitational distortions on the fabric of spacetime by massive accelerating objects (Abbott et al. "Observation" 061102). Because these gravitational distortions are extremely faint, detection requires ultra-precise instruments such as LIGO, Virgo, and KAmioka GRAvitational-wave detector (KAGRA) (Collaboration et al. 074001; Acernese et al. 024001; Akutsu et al. 05A101). However, despite their ongoing collaboration, these instruments have not yet found evidence of PBHs, since they are hard to distinguish from stellar black holes (Abbott et al. 231103). And theoretical developments can reopen mass windows previously thought closed, like the memory burden effect we will discuss in the later parts of this writing.

Are primordial black holes a good candidate for dark matter, and how can we determine this using gravitational wave observations? This question is important because it has roots in the early universe, dark matter, and quantum theory involving black holes, some of the most pronounced and difficult questions modern astrophysics has to offer. The question has become increasingly important because gravitational wave experiments may be able to determine PBH qualities that make them suitable for dark matter.

Recent theoretical frameworks have expanded observational horizons by linking inflationary models to PBH binary systems and the generation of second-order scalar-induced gravitational waves (SIGWs) (Bagui et al. Article 9; Barker et al. 123033), and the construction of new, stronger gravitational wave observatories able to detect more signals (Carr et al., "Observational Evidence" 1–68; Domènech and Sasaki 143001). SIGWs are smaller-scale gravitational waves that formed when first-order scalar perturbations (which are density and pressure ripples in the early universe, not gravitational waves) couple non-linearly (Ananda et al.; Baumann et al.; Inomata and Nakama 043511). SIGWs are important because they allow researchers to probe primordial fluctuations on small physical scales due to SIGWs being sourced by small-scale perturbations, which have amplitudes that CMB and large-scale structure can't reach (Cecchini et al. 123536).

This paper investigates whether PBHs remain a viable candidate for dark matter and evaluates how evolving gravitational wave observations can resolve our question. Unlike existing reviews that focus on advanced topics, such as inflationary models, this paper is designed to build a more general, foundational understanding of the PBH dark matter debate. By combining recent developments, reviewing established observational constraints, and evaluating the sensitivity of upcoming instruments like LISA, this paper outlines how the next generation of gravitational wave experiments will transform the PBH question into a fully testable framework.

Primordial Black Hole Origins

This section will cover the most popular theory as to how PBHs formed in the early inflationary era, and how this formation method makes them suitable dark matter candidates. We will go into how PBHs can span a large range of masses, which opens up our analysis of mass windows in later sections.

One of the major requirements that PBHs must fulfill to be considered a candidate for dark matter is that it must be non-baryonic. In other words, they cannot contribute to the strict limits placed on the universe's population of protons and neutrons. PBHs are non-baryonic because they are thought to have formed before Big Bang Nucleosynthesis (BBN) could occur (Zel'dovich and Novikov 602; Hawking, "Gravitationally Collapsed" 75-8), or before baryons came together to form the first light nuclei (Fields et al. 090001). Because PBHs formed before BBN, the mass locked inside PBHs was removed from the cosmic plasma. Therefore, PBHs can behave as non-baryonic dark matter, allowing PBHs to freely account for the universe's dark matter (Carr et al., "Primordial Black Holes" 083504).

The main mechanism driving PBH generation is the gravitational collapse of primordial density perturbations. During the inflationary epoch, quantum fluctuations were stretched, leaving behind some areas in the cosmic fluid more dense than others. Because space expanded exponentially during inflation, these fluctuations were stretched outside the cosmological horizon. Once inflation ended, the rates of growth of the fluctuations slowed down. However, the cosmological horizon continued to expand, and when these fluctuations crossed back into the horizon, they quickly collapsed (Young and Byrnes 004). The suitably dense regions, or the regions where the density exceeded a critical density threshold ( ) (Harada et al. 084051), formed into PBHs (Carr 1-19; Ivanov 7145–54).

The mass M_PBH of the black hole is dictated by the timing of this horizon re-entry. The mass of a PBH scales proportionally with the total mass contained within the cosmological horizon at the exact moment of collapse.

$$M_{PBH} \cong \gamma M_H(t) \qquad (1)$$

where M_H(t) is the mass inside the cosmological horizon at a moment t, the time elapsed since the Big Bang, and γ is a computed efficiency factor showing that around 20% of horizon mass actually collapses, so $\gamma \cong 0.2$ (Carr 1-19). Because the horizon mass grows linearly with cosmic time, PBHs formed at the Planck time ($t \cong 10^{-43}$ s) would have incredibly small masses, whereas those collapsing during the quantum chromodynamics phase transition ($t \cong 10^{-5}$ s) would reach masses on the stellar scale. This relationship demonstrates why PBHs can theoretically span an immense mass spectrum, unlike stellar black holes (SBHs).

Observational Constraints on PBH Mass and Abundance

By far the strongest pushback on PBHs as dark matter is the fact that in many mass windows, PBH abundance is far too low to be considered a viable candidate. If observations show that PBHs of a certain size would produce too much radiation or disrupt stars, that mass window is closed or restricted to only a tiny fraction of dark matter. By showing where PBHs fail versus where they survive, these constraints determine whether PBHs are a viable, complete solution or a minor contributor to the universe's dark matter. Here, we will explore different surveys such as the Subaru Hyper Suprime-Cam (HSC), LIGO-Virgo-KAGRA (LVK) collaborations, Expérience pour la Recherche d'Objets Sombres (EROS) and the Optical

Gravitational Lensing Experiment (OGLE), which have found significant restrictions on the masses and abundances of PBHs based off of gravitational waves, CMB distortions, and microlensing frequencies. We will go through each window and report findings that show that PBHs have very minimal probability of being dark matter in specific mass ranges. As a side note, M_CO, or compact object mass, is used in place of M_PBH. All PBHs are COs, but not all COs are PBHs, since white dwarves, neutron stars, and low mass red dwarves are also COs. Since the number of PBHs must be less than that of COs, the fraction of PBHs as dark matter (f_PBH) is always less than the fraction of COs as dark matter (f_CO), defined as fractions of the total dark matter density (Green, "Primordial Black Holes as a Dark Matter Candidate" 116494).

Evaporation constraints

Like SBHs, PBHs evaporate by means of Hawking radiation, and it has been calculated that any black hole lighter than ~5×10^(-19) M⊙must have evaporated by now (Hawking, "Particle Creation" 199–220; Page 3260–73). Simply put, in the quantum field, particle-antiparticle pairs are constantly popping into existence and annihilating each other. If a black hole were to suck up one of the members in this pair and the other member escapes, it would escape as radiation. Thus, the black hole would slowly lose energy and slowly shrink until it explodes. Because black holes are usually massive, this process would take unimaginable lengths of time for larger black holes. As it is, the 13.8 billion years that the universe has been existing for has been predicted to allow black holes lighter than 5×10^(-19) M⊙to evaporate. These “light” PBHs would have been perfect dark matter candidates, but as it is, this constraint completely crosses off light PBHs from the list of dark matter candidates, forcing scientists to pursue heavier PBHs, although there will be more on this later.

Microlensing surveys

Microlensing surveys scan dense star fields to detect temporary increases in starlight by microlensing events. Microlensing events are when an object passes in front of a distant object, and the distant object’s light rays are bent and magnified by the object in the foreground (Paczyński 1-5). Now, scientists have made hypothetical scenarios calculating lensing events with PBHs as foreground objects. According to Einstein and future calculations, the size of the lens is determined by the Einstein radius (R_E), which is a measure of how much a large object bends light (Oguri and Blandford 930–44). Accordingly, R_E is proportional to the Einstein crossing time (t_E), which is how long it takes for the lens to travel a distance equal to its R_E (Paczyński 1-5; Mao).

$$R_E \propto M^{(1/2)}, \quad (2)$$

and

$$R_E \propto t_E \quad (3)$$

so

$$t_E \propto M^{(1/2)} \quad (4)$$

By observing t_E, we will be able to constrain masses and their abundances by the frequency and lengths of such events.

The 10^(-11)M⊙ < M_CO < 10^(-6)M⊙ range

In a survey by the HSC on the Andromeda galaxy, while studying microlensing events, scientists assumed that if PBHs constituted 100% of dark matter, there would be around 1,000 events. Instead, they found one, single event (Niikura et al., "Microlensing Constraints" 524–34). Needless to say, PBHs can exist as dark matter at this range only at a very small fraction: f_CO ≤ 0.01, at a 95% confidence level (CL).

The 10^(-6)M⊙ < M_CO < 10^(-3)M⊙ range

In the 5-year OGLE survey of the Galactic bulge, 6 short-lived microlensing events that were of unexplainable origin were discovered (Niikura et al., "Constraints on Earth-mass" 083503). Niikura et al. made a claim using the null hypothesis, or the assumption that all 6 events were caused by random planets rather than PBHs with the mass of planets. Their claim was that the fraction of PBHs that make up dark matter (f_PBH) in this mass range cannot exceed 0.01 ≲ f_CO ≲ 0.1 at a 95% CL, or around 1% to 10%, depending on the mass. This assumption is only a hypothesis (Scholtz and Unwin 051103), so PBHs can actually sit at 1% to 10% in this mass window. However, this is still far from the 100% we are trying to reach.

The 10^(-3)M⊙ < M_CO < 10^(-1)M⊙ range

The EROS survey, specifically the EROS-2 survey, studied tens of millions of stars in the LMC, the Small Magellanic Cloud (SMC), and the Galactic bulge for microlensing events over the course of ~6.7 years (Tisserand and Milsztajn 387–404). EROS-2 found that for objects in the 10^(-3)M⊙ < M_CO < 10^(-1)M⊙ range, COs could account for only 4% of the dark matter halo, at a 95% CL (Tisserand et al. 387–404). These results rule out PBHs in the medium-to-large mass range from making up the entirety of dark matter, but they also leave open the possibility for PBHs outside of this mass window.

Summary

Figure 1 summarizes the current findings we have gone over.

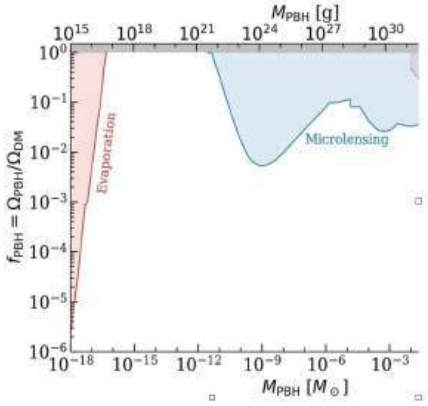

Fig 1: Constraints on f_PBH plotted against M_PBH. Adapted from Villanueva-Domingo et al. 681084.

There are still many gaps missing. For larger COs, these gaps are because microlensing surveys start to reach inefficiently long timescales. Recall that $t_E \propto M^{(1/2)}$. For $M_{CO} \gtrsim 100M_\odot$, these surveys would take years to complete, and would simply take too long. As for the other gaps, we will introduce other methods to bridge them.

Gravitational wave constraints

Instead of microlensing surveys dependent on light, we turn to gravitational wave observations, most notably the LVK collaboration. This field of gravitational waves is still somewhat new, and as such is incredibly promising.

We move on to constraints for $1M_\odot < M_{PBH} < 200M_\odot$. We now use M_{PBH} rather than M_{CO} because we are now dealing with gravitational waves, since PBH gravitational waves are easily distinguishable from other compact objects. If PBHs made up 100% of dark matter at this mass scale ($f_{PBH} = 1$), then PBH mergers would be common. It would take billions of years for orbits to decay, and such an event would manifest as a gravitational wave strain, resulting in a measurable deviation.

$f_{PBH} \neq 1$, but mergers would still occur, and cause immense gravitational waves to be formed. LVK runs look at black hole mergers, and look for any out of the ordinary "chirps", or massive amounts of energy radiated in a single event, since PBH mergers can be different from SBH mergers in terms of spin and eccentricity. As a matter of fact, SBHs cannot be of subsolar mass (Chandrasekhar 81), so if we did find a black hole of subsolar mass, it would be direct evidence of a PBH. In the third run of LVK, LVK O3, after modeling an environment of both SBHs and PBHs, they found that $f_{PBH} < 10^{-3}$ for $1M_\odot < M_{PBH} < 200M_\odot$ at a 90% CL (Andrés-Carcasona et al. "Constraints on primordial black holes from LIGO-Virgo-KAGRA" 023040), showing that the number of existing PBHs must be incredibly small.

CMB constraints

Here, we shift towards using the CMB as counterevidence. During the recombination era (~380,000 years after the Big Bang), when the universe was transitioning from an opaque plasma to a transparent state to create our CMB, PBHs with $M_{PBH} > 100M_\odot$ would be active here, and at this size, they would have a giant gravitational radii. PBHs would be consuming gases, which get compressed and heated up via friction as they fall in. These gases would release lots of high energy radiation, which would tear away electrons from neutral hydrogen atoms and create lots of free floating electrons (Ricotti et al. 829–45; Ali-Haïmoud and Kamionkowski 043534). These electrons would scatter CMB photons and create significant CMB anisotropies, which are temperature variations across different directions in the sky.

Observations of the CMB by satellites like the Planck satellite have not seen any CMB anisotropies, and instead viewed a perfectly clear CMB mapping that aligns with an early universe without such heating (Aghanim et al.). Scientists simulate different M_{CO} and f_{PBH} levels to see how the CMB would look, and found that for $M_{PBH} \approx 100M_\odot$, $f_{PBH} < 10^{-3}$ at a 95% CL (Ali-Haïmoud and Kamionkowski 043534). Any higher mass or fraction, and the

distortion caused by radiation would be noticeable by the Planck satellite. Additionally, for $M_{PBH} > 100M_\odot$, the accretion of a PBH becomes more efficient as the mass increases. An increased efficiency in accretion leads to more radiation being produced from the gasses being pulled into PBHs, leading to more extreme CMB anisotropies. At roughly $10^4M_\odot$, the constraint flattens once accretion becomes Eddington-limited (Ali-Haïmoud and Kamionkowski 043534; Carr et al. 116902). Because we don't see the radioactive aftermath of billions of massive black holes shining in the early universe's gas, we know that PBHs in that heavy mass range must be incredibly rare.

Miscellaneous constraints

Here, we will talk about recent, miscellaneous constraints that reinforce our current constraints. We look into a new kind of microlensing: cluster caustic microlensing. This is when, instead of starlight being magnified by a singular PBH, it is instead magnified, much more powerfully, by the gravity of an entire cluster of galaxies. What makes cluster caustic microlensing different is that when the galaxy cluster passes in front of the star relative to us, the resulting brightness curve can tell us about the galaxy cluster's dark matter halo. If dark matter was completely made of WIMPs, then the curve would be a smooth bell curve, as WIMPs are not gravitationally significant enough to create substantial lensing effects. However, if dark matter is not 100% made of WIMPs, and there are some MACHOs, then there will be jagged spikes in the brightness curves as MACHOs are strong enough to bend and amplify light rays (Oguri et al. 023518; Torralba and Miralda-Escudé).

In studies of the massive star Icarus located behind the MACS J1149 cluster, sharp spikes in its brightness graph were detected. It was then calculated that background stars were feasible causes of these spikes, and concluded that $0.2 \gtrsim f_{PBH}$ at a 95% CL for a $1M_\odot$PBH (Kawai and Oguri). Further studies were conducted and Earendel, the highest redshift star ever detected ($z = 6.2$), was found by the Hubble Space Telescope (HST) (Welch et al. 815–18). Eventually, Mothra, a binary star system magnified over 4,000 times by the galaxy cluster MACS J0416, was discovered by the James Webb Space Telescope (JWST) (Diego et al. A31). A specific study found that by studying Earendel, Mothra, and other cluster lensed stars (and star systems), and studying their layouts and each cluster's critical curves, $M_{CO} > 10^{(-6)}M_\odot$ leads to a constraint of $f_{CO} \lesssim 0.02$ at a 95% CL (Müller and Miralda-Escudé 1579–85).

Additionally, in a later paper on the fourth run of LVK, LVK O4, it was still found that $f_{PBH} < 10^{(-3)}$, but for $10^{(-4)}M_\odot < M_{PBH} < 10^4M_\odot$ at a 90% CL (Andrés-Carcasona et al., "Constraints on Primordial"). These findings also limit the range $10^{-1}M_\odot < M_{PBH} < 1M_\odot$, a window difficult to discern because compact objects with these masses can easily be misidentified for other stellar objects.

Thus, we have synthesized many different papers proving, at high CLs, that PBHs cannot exist in many mass windows. There are a large number of constraints, and a counterargument can be presented that it is much easier to say that PBHs are not good dark matter candidates. However, PBHs still do remain extremely compelling dark matter candidates because, again,

they do not require the discovery of a new elementary particle, and instead work very well with the modern, widely accepted inflationary model in the early universe. PBH formation follows along with our well-established gravitational physics, and continues to do so. No new physics is required, and they are non-baryonic just like dark matter is. Additionally, not all windows have been completely constrained.

Recent Advances and Future Directions

In this section, we will detail the one remaining mass window, then move on to future advancements in theory and technology, and how they could possibly constrain, or even expand, our PBH mass windows so that we can better understand PBHs as dark matter.

The asteroid-mass window

One window remains: $5\times10^{-19}\ M_\odot < M_{PBH} < 10^{-11}\ M_\odot$,

more formally known as the asteroid-mass window. These objects are too small for gravitational lensing and have not evaporated by Hawking radiation, thus meaning that technically, it is possible that $f_{PBH} = 1$ (Gorton and Green 032; Montero-Camacho et al. 031). This window is the last to remain, and with more sensitive future observatories, we might be able to probe this window and determine for once and for all whether PBHs could possibly account for 100% of dark matter in a specific range of masses. This is the most probable window, and future advancements such as new spectroscopers and telescopes will allow us to probe it. Now, we will look at how scientists are working on more advanced detection methods to better seek PBHs.

Scalar-induced gravitational waves

While PBHs cannot be seen directly at all, the dawn of gravitational wave astronomy has provided a method to confirm their existence. When two black holes orbit each other, they form a binary system. As they spin around, their immense gravity makes waves on the fabric of spacetime, radiating energy outward in the form of gravitational waves until their collision, which releases one last burst of gravitational waves. The stochastic gravitational wave background (SGWB) generated by PBH binary mergers may be distinguishable through the evolution of the merger rate (Mandic et al. 201102; Wang et al. 191102; Mukherjee and Silk 3977). These are different from the “chirps” that we saw resembling single events in Section 3; rather, these gravitational waves live in the past, so we look into the SGWB. Because these early universe binaries do not follow typical stellar paths, their overlapping signals change the standard spectrum of the cosmic background (Bagui and Clesse 101115).

A signature of PBH formation is the existence of a secondary stochastic background (Saito and Yokoyama 161101). The large scale primordial perturbations required for the gravitational collapse of PBHs act as a source for a secondary shockwave of gravitational waves (Joana 023518). These SIGWs are generated during the radiation dominated era at the exact moment the perturbation re-enters the cosmic horizon. The secondary shockwave creates a

dual-layered signature: the low frequency SIGWs, (typically in the nanohertz to millihertz band) and the high frequency SGWB (in the hertz to kilohertz bands) (Kushwaha).

Because both signals share an identical origin in the early universe's inflationary perturbations, their amplitudes and profiles are linked (Joana 023518). Currently, scientists are working on observing whether the SGWB can match the SIGW background associated with the birth of a PBH through the LVK collaboration (Kushwaha). While this is a work in progress, there has been some additional work on cutting back on some constraints.

The memory burden effect

A new quantum theory might reverse the mass window that we excluded due to Hawking radiation. The memory burden effect states that as a black hole radiates its mass from Hawking radiation, it stores information in quantum microstates. Since information cannot be destroyed, once the black hole radiates roughly half of its initial mass, the burden of all of the information heavily slows down the remaining radiation exponentially (Dvali et al. 103523). The burden slows down the decay so much, that ultralight PBHs with mass ranges $5\times10^{-28}\ M_\odot < M_{PBH} < 10^{-26}\ M_\odot$ can feasibly exist today (Alexandre et al. 036004; Thoss et al. 451-59). This could potentially open up a newer window, since these black holes are considered so tiny that our current observational techniques have no way of detecting them. This would be a huge reversal on our current constraints, and open up even more possibilities for these exotic black holes to play a role in the composition of dark matter. However, this effect will continue to remain a theory without the advent of newer, stronger observational techniques.

Future observatories to improve constraints on PBHs

Currently, our instruments have reached their maximum sensitivities. If we want to discover evidence of a PBH, we will have to rely on better interferometers and telescopes for a decisive conclusion.

Firstly, we have LISA. LISA is a space-based interferometer that is planned to launch in 2035. This interferometry will detect gravitational waves from roughly 0.1 mHz to 1 Hz, which cannot be detected by current ground-based observatories like LIGO (Amaro-Seoane et al.). This will allow us to detect faint planetary mass PBH binary systems, or even stochastic backgrounds to help us find PBH evidence.

Next is the Einstein Telescope and Cosmic Explorer. These are both ground-based telescopes, but massively improved in their own ways. The Einstein Telescope, detecting from ~1 Hz to ~10 kHz, will be 10 times more sensitive than most current setups, and unlike LIGO, the Einstein Telescope will be situated deep underground. This placement is to prevent any seismic noise from earthquakes or other ground vibrations, which LIGO, on the surface, is currently vulnerable to, allowing the telescope to view much farther into the past (Grado 012041). The Cosmic Explorer, detecting around the same range of frequencies, is a larger version of LIGO. The arms of LIGO are around 4 km long; the Cosmic Explorer's arms will be around 40 km long, allowing for gravitational signals to be stronger and easier to measure, as

they will stretch more space across the Cosmic Explorer's long arms (Evans et al. 90). The Cosmic Explorer will also use quantum technology so that it can be optimized across different frequencies (Srivastava et al. 22). Both the Einstein Telescope and the Cosmic Explorer will work as complements for each other, as the Einstein Telescope specifies in low-frequency noise isolation, and the Cosmic Explorer specializes in a high signal/noise ratio.

Finally, we have the recently opened Vera C. Rubin Observatory, which will use optical surveys to seek out tiny microlensing events, which may be able to help probe the upper outer edges of the elusive asteroid mass window (Ivezić et al. 111).

Closing

Discussion

The evidence summarized in this review suggests that despite the fact that definitive PBH evidence has not been found, they still remain incredibly attractive candidates for dark matter. Current constraints have made it unlikely that they account for all of the dark matter, and different methods have been closing off different portions of PBH mass windows. Hawking radiation has closed off ultralight masses, microlensing constrains many planetary to substellar masses, gravitational waves have strongly constrained stellar to intermediate masses, and CMB observations have wiped out massive PBHs. No observation completely takes PBHs out, but the combination of many mass windows having small fractions of PBH as dark matter make the argument difficult to sustain.

One window, the asteroid window, still remains a last hope for PBHs. It could easily be closed off by our new observatories, and the PBH argument would be difficult to sustain. Or, we could find evidence that points to f_PBH = 1, and give us our first real insight on dark matter. The new field of gravitational wave astronomy, as mentioned, is still relatively new. Yet, the field provides indirect evidence by allowing us to test primordial density fluctuations that led to PBHs in the first place. In particular, SIGWs may contain information about the primordial perturbations responsible for PBH formation. Of course, gravitational waves cannot provide definitive evidence supporting the existence of PBHs, as SBHs can also produce gravitational waves. Gravitational waves instead can provide evidence for a primordial origin if they have differences from SBH gravitational waves that are difficult to explain.

Additionally, this review has limitations. Some constraints are model dependent, while others are not. For example, CMB restraints are still based on uncertain assumptions and models, and PBH abundance can change by orders of magnitude if model abundance limits are relaxed (Facchinetti et al. 043537). And this field is rapidly evolving. We can see the impact of quantum information theory on our Hawking radiation constraints already. These constraints are not permanent, and this review may become outdated soon.

So, what is the answer to our research question? Overall, PBHs are plausible, yet constrained candidates for dark matter. Many observational constraints have left few mass windows open, most notably, the asteroid-mass window. We can constrict or relax these

constraints based on gravitational wave observations, as they provide constraints from black hole mergers and potential evidence through SIGWs. Future gravitational wave observatories can further test the remaining possibilities, but quantum theories can open up new ones. The future of PBHs lies in if future observations can find evidence for them as dark matter in the narrow windows we have left.

Acknowledgements

I would like to thank Catherine Petretti, the TA for my class, for helping me reinforce my paper and giving useful feedback, and Dr. Ryan Farber, the teacher for my class, for helping me to pick a topic, guide me on my publishing journey, and provide specialized help for finalizing my paper. I would also like to thank the YouTube channel melodysheep for getting me started in my passion for astrophysics, and of course, Google Scholar for allowing me to find all of these cool papers.

Works Cited

Aasi, Junaid et al. “Advanced Ligo.” Classical and Quantum Gravity, vol. 32, no. 7, 2015, p. 074001.

Abbott, B. P., R. Abbott, T. D. Abbott, M. R. Abernathy, et al. “Observation of Gravitational Waves from a Binary Black Hole Merger.” Physical Review Letters, vol. 116, no. 6, Feb. 2016, p. 061102. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.116.061102.

Abbott, B. P., R. Abbott, T. D. Abbott, F. Acernese, et al. “Search for Subsolar-Mass Ultracompact Binaries in Advanced LIGO’s First Observing Run.” Physical Review Letters, vol. 121, no. 23, Dec. 2018, p. 231103. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.121.231103.

Acernese, Fausto, et al. “Advanced Virgo: A Second-Generation Interferometric Gravitational Wave Detector.” Classical and Quantum Gravity, vol. 32, no. 2, 2015, p. 024001.

Aghanim, N. “Planck 2018 Results: VI. Cosmological Parameters.” Astronomy & Astrophysics, vol. 641, 2020, p. A6.

Akutsu, T., et al. “Overview of KAGRA: Detector Design and Construction History.” Progress of Theoretical and Experimental Physics, vol. 2021, no. 5, 2021, p. 05A101.

Alexandre, Ana, et al. “New Mass Window for Primordial Black Holes as Dark Matter from the Memory Burden Effect.” Physical Review D, vol. 110, no. 3, Aug. 2024, p. 036004. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.110.036004.

Ali-Haïmoud, Yacine, and Marc Kamionkowski. “Cosmic Microwave Background Limits on Accreting Primordial Black Holes.” Physical Review D, vol. 95, no. 4, Feb. 2017, p. 043534. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.95.043534.

Amaro-Seoane, Pau, et al. “Laser Interferometer Space Antenna.” arXiv:1702.00786, arXiv, 23 Feb. 2017. arXiv.org, https://doi.org/10.48550/arXiv.1702.00786.

Ananda, Kishore N., et al. “Cosmological Gravitational Wave Background from Primordial Density Perturbations.” Physical Review D, vol. 75, no. 12, June 2007, p. 123518. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.75.123518.

Andrés-Carcasona, M., A. J. Iovino, et al. “Constraints on Primordial Black Holes from LIGO-Virgo-KAGRA O3 Events.” Physical Review D, vol. 110, no. 2, July 2024, p. 023040. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.110.023040.

Andrés-Carcasona, M., A. J. Iovino, et al. “Constraints on Primordial Black Holes from the First Part of LIGO-Virgo-KAGRA Fourth Observing Run.” arXiv:2605.15749, arXiv, 15 July 2026. arXiv.org, https://doi.org/10.48550/arXiv.2605.15749.

Bagui, Eleni, et al. “Primordial Black Holes and Their Gravitational-Wave Signatures.” Living Reviews in Relativity, vol. 28, no. 1, Jan. 2025, p. 1. DOI.org (Crossref), https://doi.org/10.1007/s41114-024-00053-w.

Bagui, Eleni, and Sebastien Clesse. “A Boosted Gravitational Wave Background for Primordial Black Holes with Broad Mass Distributions and Thermal Features.” Physics of the Dark Universe, vol. 38, 2022, p. 101115.

Barker, William, et al. “Inflationary and Gravitational Wave Signatures of Small Primordial Black Holes as Dark Matter.” Physical Review D, vol. 111, no. 12, June 2025, p. 123033. DOI.org (Crossref), https://doi.org/10.1103/4hrv-zfch.

Baumann, Daniel, et al. “Gravitational Wave Spectrum Induced by Primordial Scalar Perturbations.” Physical Review D, vol. 76, no. 8, Oct. 2007, p. 084019. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.76.084019.

Bird, Simeon, et al. “Did LIGO Detect Dark Matter?” Physical Review Letters, vol. 116, no. 20, May 2016, p. 201301. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.116.201301.

Carr, Bernard, Florian Kühnel, et al. “Primordial Black Holes as Dark Matter.” Physical Review D, vol. 94, no. 8, Oct. 2016, p. 083504. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.94.083504.

Carr, Bernard J., et al. “Observational Evidence for Primordial Black Holes: A Positivist Perspective.” Physics Reports, vol. 1054, 2024, pp. 1–68.

Carr, Bernard J. “The Primordial Black Hole Mass Spectrum.” Astrophysical Journal, Vol. 201, Oct. 1, 1975, Pt. 1, p. 1-19. Research Supported by the Science Research Council of England, vol. 201, 1975, pp. 1–19.

Carr, Bernard, and Florian Kühnel. “Primordial Black Holes as Dark Matter: Recent Developments.” Annual Review of Nuclear and Particle Science, vol. 70, no. 1, Oct. 2020, pp. 355–94. DOI.org (Crossref), https://doi.org/10.1146/annurev-nucl-050520-125911.

Cecchini, Chiara, et al. “Forecasting Constraints on Scalar-Induced Gravitational Waves with Future Pulsar Timing Array Observations.” Physical Review D, vol. 111, no. 12, June 2025, p. 123536. DOI.org (Crossref), https://doi.org/10.1103/nxx5-gx7d.

Chandrasekhar, Subrahmanyan. “The Maximum Mass of Ideal White Dwarfs.” Astrophysical Journal, Vol. 74, p. 81, vol. 74, 1931, p. 81.

Chapline, George F. “Cosmological Effects of Primordial Black Holes.” Nature, vol. 253, no. 5489, 1975, pp. 251–52.

Clesse, Sebastien, and Juan García-Bellido. “The Clustering of Massive Primordial Black Holes as Dark Matter: Measuring Their Mass Distribution with Advanced LIGO.” Physics of the Dark Universe, vol. 15, 2017, pp. 142–47.

Diego, Jose M., et al. “JWST’s PEARLS: Mothra, a New Kaiju Star at Z= 2.091 Extremely Magnified by MACS0416, and Implications for Dark Matter Models.” Astronomy & Astrophysics, vol. 679, 2023, p. A31.

Domènech, Guillem, and Misao Sasaki. “Probing Primordial Black Hole Scenarios with Terrestrial Gravitational Wave Detectors.” Classical and Quantum Gravity, vol. 41, no. 14, 2024, p. 143001.

Dvali, Gia, et al. “Black Hole Metamorphosis and Stabilization by Memory Burden.” Physical Review D, vol. 102, no. 10, Nov. 2020, p. 103523. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.102.103523.

Facchinetti, Gaétan, et al. “Relaxing CMB Bounds on Primordial Black Holes: The Role of Ionization Fronts.” Physical Review D, vol. 107, no. 4, Feb. 2023, p. 043537. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.107.043537.

Fields, Brian D., et al. “Big-Bang Nucleosynthesis.” Chinese Physics C, vol. 38, no. 9, Aug. 2014, p. 090001. arXiv.org, https://doi.org/10.1088/1674-1137/38/9/090001.

Gorton, Matthew, and Anne Green. “How Open Is the Asteroid-Mass Primordial Black Hole Window?” SciPost Physics, vol. 17, no. 2, 2024, p. 032.

Grado, Aniello. “Einstein Telescope, the Future Generation of Ground Based Gravitational Wave Detectors.” Journal of Physics: Conference Series, vol. 2429, no. 1, 2023, p. 012041. Google Scholar, https://iopscience.iop.org/article/10.1088/1742-6596/2429/1/012041/meta.

Green, Anne M. “Primordial Black Holes as a Dark Matter Candidate-a Brief Overview.” Nuclear Physics B, vol. 1003, 2024, p. 116494.

Green, Anne M., and Bradley J. Kavanagh. “Primordial Black Holes as a Dark Matter Candidate.” Journal of Physics G: Nuclear and Particle Physics, vol. 48, no. 4, 2021, p. 043001.

Hall, Evan D. “Cosmic Explorer: A next-Generation Ground-Based Gravitational-Wave Observatory.” Galaxies, vol. 10, no. 4, 2022, p. 90.

Harada, Tomohiro, et al. “Threshold of Primordial Black Hole Formation.” Physical Review D, vol. 88, no. 8, Oct. 2013, p. 084051. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.88.084051.

Hawking, S. W. “Particle Creation by Black Holes.” Communications In Mathematical Physics, vol. 43, no. 3, Aug. 1975, pp. 199–220. DOI.org (Crossref), https://doi.org/10.1007/BF02345020.

Hawking, Stephen. “Gravitationally Collapsed Objects of Very Low Mass.” Monthly Notices of the Royal Astronomical Society, vol. 152, no. 1, 1971, pp. 75–78.

Inomata, Keisuke, and Tomohiro Nakama. “Gravitational Waves Induced by Scalar Perturbations as Probes of the Small-Scale Primordial Spectrum.” Physical Review D, vol. 99, no. 4, Feb. 2019, p. 043511. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.99.043511.

Ivanov, P. “Nonlinear Metric Perturbations and Production of Primordial Black Holes.” Physical Review D, vol. 57, no. 12, June 1998, pp. 7145–54. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.57.7145.

Ivezić, Željko, et al. “LSST: From Science Drivers to Reference Design and Anticipated Data Products.” The Astrophysical Journal, vol. 873, no. 2, 2019, p. 111.

Joana, Cristian. “Scalar Induced Gravitational Waves Signaling Primordial Black Hole Dark Matter.” arXiv:2504.18237, arXiv, 25 Apr. 2025. arXiv.org, https://doi.org/10.48550/arXiv.2504.18237.

Kawai, Hiroki, and Masamune Oguri. “Constraints on Primordial Black Holes from the Observed Number of Icarus-like Ultrahigh Magnification Events.” arXiv Preprint arXiv:2411.13816, 2024. Google Scholar, https://arxiv.org/abs/2411.13816.

Kushwaha, Ashu. “Gravitational Waves from Primordial Black Holes: Connecting Low-Frequency Scalar-Induced Signatures to High-Frequency Binary Mergers.” arXiv:2607.01818, arXiv, 17 July 2026. arXiv.org, https://doi.org/10.48550/arXiv.2607.01818.

Mandic, Vuk, et al. “Stochastic Gravitational-Wave Background Due to Primordial Binary Black Hole Mergers.” Physical Review Letters, vol. 117, no. 20, Nov. 2016, p. 201102. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.117.201102.

Mao, Shude. “Introduction to Gravitational Microlensing.” arXiv:0811.0441, arXiv, 4 Nov. 2008. arXiv.org, https://doi.org/10.48550/arXiv.0811.0441.

Montero-Camacho, Paulo, et al. “Revisiting Constraints on Asteroid-Mass Primordial Black Holes as Dark Matter Candidates.” Journal of Cosmology and Astroparticle Physics, vol. 2019, no. 08, 2019, pp. 031–031.

Mukherjee, Suvodip, and Joseph Silk. “Can We Distinguish Astrophysical from Primordial Black Holes via the Stochastic Gravitational Wave Background?” Monthly Notices of the Royal Astronomical Society, vol. 506, no. 3, 2021, pp. 3977–85.

Müller, Claudi Vall, and Jordi Miralda-Escudé. “Limits on Dark Matter Compact Objects Implied by Supermagnified Stars in Lensing Clusters.” Monthly Notices of the Royal Astronomical Society, vol. 536, no. 2, 2025, pp. 1579–85.

Niikura, Hiroko, Masahiro Takada, Shuichiro Yokoyama, et al. “Constraints on Earth-Mass Primordial Black Holes from OGLE 5-Year Microlensing Events.” Physical Review D, vol. 99, no. 8, Apr. 2019, p. 083503. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.99.083503.

Niikura, Hiroko, Masahiro Takada, Naoki Yasuda, et al. “Microlensing Constraints on Primordial Black Holes with Subaru/HSC Andromeda Observations.” Nature Astronomy, vol. 3, no. 6, 2019, pp. 524–34.

Oguri, Masamune, et al. “Understanding Caustic Crossings in Giant Arcs: Characteristic Scales, Event Rates, and Constraints on Compact Dark Matter.” Physical Review D, vol. 97, no. 2, Jan. 2018, p. 023518. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.97.023518.

Oguri, Masamune, and Roger D. Blandford. “What Is the Largest Einstein Radius in the Universe?” Monthly Notices of the Royal Astronomical Society, vol. 392, no. 2, 2009, pp. 930–44.

Paczynski, Bohdan. “Gravitational Microlensing by the Galactic Halo.” Astrophysical Journal, Part 1 (ISSN 0004-637X), Vol. 304, May 1, 1986, p. 1-5., vol. 304, 1986, pp. 1–5.

Page, Don N. “Particle Emission Rates from a Black Hole. II. Massless Particles from a Rotating Hole.” Physical Review D, vol. 14, no. 12, Dec. 1976, pp. 3260–73. DOI.org (Crossref), https://doi.org/10.1103/PhysRevD.14.3260.

Ricotti, Massimo, et al. “Effect of Primordial Black Holes on the Cosmic Microwave Background and Cosmological Parameter Estimates.” The Astrophysical Journal, vol. 680, no. 2, 2008, pp. 829–45.

Saito, Ryo, and Jun'ichi Yokoyama. "Gravitational-Wave Background as a Probe of the Primordial Black-Hole Abundance." Physical Review Letters, vol. 102, no. 16, Apr. 2009, p. 161101. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.102.161101.

Sasaki, Misao, et al. "Primordial Black Hole Scenario for the Gravitational-Wave Event GW150914." Physical Review Letters, vol. 117, no. 6, Aug. 2016, p. 061101. arXiv.org, https://doi.org/10.1103/PhysRevLett.117.061101.

Scholtz, Jakub, and James Unwin. "What If Planet 9 Is a Primordial Black Hole?" Physical Review Letters, vol. 125, no. 5, July 2020, p. 051103. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.125.051103.

Srivastava, Varun, et al. "Science-Driven Tunable Design of Cosmic Explorer Detectors." The Astrophysical Journal, vol. 931, no. 1, 2022, p. 22.

Thoss, Valentin, et al. "Breakdown of Hawking Evaporation Opens New Mass Window for Primordial Black Holes as Dark Matter Candidate." Monthly Notices of the Royal Astronomical Society, vol. 532, no. 1, 2024, pp. 451–59.

Tisserand, Patrick, et al. "Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds." Astronomy & Astrophysics, vol. 469, no. 2, 2007, pp. 387–404.

Tisserand, Patrick, and Alain Milsztajn. "Preliminary Results on Galactic Dark Matter from the Complete EROS2 Microlensing Survey." arXiv:astro-ph/0501584, arXiv, 8 Mar. 2005. arXiv.org, https://doi.org/10.48550/arXiv.astro-ph/0501584.

Torralba, Gabriel, and Jordi Miralda-Escudé. "Supermagnified Stars in Lensing Clusters and Small-Scale Structure in the Dark Matter." arXiv Preprint arXiv:2607.14034, 2026. Google Scholar, https://arxiv.org/abs/2607.14034.

Villanueva-Domingo, Pablo, et al. "A Brief Review on Primordial Black Holes as Dark Matter." Frontiers in Astronomy and Space Sciences, vol. 8, 2021, p. 681084.

Wang, Sai, et al. "Constraints on the Primordial Black Hole Abundance from the First Advanced LIGO Observation Run Using the Stochastic Gravitational-Wave Background." Physical Review Letters, vol. 120, no. 19, May 2018, p. 191102. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.120.191102.

Welch, Brian, et al. "A Highly Magnified Star at Redshift 6.2." Nature, vol. 603, no. 7903, 2022, pp. 815–18.

Young, Sam, and Christian T. Byrnes. "Initial Clustering and the Primordial Black Hole Merger Rate." Journal of Cosmology and Astroparticle Physics, vol. 2020, no. 03, 2020, pp. 004–004.

Zel'Dovich, Ya B., and I. D. Novikov. "The Hypothesis of Cores Retarded during Expansion and the Hot Cosmological Model." Soviet Astronomy, Vol. 10, p. 602, vol. 10, 1967, p. 602.

Artificial Intelligence in Gastritis: Current Applications, Challenges, Clinical Readiness, and Future Directions Fuad Sudheer Kannu

Abstract

Gastritis is an inflammatory condition of the stomach lining and may be difficult to recognize, as symptoms are often vague and nonspecific and can be seen with other conditions. As AI and ML are being utilized more extensively in the healthcare system, researchers are exploring the possibilities of AI for gastritis diagnosis and management. This review will discuss the current uses of AI in gastritis and what possible areas for further research exist, mainly related to stress-associated gastritis. Relevant studies were identified through literature searches and reference screening. The selected studies were grouped into five application areas: endoscopic diagnosis, digital pathology and biopsy analysis, severity assessment, detection of Helicobacter pylori-associated gastritis (a form of gastritis linked to bacterial infection), and prediction of progression from atrophic gastritis to gastric cancer. Studies were then compared based on the data used, AI methods, reported performance, and clinical relevance. The studies reviewed in this paper reported promising performance across different areas of gastritis care, including endoscopic diagnosis, digital pathology, severity assessment, H. pylori-related applications, and gastric cancer risk prediction. However, the studies show different levels of evidence, and there are several outstanding issues such as data quality, external validation, explainability, and real-world readiness for routine clinical care. There is a significant number of studies conducted on the relationship between stress, gastric health, and the gut-brain axis, but there are no major attempts at applying AI in stress-associated gastritis. Based on the available evidence, AI appears useful as a clinical decision support tool in the diagnosis and management of gastritis. While many findings are promising, additional testing in diverse clinical settings is still needed. Future research may also examine whether AI can help identify patterns associated with stress-associated gastritis by combining clinical, lifestyle, and psychological factors.

Keywords Gastritis, Artificial Intelligence, Machine Learning, Digital Pathology, Endoscopy, Helicobacter pylori, Gastric Cancer, Clinical Decision Support, Gut-Brain Axis, Stress-Associated Gastritis.

1. Introduction

Gastritis is a common inflammatory condition of the stomach lining and remains an important health concern worldwide. Although some individuals experience symptoms such as abdominal discomfort, nausea, and indigestion, others may have few or no significant symptoms at all. Symptoms may vary considerably from one patient to another, making it difficult to recognize gastritis based on symptoms alone. For that reason, diagnosis often requires additional clinical evaluation, endoscopy, and biopsy findings. There were approximately 27.2 million incident cases and 38.3 million total cases of the combined category of gastritis and duodenitis worldwide, according to the Global Burden of Disease (GBD) analysis (L. Wang et al. 2570).

Although these numbers reflect the combined burden of gastritis and duodenitis, regional studies make it clear that gastritis remains very common. According to a population study conducted in Sweden, 40.2% of the adults in the cohort who had undergone endoscopy and biopsies were diagnosed with gastritis (Zuzek et al. 528). Chronic gastritis increases the chances of progression into precancerous gastric changes, including gastric atrophy and intestinal metaplasia, which in turn increases the risk of gastric cancer if left undetected or untreated. In fact, Song et al.'s study showed that among patients who undergo gastroscopy (a type of endoscopy that focuses only on the upper digestive tract) with biopsy for clinical indications, approximately 1 in 50 with atrophic gastritis and 1 in 39 with intestinal metaplasia will develop gastric cancer within 20 years (Song et al.). Particular attention is often given to chronic atrophic gastritis and H. pylori-associated gastritis because both are associated with an increased risk of gastric cancer. Early detection and appropriate management are therefore very important (Kim et al.; Turtoi et al.; Zou et al.; Kumar et al. 527). These factors support the importance of accurate diagnosis and demonstrate why AI and ML are being researched as aids for more consistent clinical decision-making. AI is driving the healthcare revolution.

AI systems can evaluate large datasets to identify patterns that assist in clinical decision-making as well as improve diagnosis. Gastritis serves as a valuable environment for this type of research, as it often involves the integration of endoscopy, pathology, radiology, lab tests, and patient health records. Technology can be used to make more efficient decisions from combined information that traditional systems would do separately. The earlier detection of a disease and consistent clinical decision-making may be made easier by looking at these data in conjunction.

Many of the studies reviewed in this paper used AI to analyze endoscopic images and biopsy-related findings. The studies reviewed in this paper reported quantitative performance across different gastritis applications, including sensitivity up to 94%, specificity up to 96%, accuracy up to 95%, and AUC values up to 0.98, as summarized in Tables 2-5 (Turtoi et al.). These findings are promising, even though their clinical readiness requires further examination, as discussed in section 5.

There has been growing interest in the gut-brain axis in recent years. The gut-brain axis is the two-way communication network that connects the gastrointestinal system and the brain, allowing the brain to influence gastrointestinal health and vice versa. This explains how psychological aspects can affect digestive health. Multiple studies reviewed in this paper help to understand the impact of stress on gastrointestinal health and its potential role in the development of digestive disorders (Leigh et al. 4492). Stress-associated gastritis is gastritis that may be influenced by psychological stress through a combination of biological, behavioral, and gut-brain-axis mechanisms. Even though stress is not considered the sole cause of gastritis, growing evidence suggests that it may contribute to gastric disease severity and gastric health. Psychogastroenterology is therefore relevant because it examines how mental wellbeing and digestive health interact.

Even though AI research in gastroenterology and psychogastroenterology has advanced considerably, they have mostly developed separately. AI studies in gastritis focus on diagnosis and disease management. But psychogastroenterology studies show how stress and gut-brain interactions affect digestive health. This gap creates an opportunity to explore how advances in AI and ML for gastritis may inform future research into stress-associated gastritis.

Some medical and pathology terms used in this review may be unfamiliar to readers without a healthcare background. To address that, such terms have been provided with medical definitions along with plain language explanations in Table 1.

Term	Medical Description	Plain Language Explanation
H. pylori (Helicobacter pylori)	A spiral-shaped, Gram-negative bacterium that colonizes the gastric mucosa and causes chronic active gastritis. Infection is also associated with peptic ulcer disease, gastric adenocarcinoma and gastric mucosa-associated lymphoid tissue lymphoma.	Common stomach bacteria that can damage the stomach lining over time.
Gastric Atrophy (Atrophic Gastritis)	Loss of the normal gastric glands in the stomach mucosa, with or without their replacement by metaplastic epithelium, as a result of chronic inflammation.	The stomach lining becomes weaker and thinner after years of damage, from various infections.
Intestinal Metaplasia	A metaplastic change in which the normal gastric epithelium is replaced by intestinal-type epithelium, including mucin-producing goblet cells.	The stomach lining starts changing into a different type of tissue that is not normally found there.
Gut-Brain Axis	A bidirectional communication network between the gastrointestinal tract and the central nervous system involving neural, endocrine, immune and metabolic signaling pathways.	The two-way communication mechanism between brain and digestive system.
Intestinal Permeability	The functional property of the intestinal epithelial barrier that regulates the passage of substances from the intestinal lumen across the mucosa into the internal environment.	How easily materials can pass through the intestinal lining and enter body and blood stream.

Gut Microbiota Composition	The types, relative abundance and diversity of microorganisms, including bacteria, archaea, fungi and viruses, that inhabit the gastrointestinal tract.	Mix of microorganisms living in the digestive tract, including healthy bacteria (probiotic), harmful bacteria, fungi, etc..
Endoscopy	A procedure used to examine the interior of the gastrointestinal tract with an endoscope, a flexible tube equipped with a light and camera.	A doctor uses a small camera to look inside the stomach.
Histopathology	The study and diagnosis of disease through microscopic examination of cells and tissues, usually obtained by biopsy or surgery.	Looking at tissues collected from the human body under a microscope to find signs of disease.
Digital Pathology	The practice of pathology using digitized glass slides and associated data that can be viewed, managed, analyzed and shared electronically.	Converting microscope slides into digital images that can be viewed and analyzed on a computer.

Table 1. Key Medical Terms Used in this Review

This review paper aims to examine how AI and ML systems are used in the diagnosis and management of gastritis, with attention to clinical applications, explainability, limitations, and how these technologies could be extended to stress-associated gastritis as a future research direction.

In this review paper, I have reviewed studies on; 1) Current Applications of AI and ML in gastritis care; 2) Whether the AI models used in gastritis are ready for routine clinical care by focusing on factors that affect their real-world use; 3) The role of stress and gut-brain axis in gastritis; and 4) Future research opportunities that involves the use of AI in stress-associated gastritis specifically.

2. Methods

This review paper examines applications and limitations of AI and ML in gastritis research and clinical care. Relevant studies conducted in this field were found using scholarly search and bibliography screening, which began with the finding of a systematic review paper focused on the role of AI in automated diagnosis of gastritis (Turtoi et al.). Additional studies were identified through reference screening and keyword searches using terms like gastritis, AI, ML, deep learning, endoscopy, digital pathology, H. pylori, gastric cancer, psychological stress, and gut-brain axis.

The studies on AI for gastritis, stress-related gastrointestinal health, and stress-associated gastritis were reviewed. The studies were grouped according to the clinical area they addressed, such as gastritis detection from endoscopic images or prediction of progression from atrophic gastritis to gastric cancer. Datasets, AI methods, and reported results were compared across the selected studies. This review examined whether the models had also been tested beyond the hospitals and patient populations in which they were developed. Some studies raised concerns mostly across the area of validation and real-world applicability. Special focus was also given to stress-associated gastritis, an area where attempts to apply AI have been limited.

Data collected from the selected studies consist of sample size, data source, type of medical data, AI approach, targeted gastritis condition, reported outcomes, and study interpretation. The paper also examined model effectiveness, clinical relevance, and limitations reported by the studies reviewed. Since the previously conducted studies used different datasets, methodologies, and evaluation approaches, findings were summarized in a descriptive manner rather than following a formal statistical comparison method.

3. Medical Background of Gastritis: Symptoms, Causes, Diagnosis and Treatment

Prior to investigating various AI applications, it is important to understand every detail about gastritis. This section provides a brief introduction to gastritis and its symptoms, causes, diagnosis, and treatment.

Gastritis is an inflammatory condition of the stomach mucosa (which is the stomach lining) that poses itself as an important public health concern worldwide. It can be acute or chronic (Khandelwal et al. 93). Gastritis could end up causing significant pain and discomfort, and also impact digestion and quality of life, and in some cases even lead to dangerous forms, including gastric cancer. Notably, among patients who undergo gastroscopy and biopsy, 1 in 85 with gastritis will develop gastric cancer within 20 years (Song et al.). Determining the exact global numbers for gastritis alone is challenging because many studies on larger populations report gastritis alongside duodenitis. In a recent study, the Global Burden of Diseases estimated that there were 27.2 million incident cases of gastritis and duodenitis in 2021. The analysis also indicated that there are currently 38 million cases globally of gastritis and duodenitis (L. Wang et al. 2570). Regional studies further demonstrate the impact of gastritis. A cross-sectional study carried out at Saint Paul Hospital Millennium Medical College in Addis Ababa showed that 78.8% of the studied patients had gastritis, including 48.9% acute gastritis and 29.9% chronic gastritis (Feyisa and Woldeamanuel). Another study was conducted at District General Hospital in Amravati, India, with 150 gastritis patients. It revealed an association between gastritis and factors like lifestyle, diet, medication use, and stress (Khandelwal et al. 93). Furthermore, a Swedish population study found that 40.2% of adults who underwent endoscopy and biopsy had histological gastritis (Zuzek et al. 528).

Gastritis patients may have a range of symptoms, even though some patients may not have noticeable symptoms at first. Common symptoms include upper abdominal pain, nausea, and discomfort after meals. In chronic cases, gastritis could present with gastrointestinal

bleeding, which can appear as vomiting blood or black stools. Symptoms alone are not sufficient to confirm diagnosis since these symptoms may overlap with other diseases as well. Various factors may lead to gastritis, including long-term NSAID use, drinking, smoking, lifestyle factors, psychological factors, and *H. pylori* infection.

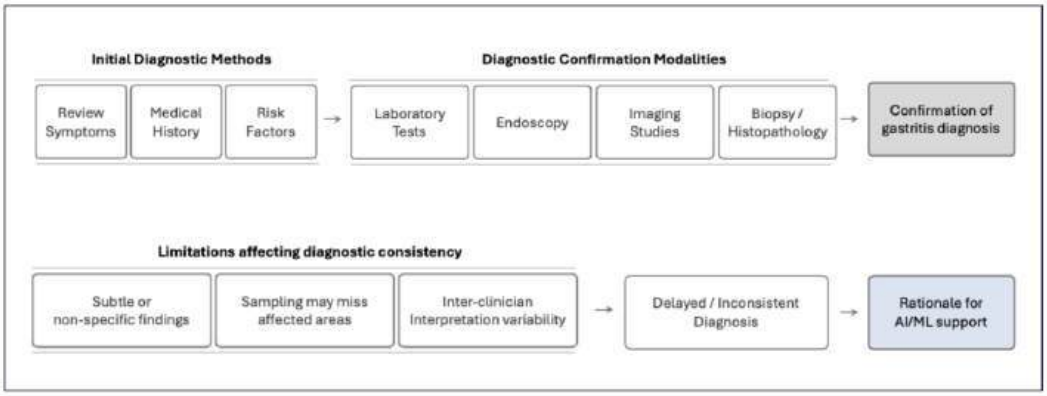

Figure 1. Gastritis Diagnosis Flow & Factors affecting diagnostic consistency

As shown in Figure 1, the diagnosis starts in a physician's office with a review of symptoms, medical history, and possible risk factors. Depending on specific cases, doctors may use lab tests, endoscopy, imaging studies, and biopsy-based histopathological examinations to confirm the diagnosis. Gastritis diagnosis is not easy, even with these tools. Although the findings may be subtle and biopsy samples may miss most affected areas, interpretation can vary between clinicians. Due to these challenges, delayed or inconsistent diagnoses may take place, and this may help explain the growing interest in AI and ML in gastritis care.

The treatment will depend on what causes gastritis. For people infected with H. pylori, doctors will likely prescribe strong antacids and antibiotics to kill the infection. In some cases, simply reducing the irritation to the stomach may work well for people. This could include stopping some medications, limiting alcohol, stopping smoking, or adjusting eating habits. Patients with chronic gastritis may also need ongoing monitoring and lifestyle changes. Accurate diagnosis and assessment are therefore important because treatment may differ considerably between patients based on the underlying cause and severity (Kumar et al. 527).

4. Current Applications of AI in Gastritis

This section reviews how AI is being used in gastritis research and clinical practice. Since gastritis diagnosis often involves symptoms, endoscopy, histopathology, and other clinical data, it generates large amounts of information that may be suitable for AI analysis. Previous studies have applied AI to tasks such as disease grading, severity assessment, H. pylori detection, and prediction of progression to gastric cancer. We will look into these applications, their results, and limitations in the following sections.

4.1. Automatic Detection of Gastritis from Endoscopic Images and Videos

AI has increasingly been used to analyze endoscopic images and videos for signs of gastritis, either during or after endoscopy. The goal is to help identify mucosal changes linked to chronic atrophic gastritis, gastritis-related inflammation, etc. As endoscopy is one of the main diagnostic tools used to diagnose gastritis, image interpretation can vary based on the quality of the image used, the severity of mucosal changes, and the experience of the clinician (Shi et al., "Diagnostic Value"; Zhao et al.; Mu et al. E956). In light of this, AI is viewed as an instrument that can aid physicians and not replace them.

As part of this paper, various studies conducted on AI applications in the automatic detection of gastritis from endoscopic images have been reviewed. To facilitate comparisons, the data used, performance outcomes reported, and the interpretation of each study are presented in Table 2, which compiles the main studies on automated gastritis detection.

Study	Dataset Summary	Sensitivity	Specificity	Accuracy	AUC	Key Interpretation
Shi et al.'s Meta-analysis (Shi et al., "Diagnostic Value")	8 AI-assisted endoscopy studies; 25216 patients; 84678 training images; 10937 test images/videos	94%	96%	N/A	0.98	Overall high performance. But results varied substantially across studies, so the pooled result should be read cautiously.
Shi et al.'s GAM-EfficientNet (Shi et al., "Deep Learning-Assisted Diagnosis")	10961 endoscopic images and 118 video clips from 4050 patients	93%	94%	93.5%	N/A	Performed well on an external test, meaning useful support for chronic atrophic gastritis detection.
Zhao et al.'s U-Net Model (Zhao et al.)	1306 patients; real-time endoscopic diagnosis and OLGA staging	89.31%	90.46%	89.89%	0.919	Showed good diagnostic and OLGA staging performance in a real endoscopy setting.
Mu et al.'s Multicenter Classification Model (Mu et al. E955)	8141 endoscopic images from 4587 cases	N/A	N/A	85.81%–95 .00%	N/A	Accuracy varied across test sets and tasks, showing why external and video-based testing matters.

Note: N/A indicates that the metric was not reported in the study.

Table 2. Comparison of Studies Conducted on AI-Based Gastritis Detection Using Endoscopy

As depicted in Table 2, reported performance varied across studies. Shi et al.'s meta-analysis reported 94% sensitivity, 96% specificity, and an AUC of 0.98, while individual studies reported accuracy ranging from 85.81% to 95.00% (Shi et al., "Diagnostic Value"; Mu et al. E961). In their research, Mu et al. observed differences among testing environments (Mu et al. E961). Zhao et al. examined real-time diagnosis and OLGA staging (Zhao et al.). As evident

from the above table, performance varied across studies and testing environments, likely reflecting the differences in datasets and evaluation methods.

One challenge is that AI models learn from the images and data they are trained on. A model that performs well in one setting may not perform the same way in another. Hospitals use different equipment, serve different patient populations, and encounter different disease patterns. For that reason, reported accuracy should be considered together with other important issues such as dataset diversity, validation, explainability, bias, and clinical implementation (Turtoi et al.; Mehta et al.; Patel et al. 10-11).

4.2. Digital Pathology and Biopsy Slide Analysis

Biopsy samples are used to analyze tissue changes that may not be detected or visible during endoscopy. Traditionally, pathologists test these tissues under microscopes to understand tissue changes. This process may take time and vary among pathologists. Digital pathology converts tissues into high-resolution whole-slide images. AI models can be used to assess these images to detect patterns of inflammation, gastric atrophy, intestinal metaplasia, and H. pylori infection. Pathologists assisted by AI models experienced better diagnostic accuracy and speed (Fang et al. 348; Y. Lin et al. 771; Gonçalves et al.).

Multiple studies of AI applications in digital pathology and biopsy slide analysis were reviewed as part of this paper. To make comparisons easier, Table 3 summarizes the dataset used, results reported in each study, along with interpretations of the results.

Study	Dataset Summary	Sensitivity	Specificity	Accuracy	AUC	Key Interpretation
SFang et al.'s GasMIL Assisted Pathologist Observer Study (Fang et al. 343)	Observer study using 60 random patients from external validation set, comparing pathologists with and without GasMIL assistance.	80%	85%	N/A	0.953 for gastric atrophy diagnosis.	GasMIL assistance improved gastric atrophy diagnosis in pathologist study
Yi-Jyun Lin et al.'s Two-Tier Deep Learning Model (Y. Lin et al. 771)	885 Whole-Slide Images for histologic detection of H. pylori-associated gastritis.	93.3%	90.1%	N/A	0.9739	Strong performance in detecting H. pylori-associated gastritis from pathology slides.
Ma et al.'s AMMNet (Ma et al. 1538)	1096 patients from seven medical centers; pathology-based multilabel assessment.	N/A	N/A	N/A	0.93 for inflammatory activity. 0.97 for atrophy. 0.93 for intestinal metaplasia.	Handled several pathology indicators, including inflammation, atrophy and intestinal metaplasia.

Note: N/A indicates that the metric was not reported in the study.

Table 3. Comparison of Studies Conducted on AI-Based Gastritis Detection Using Digital Pathology Studies

The studies summarized in Table 3 reported AUC values ranging from 0.93 to 0.9739 across the evaluated pathology tasks; Yi-Jyun Lin et al. additionally reported 93.3% sensitivity and 90.1% specificity for detecting H. pylori-associated gastritis. (Yi-Jyun Lin et al.). These findings suggest that AI may help pathologists evaluate biopsy samples more consistently and efficiently.

The studies reviewed here were different from one another. One was focused on the diagnosis of gastric atrophy, while the other two were focused on H. pylori-associated gastritis and changes. Therefore, direct comparison is not easy. AI will be a useful tool in digital pathology if it continues to perform strongly in more extensive and diverse patient populations.

There is still a challenge, as hospitals do not prepare the biopsy slides in the same way. Whole-Slide Images (WSIs) can look different due to the methods used for tissue preparation, staining, and scanning. Variation in tissue preparation, staining, and slide scanning across different hospitals may affect the consistency of digital-pathology models.

4.3. Grading, Staging and Severity Assessment

The aim is not only to detect the presence of gastritis but also to grade, stage, and assess its severity. Doctors must also gauge how severe the underlying changes are. Gastric atrophy and intestinal metaplasia are good examples because the severity can influence treatment options, follow-up, and cancer risk monitoring. As a result, disease staging systems such as OLGA can be useful to determine the severity of disease and plan follow-up. AI could assist in making more consistent assessments and provide physicians with additional assistance as they plan follow-up care (Fang et al. 350-351; Zhao et al.).

Study	Dataset Summary	Sensitivity	Specificity	Accuracy	AUC	Key Interpretation
Zhao et al.'s U-Net Model (Zhao et al.)	1306 patients. chronic atrophic gastritis diagnosis and OLGA staging.	89.31%	90.46%	89.89%	0.919	Good Performance for chronic atrophic gastritis diagnosis and OLGA staging, with results reported above endoscopists on tested diagnostic indexes.
Fang et al.'s GasMIL (Fang et al. 343)	2725 Whole-Slide Images from 545 patients across 13 tertiary hospitals in China. Gastric atrophy and intestinal metaplasia grading.	N/A	N/A	N/A	0.877 for gastric atrophy. 0.884 for intestinal metaplasia.	Graded atrophy and intestinal metaplasia with strong AUC values.
Tao et al.'s White-Light Endoscopy System (Tao et al. 1319)	White-light endoscopy-based gastric atrophy risks stratification across non-GA, mild GA, moderate GA, and severe GA categories.	N/A	N/A	Overall accuracy: 81.37%. Category-specific accuracy: 77.42%–85. 71%	N/A	Helped group patients by gastric atrophy risk, but accuracy differed between severity levels.

Note: N/A indicates that the metric was not reported in the study.

Table 4. Comparison of AI Studies on Grading, Staging, and Severity Assessment of Gastritis

Table 4. Comparison of AI Studies on Grading, Staging, and Severity Assessment of Gastritis Table 4 summarizes the dataset used and key results reported, along with the interpretations of the studies that were reviewed in AI applications on grading, staging, and severity assessment. These studies show that AI can be helpful in this area. It is important because severity can drive doctors' decisions on monitoring and treatment plans. Comparing studies with one another was difficult, as they targeted different outcomes and used different datasets and evaluation methods.

One thing that became clear from these studies was how much the results depended on the data used to develop, train, and test models. Differences in labeling, patient populations, clinical practices, and targeted outcomes may impact model performance and make direct comparison across severity assessment studies difficult (Fang et al. 350; Mehta et al.; Turtoi et al.).

4.4. Detection of Helicobacter pylori Gastritis

H. pylori-associated gastritis is one of the most common forms of gastritis. As studies were reviewed, it became clear that detecting H. pylori is not the main challenge. Regular tests like blood tests, stool tests, and urea breath tests can confirm it, but they don't reveal changes taking place within the stomach tissue. Hence, researchers started exploring AI models that analyze pathology slides and endoscopic images to detect inflammation and damage. This could help doctors better understand how infection is affecting the stomach walls and plan treatment (Y. Lin et al. 771; Shi et al., "Deep Learning-Assisted Diagnosis"; Zhao et al.).

Study	Dataset Summary	Sensitivity	Specificity	Accuracy	Accuracy / Other Metrics	AUC	Key Interpretation
Yi-Jyun Lin et al.'s Two-Tier Deep Learning Model (Y. Lin et al. 771)	885 Whole-Slide Images (WSIs) for histologic detection of H. pylori-associated gastritis.	Pathology-based detection of H. pylori-associated gastritis.	93.3%	90.1%	Same sensitivity as pathologist; higher specificity than pathologists.	0.9739	Matched pathologist sensitivity and improved specificity for H. pylori-associated gastritis detection.
Yingying Wang et al.'s XGBoostBased Antibiotic Resistance Model (Y. Wang et al.)	Genomic data from 296 H. pyloriisolates	Genomic prediction of H. pyloriantibiotic resistance to support treatment decisions.	N/A	N/A	For selected resistance targets: accuracy 0.8500, recall 0.8500, precision 0.8474, F1-score 0.8462.	0.8466 for metronidazole resistance; 0.9320 for clarithromycin resistance.	Supports antibiotics resistance prediction for treatment planning, not direct gastritis detection.

Note: N/A indicates that the metric was not reported in the study. The two studies address different objectives, so they should not be interpreted as directly comparable.

Table 5. AI Applications in H. pylori Gastritis

The studies listed in Table 5 show various applications of AI in H. pylori gastritis. These studies show that AI could be helpful in the management of H. pylori-associated gastritis. Yi-Jyun Lin et al. focused on detecting H. pylori-associated gastritis from pathology slides (Y. Lin et al. 771). Yingying Wang et al. used genomic data to predict antibiotic resistance (Y. Wang et al.). Although the studies were on different clinical problems, both explored the ways in which AI could support decision-making in H. pylori-related gastritis.

Detection of H. pylori-associated gastritis from biopsy samples is not easy. The bacteria can be small and unevenly distributed within biopsy samples, which makes detection tougher even for experienced pathologists. Differences in biopsy slide preparation, staining, and imaging methods across different hospitals, and uneven distribution of bacteria may affect model performance across different datasets (Mehta et al.; Turtoi et al.).

4.5. Predicting Progression from Atrophic Gastritis to Gastric Cancer

Predicting progression from atrophic gastritis to gastric cancer focuses on identifying patients who are more likely to develop gastric cancer in the future. Not all patients with atrophic gastritis or other precancerous gastric changes will develop gastric cancer, but some face a much higher risk than others. Identifying such cases early can help doctors in monitoring and follow-up decisions. AI could analyze complex data elements from biopsy results, endoscopy findings, and medical history together to identify patterns associated with disease progression (Kim et al.; Fang et al. 343; Mehta et al.).

Very few papers have focused on this area, and they approached the problem in different ways. Kim et al. developed DeepPrevention to estimate whether patients with atrophic gastritis might later progress to gastric cancer and reported an AUC of 0.84 (Kim et al.). Fang et al. took a different approach, focusing on precancerous gastric changes and reporting AUC values of 0.877 for gastric atrophy and 0.884 for intestinal metaplasia (Fang et al. 343). Zhao et al. tried to address chronic atrophic diagnosis and OLGA staging, achieving an AUC of 0.919 and an accuracy of 89.89% (Zhao et al.). However, Kim et al.'s was the only study that directly addressed gastric cancer progression, while Fang et al.’s and Zhao et al.’s studies assessed the current disease stage. Based on these studies, it can be inferred that AI can be used to identify important gastric changes and help doctors assess future disease risk earlier and consistently.

One thing that became clear while analyzing these studies is that reliable prediction depends heavily on the data available for model development and training. Predicting cancer progression is difficult when datasets are small or when they are from a small patient population. The limited number of progression studies and the corresponding narrow datasets restrict generalizability and strength of conclusions about future cancer progression risk. Accordingly, these systems are useful when used along with clinical judgment (Kim et al.; Mehta et al.).

4.6. Summary of Findings

After reviewing these studies, it’s clear that AI is being explored in several different areas of gastritis care. The applications ranged from diagnosis and severity assessment to H.

pylori-related areas and prediction of progression towards gastric cancer. Across the studies summarized in Tables 2–5, reported performance included sensitivity up to 94%, specificity up to 96%, accuracy up to 95%, and AUC values up to 0.98, although the studies evaluated different clinical tasks and are not directly comparable. Reported results show that AI may be able to support tasks that usually require extensive time, expertise, or consistent interpretation across pathologists and doctors. In some cases, technology was being used in risk assessment and treatment-related decisions that are beyond just diagnosis.

Despite the reported performance summarized in Tables 2-5, clinical readiness cannot be concluded from performance alone. Section 5 discusses broader issues affecting clinical use in greater detail.

5. Explainability, Validation, Dataset Quality & Real-world Readiness

Building an accurate model is one thing. Knowing whether it can be trusted in routine clinical care is another. Although many AI systems perform encouragingly during development, questions about explainability, validation, dataset quality, and real-world performance remain. These issues play a major role in determining whether a model is ready for routine use in gastritis care and are examined in the following sections.

5.1. Explainability and Clinical Decision Support

Doctors do not make decisions based on predictions from an AI model alone. They need to understand how those predictions were reached, especially in healthcare where decisions can affect diagnosis, treatment, and patient outcomes. If an AI system identifies gastritis, precancerous gastric changes, or an increased cancer risk, doctors need to know which findings influenced that. Otherwise, it becomes difficult to judge whether it should be trusted. This is where explainability becomes important. The clearer a model's reasoning is, the more likely doctors feel comfortable using it in daily practice.

Explainability is considered a critical parameter in the reviewed studies. Model explainability improves confidence levels of doctors and clinicians in using AI models. Ma et al. developed an interpretable multilabel deep learning model capable of assessing multiple gastritis-related indicators simultaneously (Ma et al. 1538). Kim et al. proposed an explainable deep learning model to predict progression from atrophic gastritis to gastric cancer (Kim et al.). Yingying Wang et al. also incorporated interpretability into a model for predicting H. pylori antibiotic resistance (Y. Wang et al.).

These studies suggest that explainability is very useful across different AI applications. These studies suggest that the role of AI in gastritis is not to replace doctors, but to support them in decision-making. Explainable AI can help doctors by providing the reasons why AI models came to their decisions, which would highlight important findings and present information more consistently. However, the final interpretation and decision still depend on clinical judgment, medical history, and expert evaluation. For that reason, explainable AI in gastritis should be used to strengthen clinical decision-making instead of trying to replace it.

5.2. Dataset Quality and Diversity

Irrespective of the field in which AI models are implemented, when it comes to the performance of AI models, it follows a simple principle: “garbage in, garbage out”. The quality of the output generated by AI models heavily depends on the data used to train them, since AI systems learn by identifying patterns within training data. If it’s trained with high-quality and diverse datasets, the outputs generated and its performance would be very promising; however, if low-quality datasets with poor-quality images, incorrect labels, small samples, and missing information are used, the model's performance is likely to suffer irrespective of how advanced the technology itself is. Especially when AI systems are deployed in healthcare, specifically for gastritis, the datasets used to train the models must be high quality and diverse, as these AI models would have to identify changes that are subtle, inconsistent, and difficult to locate. Additionally, the data for gastritis consists of a lot of medical images such as endoscopic and pathology images, which would be very difficult to interpret if they are of low quality. As a result, the quality of the data becomes just as important as the design of the AI model itself.

Some AI models may exhibit excellent performance when used to diagnose, detect, or assess the severity of gastritis at certain hospitals or populations. However, when they are used outside those hospitals or populations, their performance declines. One possible reason for this decline in performance is that the data used for training may not fully represent the diversity of patients, devices, and hospital settings. Single-hospital training data is very useful when developing AI models that are made specifically for that hospital. However, they may not perform well at other hospitals or clinics when exposed to different datasets consisting of different patients, devices, or imaging styles. Mu et al. and Ne Lin et al. tested their models using datasets from multiple hospitals, which allowed them to assess the models’ performance in broader clinical settings, patients, and image/device differences (Mu et al. E956-E957; N. Lin et al.). The models trained/tested on multi-hospital datasets turned out to be more efficient, as these models can be used across different hospitals. Furthermore, models that continue to perform well when tested on separate datasets are generally more likely to be used in real clinical settings. However, we must keep in mind that these AI models should be validated carefully before implementing them in clinical care and may require further training as new data becomes available.

5.3. Validation and Generalizability

One thing that became clear across the studies is that validation plays a major role in determining whether an AI model is ready for clinical use. A model may perform well during development but produce extremely different results when tested on new patients or unfamiliar data. Because of this, researchers usually look beyond training performance and evaluate how models perform on different data. Models that continue to perform well across different hospitals, patient populations, or geographies are more likely to remain reliable outside the setting in which they were developed and tested. Validation is especially important in gastritis research because data can come from different hospitals, patient populations, and clinical

settings. Endoscopic findings, biopsy preparation, image quality, and labeling practices may vary considerably between hospitals. A model that performs well in one hospital should not automatically be assumed to perform well in another. Consistent performance across different patient groups and healthcare environments is ultimately what determines whether a model is ready for routine healthcare use.

Multiple reviewed studies have used external validation to prove models' performance and accuracy. Shi et al. developed a GAMEfficientNet deep-learning model for the diagnosis of chronic atrophic gastritis in endoscopy. Even though an internal image dataset from Zibo Central Hospital was used for the training and internal testing of the model, an additional external image test set from the Zhangdian Maternal and Child Health Care Hospital was also used for testing. The sensitivity, specificity, precision, AUC, and accuracy of the model were 93%, 94%, 93.94%, 99.46%, and 93.5%, respectively, on the external test set (Shi et al., "Deep Learning-Assisted Diagnosis"). Similarly, Cao et al. developed and validated a CAG risk prediction model to accurately identify patients who are at an elevated risk for chronic atrophic gastritis, in which the model achieved AUCs of 0.8542 and 0.8073 in the training and testing sets, respectively. Additionally, the model achieved an AUC of 0.8505 in the external validation, which shows good generalizability (Cao et al.). Such studies that exhibited their models' promising performance by externally validating them demonstrate their potential clinical utility and the importance of external validation (Shi et al., "Deep Learning-Assisted Diagnosis"; Cao et al.).

5.4. Real World Readiness and Recurring Limitations

The next step, which is crucial to the successful use of AI in routine clinical care, is real-world testing. Studies that have undergone real-world clinical trials tend to be stronger than laboratory or controlled testing, as they demonstrate how AI systems perform under conditions that more closely reflect everyday clinical practice. A real-time video-monitoring model was developed by Zhao et al. for endoscopic diagnosis of chronic atrophic gastritis and OLGA staging based on the U-Net deep learning architecture. Furthermore, Zhao et al. tested the model in a real-world setting with 1,306 patients, where their model achieved sensitivity, specificity, accuracy, and AUC scores of 89.31%, 90.46%, 89.89%, and 0.919, respectively (Zhao et al.). Another AI-assisted endoscopy system was developed by Zou et al. to diagnose H. pylori infection using 7,377 endoscopic images from 639 patients, achieving accuracy, sensitivity, and specificity scores of 89.6%, 90.9%, and 88.9%, respectively (Zou et al.). Both studies provide examples of AI Systems evaluated under real-world testing conditions. Such studies show the path for implementing AI systems in daily clinical care (Zhao et al.; Zou et al.).

Despite the performance reported in these studies, several limitations remain. Some of these limitations include limited training data, narrow patient populations, and lack of real-world testing. Such limitations affect the performance of these systems/models in the real world (Shi et al., "Diagnostic Value"; Mu et al. E961; Zhao et al.). As a result, a model that performs well during development may not always perform equally well when used in routine healthcare settings.

Some of the studies reviewed in this paper show a clear move of AI models from testing in the research stage to real hospital environments. This matters because strong results in the testing phase do not always get carried on to routine healthcare practice. Challenges remain, but testing these systems in real hospital settings helps researchers better understand where they work well and where improvements are still needed.

5.5. Summary of Findings

What stood out from these studies was that building an accurate AI model is only one part of the process. Whether a model can be explained, validated, and used reliably across different healthcare settings is just as important as its reported performance. The studies summarized in Section 4 reported measurable performance across different AI applications, but there is still uncertainty about how these systems will perform in real healthcare practice. The discussion now shifts to stress-associated gastritis and the gut-brain axis, an area that developed largely outside the literature reviewed so far.

6. Stress, Gut-Brain Axis and Future Opportunities

After reviewing the studies discussed in this paper, one thing became very clear. AI is already being used in several areas of gastritis care, from diagnosis and disease classification to severity assessment and cancer risk prediction. The review also showed that good performance alone is not enough. Explainability, dataset quality, validation, and real-world testing still matter. Very little research, however, appears to have explored how AI could be applied specifically to stress-associated gastritis. That was interesting because many studies already point to a connection between psychological stress, the gut-brain axis, and gastric health. This section looks at those findings more closely and considers how AI might be applied in stress-associated gastritis.

6.1. The Gut-Brain Axis and Gastric Health

The gut-brain axis (GBA) is a two-way communication system that connects the gastrointestinal tract and the central nervous system, through which they keep exchanging signals even though the digestive system and the brain are different parts of the body. The brain can influence gastrointestinal function, and changes in the gut can also affect the brain. This communication between the two parts happens through several interconnected biological pathways, which include neural and microbiome-related mechanisms (Leigh et al. 4492; Konturek et al. 592). Because the stomach is a part of this system, psychological stress may influence gastric health and could contribute to gastritis.

Psychological stress may affect gastritis through several pathways. Stress can influence gastrointestinal secretion, mucosal blood flow and repair, gut motility, intestinal permeability, microbiota composition, and immune activity(Konturek et al. 591; Leigh et al. 4520). Gyires and Feher also reported that stress may contribute to gastric mucosal injury by affecting circulation within the stomach lining and other stress response processes (Gyires and Feher 3928). These

findings do not mean that stress is the only cause of gastritis. Instead, stress appears to be one possible contributor.

6.2. Evidence Linking Stress and Gastritis

Evidence linking stress and gastritis has started to emerge from several areas of research. Feyisa and Woldeamanuel identified stress as one of the major factors associated with chronic gastritis in their study population, with an adjusted odds ratio of 2.168 (Feyisa and Woldeamanuel). In a separate study involving more than 168,000 full-time workers in South Korea, employees working more than 55 hours per week had a significantly higher prevalence of endoscopic gastritis than those working 35 to 40 hours per week. These findings suggest that stress and work-related pressures may influence gastric health (Mun et al.).

Research in this area has found connections between gastritis and factors such as anxiety, depression, sleeping habits, lifestyle, and psychological wellbeing (Saba et al. 2413; Wu et al. 3147). Some studies have also examined whether these factors are related to chronic atrophic gastritis and progression towards early gastric cancer (Wu et al. 3148; Liang et al.). This does not mean that stress directly causes gastritis, but it does indicate that stress may have a role in gastric health. Other factors, including H. pylori infection, medication use, lifestyle, and underlying medical conditions, remain important (Feyisa and Woldeamanuel). In this paper, stress-associated gastritis refers to gastritis that may be influenced or worsened by psychological stress through a combination of biological responses, lifestyle factors, and gut-brain interactions.

6.3. A Future Opportunity for AI Research

Research on AI in gastritis and research on stress-related gastrointestinal health have largely been studied separately. Current AI studies have concentrated on disease detection, classification, severity assessment, H. pylori-related applications, and gastric cancer risk prediction (Turtoi et al.). In contrast, stress-related research has focused on how factors such as stress, lifestyle, and gut-brain interactions may impact digestive health (Leigh et al. 4492; Ballou and Feingold 697; Konturek et al. 592). Very few studies have examined whether AI could be used to better understand or predict stress-associated gastritis.

One question that comes out of this gap is whether AI might uncover patterns that are difficult to spot when medical findings are considered on their own. This is not about proving that stress causes gastritis. Instead, researchers could examine how stress, sleeping habits, lifestyle, symptoms, and gastritis interact in real patients. This may help identify individuals who are at higher risk and could benefit from closer monitoring or further evaluation.

6.4. Data, Validation and Ethical Requirements

Studying stress-associated gastritis with AI would require more medical data than the other gastritis studies. Researchers would still need confirmed diagnoses, endoscopy findings, biopsy results, and H. pylori status. Additional information about stress, sleep patterns, symptoms, lifestyle, and medication use would also be needed. Since stress and gastric health

can change over time, data from one medical visit will not be enough. This adds additional complexity as well. Another challenge would be determining whether symptoms are caused by gastritis itself or by stress-related digestive problems. Information from different hospitals, geographies, and patient groups would also need to be used in the study. Otherwise, the findings might reflect only one healthcare setting or population (Mehta et al.; Wu et al. 3157; Liang et al.).

There are also important ethical questions that need to be addressed. Research on stress-associated gastritis involves information about physical health, psychological health, sleeping habits, lifestyle, and medical history. Since much of this information is personal, protecting patient privacy is important. Appropriate IRB (Institutional Review Board) approval and, where required, patient consent will be needed before any data can be collected or analyzed. Patient information needs to be stored securely, handled responsibly, and evaluated carefully across patient populations. Another concern is that a model developed using data from a single group of patients may not perform the same in other settings. To avoid this, future studies need information from different hospitals, regions, and patient groups. For these reasons, privacy, fairness, transparency, external validation, and clinician oversight should be considered from the very beginning of any future research effort (Mehta et al.).

6.5. Potential Value and Summary

If future studies can overcome these data, validation, and ethical challenges, AI may provide a new way to study how stress and gastritis interact over time. Looking at confirmed gastritis alongside stress levels, sleep patterns, lifestyle habits, and symptom changes may reveal patterns that are easy to miss when each factor is examined separately. Over time, this could help researchers build models that predict the likelihood of stress-associated gastritis using a combination of medical, psychological, and lifestyle-related information. Developing a model like this could be useful for the earlier detection of patients who could benefit from additional support, including closer monitoring, stress-management support, lifestyle changes, or further clinical evaluation.

At this stage, stress-associated gastritis is best understood as a possible future area of research rather than a current AI application. The findings reviewed in this paper suggest that the idea is worth exploring, but there would be a lot of challenges to be tackled by the researchers because of the type and amount of data needed and potential ethical approvals. More evidence, better data, and careful validation would be needed before any AI system could be developed or tested in routine clinical care. Even so, studying stress, gastric health, and AI together may help researchers better understand how psychological factors influence digestive diseases.

7. Conclusion

This review examined the applications of AI and ML in the diagnosis and management of gastritis. The studies showed that AI can already assist with disease detection, severity assessment, and gastric cancer risk prediction. This review examined applications of AI and ML

in the diagnosis and management of gastritis. Across the studies summarized in Tables 2–5, reported performance included sensitivity up to 94%, specificity up to 96%, accuracy up to 95%, and AUC values up to 0.98, although the studies addressed different clinical tasks and were not directly comparable. These findings also show that performance alone is insufficient to determine readiness for everyday clinical care.

Currently, AI can be considered useful when it supports doctors in decision-making instead of replacing their judgment. For that to happen, doctors need to understand how the model arrived at its decision. The information used to train the model must also be accurate, diverse, and from different populations and geographies. Models should be tested in different hospitals and patient groups, not only in the setting where they were first developed. These factors determine if an AI model can work reliably outside a research setting.

The review also found a gap between AI research on gastritis and research on stress and gastric health. Existing studies have linked psychological stress, lifestyle, and gut-brain interactions to gastric health (Leigh et al. 4492, 4519; Ballou and Feingold 697; Konturek et al. 592, 594; Mun et al.). But very few have examined how AI could be applied to stress-associated gastritis.

Future studies will need information from patients with confirmed gastritis, along with stress, sleep, lifestyle, and other health information. Since this information is very personal, it would require privacy measures and ethical approval. If researchers can overcome these issues, AI can analyze these psychological, lifestyle, and medical factors together and make quicker decisions. This would end up providing a better understanding of stress-associated gastritis and, at the same time, help deliver personalized care to certain patients (Mehta et al.; Wu et al. 3147, 3157; Liang et al.). However, applying AI to various aspects of stress-associated gastritis isn't an established application of AI yet. It is a promising area for future research that builds on the progress already made in AI-assisted gastritis care and the growing evidence linking stress to gastric health.

8. Acknowledgments

I am very grateful to Dr. Siddharth Krishnan (UNC Charlotte) and Dr. Samuel Lefcourt (Johns Hopkins University) for the time they invested in this project. Their guidance, feedback, and willingness to answer my questions helped me improve both the paper and my understanding of AI, gastritis, and research itself as a high schooler. Every discussion gave me something new to think about and pushed me to look at the topic from a different perspective.

I am grateful for their patience throughout this project from day one to completion. Their comments challenged me to think more carefully about the evidence, question my assumptions, and strengthen my arguments. Working under their guidance has been one of the most valuable parts of this experience, and this paper would not be what it is today without their support.

I would also like to thank my parents for supporting me throughout this project. They encouraged me during long periods of reading, writing, and revising, and they always reminded me to keep going whenever it felt overwhelming and helped me complete it successfully.

Works Cited

Ballou, Sarah, and Jordyn H. Feingold. “Stress, Resilience, and the Brain-Gut Axis: Why Is Psychogastroenterology Important for All Digestive Disorders?” Gastroenterology Clinics of North America, vol. 51, no. 4, Dec. 2022, pp. 697–709. https://doi.org/10.1016/j.gtc.2022.07.001.

Cao, Hui, et al. “Risk Prediction for Chronic Atrophic Gastritis Using a Random Forest Model: A Multicenter Study.” World Journal of Gastrointestinal Oncology, vol. 18, no. 2, Feb. 2026, p. 113959. https://doi.org/10.4251/wjgo.v18.i2.113959.

Fang, Shuangshuang, et al. “Diagnosing and Grading Gastric Atrophy and Intestinal Metaplasia Using Semi-Supervised Deep Learning on Pathological Images: Development and Validation Study.” Gastric Cancer, vol. 27, no. 2, Mar. 2024, pp. 343–54. https://doi.org/10.1007/s10120-023-01451-9.

Feyisa, Zelalem Tadese, and Berhanu Teshome Woldeamanuel. “Prevalence and Associated Risk Factors of Gastritis among Patients Visiting Saint Paul Hospital Millennium Medical College, Addis Ababa, Ethiopia.” PLOS ONE, vol. 16, no. 2, Feb. 2021, p. e0246619. https://doi.org/10.1371/journal.pone.0246619.

Gonçalves, Wanderson Gonçalves E., et al. “Deep Learning in Gastric Tissue Diseases: A Systematic Review.” BMJ Open Gastroenterology, vol. 7, no. 1, Mar. 2020, p. e000371. https://doi.org/10.1136/bmjgast-2019-000371.

Gyires, Klara, and Agnes Feher. “Stress, Neuropeptides and Gastric Mucosa.” Current Pharmaceutical Design, vol. 23, no. 27, 2017, pp. 3928–40. https://doi.org/10.2174/1381612823666161118144216.

Khandelwal, V., et al. “Epidemiology and Associated Risk Factors of Gastritis in Patients at District General Hospital, Amravati.” International Journal of Health Sciences and Research, vol. 14, no. 2, Feb. 2024, pp. 93–100. https://doi.org/10.52403/ijhsr.20240212.

Kim, Hyon, et al. “A Deep Recurrent Neural Network-Based Explainable Prediction Model for Progression from Atrophic Gastritis to Gastric Cancer.” Applied Sciences, vol. 11, no. 13, July 2021, p. 6194. https://doi.org/10.3390/app11136194.

Konturek, Peter C., et al. “Stress and the Gut: Pathophysiology, Clinical Consequences, Diagnostic Approach and Treatment Options.” Journal of Physiology and Pharmacology: An Official Journal of the Polish Physiological Society, vol. 62, no. 6, Dec. 2011, pp. 591–99.

Kumar, Shria, et al. “Risk Factors and Incidence of Gastric Cancer After Detection of Helicobacter pylori Infection: A Large Cohort Study.” Gastroenterology, vol. 158, no. 3, Feb. 2020, pp. 527–536.e7. https://doi.org/10.1053/j.gastro.2019.10.019.

Leigh, Sarah‐Jane, et al. “The Impact of Acute and Chronic Stress on Gastrointestinal Physiology and Function: A Microbiota–Gut–Brain Axis Perspective.” The Journal of Physiology, vol. 601, no. 20, Oct. 2023, pp. 4491–538. https://doi.org/10.1113/JP281951.

Liang, Mengmeng, et al. “The Association between Psychological Status and the Development of Early Gastric Cancer from Atrophic Gastritis.” Medicine, vol. 104, no. 45, Nov. 2025, p. e45653. https://doi.org/10.1097/MD.0000000000045653.

Lin, Ne, et al. “Simultaneous Recognition of Atrophic Gastritis and Intestinal Metaplasia on White Light Endoscopic Images Based on Convolutional Neural Networks: A Multicenter Study.” Clinical and Translational Gastroenterology, vol. 12, no. 8, Aug. 2021, p. e00385. https://doi.org/10.14309/ctg.0000000000000385.

Lin, Yi‐Jyun, et al. “Two‐tiered Deep‐learning‐based Model for Histologic Diagnosis of Helicobacter Gastritis.” Histopathology, vol. 83, no. 5, Nov. 2023, pp. 771–81. https://doi.org/10.1111/his.15018.

Ma, Mengke, et al. “Advancing Automatic Gastritis Diagnosis: An Interpretable Multilabel Deep Learning Framework for the Simultaneous Assessment of Multiple Indicators.” The American Journal of Pathology, vol. 194, no. 8, Aug. 2024, pp. 1538–49. https://doi.org/10.1016/j.ajpath.2024.04.007.

Mehta, Yagna, et al. “Gastroenterology in the Age of Artificial Intelligence: Bridging Technology and Clinical Practice.” World Journal of Gastroenterology, vol. 31, no. 36, Sept. 2025, p. 110549. https://doi.org/10.3748/wjg.v31.i36.110549.

Mu, Ganggang, et al. “Expert-Level Classification of Gastritis by Endoscopy Using Deep Learning: A Multicenter Diagnostic Trial.” Endoscopy International Open, vol. 9, no. 6, June 2021, pp. E955–64. https://doi.org/10.1055/a-1372-2789.

Mun, Eunchan, et al. “Cross-Sectional Association between Long Working Hours and Endoscopic Gastritis: The Kangbuk Samsung Health Study.” BMJ Open, vol. 11, no. 9, Sept. 2021, p. e050037. https://doi.org/10.1136/bmjopen-2021-050037.

Patel, Vatsal, et al. “Artificial Intelligence Applied to Gastrointestinal Diagnostics: A Review.” Journal of Pediatric Gastroenterology and Nutrition, vol. 70, no. 1, Jan. 2020, pp. 4–11. https://doi.org/10.1097/MPG.0000000000002507.

Saba, Malka I., et al. “The Impact of Psychological Stress on Gastrointestinal Health: A Focus on Gastritis.” Journal of Population Therapeutics and Clinical Pharmacology, vol. 31, no. 11, Dec. 2024, pp. 2413–18. https://doi.org/10.53555/45sncb93.

Shi, Yanting, et al. “Deep Learning-Assisted Diagnosis of Chronic Atrophic Gastritis in Endoscopy.” Frontiers in Oncology, vol. 13, Mar. 2023, p. 1122247. https://doi.org/10.3389/fonc.2023.1122247.

Shi, Yanting, et al. “Diagnostic Value of Artificial Intelligence-Assisted Endoscopy for Chronic Atrophic Gastritis: A Systematic Review and Meta-Analysis.” Frontiers in Medicine, vol. 10, May 2023, p. 1134980. https://doi.org/10.3389/fmed.2023.1134980.

Song, Huan, et al. “Incidence of Gastric Cancer among Patients with Gastric Precancerous Lesions: Observational Cohort Study in a Low Risk Western Population.” BMJ, vol. 351, July 2015, p. h3867. https://doi.org/10.1136/bmj.h3867.

Tao, Xiao, et al. “An Artificial Intelligence System for Chronic Atrophic Gastritis Diagnosis and Risk Stratification under White Light Endoscopy.” Digestive and Liver Disease, vol. 56, no. 8, Aug. 2024, pp. 1319–26. https://doi.org/10.1016/j.dld.2024.01.177.

Turtoi, Daria Claudia, et al. “Artificial Intelligence for the Automatic Diagnosis of Gastritis: A Systematic Review.” Journal of Clinical Medicine, vol. 13, no. 16, Aug. 2024, p. 4818. https://doi.org/10.3390/jcm13164818.

Wang, Lufei, et al. “Global, Regional, and National Burden of Gastritis and Duodenitis from 1990 to 2021 with Projections to 2050: A Systematic Analysis of the Global Burden of Disease Study 2021.” International Journal of Medical Sciences, vol. 22, no. 11, May 2025, pp. 2570–82. https://doi.org/10.7150/ijms.109762.

Wang, Yingying, et al. “Assessment for Antibiotic Resistance in Helicobacter pylori: A Practical and Interpretable Machine Learning Model Based on Genome-Wide Genetic Variation.” Virulence, vol. 16, no. 1, Dec. 2025, p. 2481503. https://doi.org/10.1080/21505594.2025.2481503.

Wu, Hongzhen, et al. “The Status and Risk Factors for Anxiety/Depression in Patients with Atrophic Chronic Gastritis: A Cross-Sectional Study.” Annals of Palliative Medicine, vol. 11, no. 10, Oct. 2022, pp. 3147–59. https://doi.org/10.21037/apm-22-730.

Zhao, Quchuan, et al. “U-Net Deep Learning Model for Endoscopic Diagnosis of Chronic Atrophic Gastritis and Operative Link for Gastritis Assessment Staging: A Prospective Nested Case-Control Study.” Therapeutic Advances in Gastroenterology, vol. 16, 2023, p. 17562848231208669. https://doi.org/10.1177/17562848231208669.

Zou, Pei-Ying, et al. “Development and Application of an Artificial Intelligence-Assisted Endoscopy System for Diagnosis of Helicobacter pylori Infection: A Multicenter Randomized Controlled Study.” BMC Gastroenterology, vol. 24, Sept. 2024, p. 335. https://doi.org/10.1186/s12876-024-03389-3.

Zuzek, Rachael, et al. “Prevalence of Histological Gastritis in a Community Population and Association with Epigastric Pain.” Digestive Diseases and Sciences, vol. 69, no. 2, Feb. 2024, pp. 528–37. https://doi.org/10.1007/s10620-023-08170-2.

The Influence of Algorithmically Curated Social Media Content on Consumer Purchasing Behaviour Among Generation Z and Millennials: A Literature Review Paper

By Yashita Jain

Abstract

Social media now dictates a large amount of what is purchased, or at least what can be purchased due to a person's browsing data as the platforms rely on recommendation algorithms and not static ad placements. The recommendation systems take a person's browsing activity, purchases and other online actions into consideration. What the literature does not dispute is that personalisation alone is not responsible for a purchase; trust, electronic word of mouth, user generated content, social proof and recommendations from influencers also contribute to a recommendation being taken up by a consumer. This review collates twelve papers, eight literature reviews and four peer reviewed works to assess how algorithmically driven recommendations affect purchasing behaviour for Millennials and Generation Z in particular. Both generations are susceptible to recommendation algorithms but to different extents as the former views algorithmic feeds as a discovery tool and spends more time on product reviews and comparison than the latter which prefers video, word of mouth, user generated content and influencers. The design of the platform and technological advancement plays a part in how these generations interact with the technology. Social influence runs concurrently to recommendation algorithms across the reviewed works with one paper suggesting that purchasing behaviour informed by recommendation algorithms cannot be understood without taking into account technological developments and how they interact with social influence. This review goes one step further and argues that algorithmic curation and social proof, which the literature treats as separate mechanisms, are better understood as one, because the recommendation algorithm decides which reviews, influencer videos and comments a consumer sees in the first place.

Keywords Recommendation algorithms, Consumer purchasing behaviour, Social media marketing, Generation Z, Millennials.

Introduction

As social media keeps on growing more widespread and accessible, it also becomes easier to buy and sell on it. Instagram and TikTok, for example, use algorithms to recommend products. These are based on the browsing and purchasing history of users rather than the basic advertisements shown to everyone using the site. Online forums and peer interactions are also a source of reviews and feedback which shapes a product's image and brand reputation (Fadilah et al., 2025; Mishra, 2024; Zhao et al., 2025). Millennials and Generation Z are both active social media users, but they do it in different ways. Their digital habits, content preferences and levels of trust differ a lot (Paliwal et al., 2025; Al-Amin & Scorita, 2025; Agrawal, 2022).

The question which this review tackles is deceptively simple: how does algorithmically curated social media content influence purchasing behaviour of Millennials and Generation Z,

and how do social proof and generational differences shape that relationship? The role of recommendation algorithms on social commerce cannot be underestimated, as they help find products which suits a user's needs (Fadilah et al., 2025; Mishra, 2024). However, the majority of the reviewed articles consistently conclude on the same point from different angles: personalisation of social media content alone is not sufficient in shaping purchase intention. Consumer confidence is reliant on trust, electronic word of mouth (E-WOM), user-generated content (UGC) and social proof (Rahman et al., 2025; Akbar & Dellyana, 2023; Al-Amin & Scorita, 2025; Ngo et al., 2024). Millennials and Generation Z also prefer different types of content in terms of their information acquisition: the former prefers informative content, the latter short-form video and peer interaction and content provided by users or influencers (Al-Amin & Scorita, 2025; Paliwal et al., 2025; Yadav & Rai, 2017; Agrawal, 2022). Algorithmic personalisation thus aids in discovery but what truly influences a purchase decision is the combination of social influence and consumer content consumption, and how that differs between the two generations. It is these points which are argued upon here: algorithms are not the singular driver of a consumer's purchasing decisions. Rather, it is the combination of personalisation and social influence which has an impact, and the ways in which this interaction differs between the two generations. This review also makes a further argument: that the literature treats algorithmic curation and social proof as two separate mechanisms when they are better understood as one. The recommendation algorithm decides which reviews, influencer videos and comments a consumer sees, so the social proof that is meant to check the algorithm is itself shaped by it.

While there is a wealth of research into recommendation algorithms, consumer behaviour and generational differences taken individually, there is little which combines these three areas together. Fadilah et al. (2025) and Rahman et al. (2025) both investigate recommendation algorithms (the former focused around impulse buying and the latter looking more broadly at social networks, reviews and social influence). Paliwal et al. (2025) and Akbar and Dellyana (2023) investigate generational differences in social media use and purchasing between Millennials and Generation Z, while Mishra (2024) tackles more generally the area of consumer behaviour in the contemporary digital age. It is of utmost importance to synthesise this information, as none of the individual pieces examine how algorithmic personalisation turns into an actual purchase. Understanding how personalisation, consumer confidence and generational differences interact is essential to marketers developing recommendation algorithms, in their design, to what extent they utilise social proof, and in their understanding of how younger consumers engage with both.

This review is based on twelve studies published between 2017 and 2025 which take particular aim at social commerce platforms such as Instagram, TikTok and YouTube, as well as at the purchasing behaviour of Millennial and Generation Z consumers specifically. The argument is made concisely with the paper containing brief findings which have been thematically grouped together. The methodology is explained followed by the findings discussed in three sections: algorithmic curation and product discovery, social proof and consumer trust,

and generational purchasing dynamics. The review then concludes and highlights both the limitations of the literature and the potential future directions in which research could develop.

Methodology

To collect papers for this study, Google Scholar was searched for scholarly articles published between 2017 and 2025, using keywords like "recommendation algorithms," "social media consumer behaviour," "Gen Z impulse buying," "Millennial purchasing habits," "social proof in e-commerce," "personalised recommendation purchase intention," "influencer credibility trust" and "electronic word-of-mouth purchase intention."

Twelve studies met the criteria. Eight are literature or bibliometric reviews that pull together existing research on the algorithms, consumer behaviour and generational differences (Akbar & Dellyana, 2023; Al-Amin & Scorita, 2025; Fadilah et al., 2025; Mishra, 2024; Paliwal et al., 2025; Rahman et al., 2025; Souza, 2024; Yadav & Rai, 2017). To avoid the risk of relying on the information contained in other reviews, four peer-reviewed empirical articles, each of which employed a distinctive methodological design (survey, experimental, and structural equation modelling), were included in the study: Agrawal (2022) (behavioural differences between Generation Y and Generation Z online shoppers), Alcántara-Pilar et al. (2024) (influencer credibility and intention to purchase on TikTok), Ngo et al. (2024) (electronic word of mouth and Generation Z's purchase intention), and Zhao et al. (2025) (personal recommendation and purchase intention). Instead of highlighting the key points of each paper, all documents were read together to extract recurring themes, trends, contradictions, and gaps in the research field, which were subsequently grouped together into the three themes highlighted below.

Theme 1: Algorithmic Curation and Product Discovery

Recommendation algorithms have largely replaced the traditional form of mass advertising in social media feeds. Instead of using the same advertisements for all users, recommendation algorithms utilize their knowledge of what has interests a specific user to generate personalized recommendations (Fadilah et al., 2025; Mishra, 2024; Zhao et al., 2025). As a result, recommendation algorithms narrow the search scope for any given user to only those products that have piqued their interest, removing a substantial amount of time from the traditional discover-find-buy cycle. The reviewed articles point to three limitations in how far this can be relied on: feeds draw attention but do not reliably turn it into sales, more accurate targeting can work against itself, and algorithms often fail to recommend what is actually relevant.

The first limitation is that feeds draw attention but do not reliably turn it into sales. The reviewed articles demonstrate conflicting evidence as to whether the feeds of recommended products drive impulsive consumption while a consumer is engaged in the discovery (Fadilah et al., 2025). In contrast, recommendation algorithms limit brand recall-driven purchases due to the recommendation feeds feeling similar to traditional brand advertising (Paliwal et al., 2025).

Taken together, the reviewed articles demonstrate that while feeds draw attention, they are inconsistent in their ability to drive sales, thus, limiting the ability of recommendation algorithms to be relied upon as a sales tool. The one experiment in this review points the same way. In a scenario experiment with 705 Chinese university students, more accurate recommendations did raise students' sense of competence in judging the product ($\beta = 0.325$, $p < 0.01$), but that sense of competence explained only 11.79% of the effect of accuracy on purchase intention, below the 20% the authors need before they count competence as a mediator between the two (Zhao et al., 2025).

The second limitation is that recommendation algorithms impose a limitation on their ability to target consumers due to the same mechanisms that drive their success. Some consumers actively avoid algorithmic targeting altogether. The more accurately an algorithm knows a consumer, the more opportunities it has to recommend products that the consumer has no interest in purchasing (Fadilah et al., 2025). This is not an opposing concept to the first limitation, rather it is closely related to it. The only significant difference between targeting and traditional advertising is the scale, with Fadilah et al. (2025) arguing that targeting fails due to the sheer volume of data it collects, while others offer no explicit distinction. Across the reviewed articles, little attention is paid to identifying the threshold between targeting that aids consumer discovery and targeting that limits the consumer's attention.

The third limitation is that algorithms often fail to recommend something relevant. This limitation is directly connected to the second. Recommendation algorithms rely on their ability to narrow the field of search, and failure to do so defeats their primary purpose. Thus, the effectiveness of recommendation algorithms is entirely dependent on the ability to correctly and timely suggest a relevant product. Additionally, they fail when presenting the consumer with products that they already own, have declined, or that have no relevance to their needs at the given moment. Their effectiveness is driven almost exclusively by convenience, as they reduce the amount of effort a consumer must expend on discovery. The final stages of a purchase, according to the reviewed literature, still depend on the consumer's ability to compare and deliberate, which feeds are unable to facilitate.

Theme 2: Social Proof and Consumer Trust

While the recommendation algorithms play a key role in determining what a consumer will see, it is the role of social proof in influencing trust that determines whether a consumer will purchase any given product. Reviews, ratings, and other forms of social proof allow consumers to reduce the risk of an unwarranted purchase recommended by algorithms (Ngo et al., 2024; Rahman et al., 2025; Yadav & Rai, 2017). In a survey of 280 Vietnamese Generation Z social media users, the quality of eWOM information had the strongest indirect effect on purchase intention ($\beta = 0.200$, $p = 0.006$), while the amount of eWOM ($\beta = 0.033$, $p = 0.484$) and how credible it looked ($\beta = 0.052$, $p = 0.368$) had no significant effect (Ngo et al., 2024). In other words, what the content said mattered more than how many people said it or how credible they looked. Social proof has long predated recommendation algorithms, with written reviews being

the primary means of reducing uncertainty prior to algorithmic curation (Yadav & Rai, 2017). However, social proof has since expanded to include video content made by influencers which serve to reduce suspicion and provide genuine opinions of the products that the algorithms suggest (Mishra, 2024; Rahman et al., 2025).

Influencers provide a similar function to community marketing which uses a relatable figurehead to convince consumers to purchase a product despite their misgivings advised by the recommendation algorithms (Al-Amin & Scorita, 2025; Alcántara-Pilar et al., 2024). A study of 880 active TikTok users used structural equation modelling and then a neural network check to test which predictors mattered most, and it supported every relationship in the authors' model linking influencer credibility, trust, loyalty and purchase intention (Alcántara-Pilar et al., 2024). Community marketing is basically an enhanced version of influencer marketing, with an entire community of trustworthy figures supporting a given product rather than a single influencer doing so alone (Akbar & Dellyana, 2023). Influence in both influencer and community marketing are the driving force of consumer trust, which physically facilitates a purchase decision despite the consumer's suspicions seeded by the algorithmic curation.

Across the reviewed articles, algorithmic curation and social proof are largely considered to be separate concepts. However, there is reasonable reasoning to consider it a single mechanism. A recommendation feed is an algorithmic mechanism which ultimately determines what the consumer will see. Every interaction within the platform is driven by the said recommendation algorithm including which review appears first or which influencer content will be suggested after engaging with a product listing. Therefore, when a consumer actively browses a recommendation feed and reads reviews, they are interacting with the recommendation algorithm.

The same dynamic applies to influencers; they must interact with the recommendation algorithm in order to promote their content before the consumers who are likely to engage with them. This is why The reviewed literature suggests that consumer trust is not explicitly driven by influencer endorsement but rather the organically appearing comments on the same page (Rahman et al., 2025). In truth, the mechanism behind organic comments are similar to influencer marketing as they must engage with the recommendation algorithm in order to appear prominently in the feed. Therefore, across both influencer and community marketing, the recommendation algorithm determines engagement which determines visibility and trust.

Theme 3: Generational Purchasing Dynamics

The way that algorithmic curation and social proof affect Millennial and Generation Z consumers are different, though the reviewed literature does not explicitly clarify what these differences are apart from the obvious generational differences in purchasing behavior. While Millennials engage with recommendation feeds more deliberately, seeking branded content and relying on reviews before making a purchase decision (Agrawal, 2022; Mishra, 2024; Souza, 2024), younger consumers are more impulsive, engaging with the recommendation feeds on

short-form platforms and relying on the advice of ordinary users (Al-Amin & Scorita, 2025; Fadilah et al., 2025; Yadav & Rai, 2017).

The difference between the two similar marketing approaches is largely reflected in the tendency of the generations to engage with the recommendation feeds. For Millennials, the purchase decision comes significantly after discovering a product on a recommendation feed, allowing them to weigh their options and consider alternatives (Agrawal, 2022; Mishra, 2024; Souza, 2024). In contrast, the Generation Z purchase decisions are significantly more impulsive, driven by the recommendation feeds on the same platforms which they enjoy as users (Al-Amin & Scorita, 2025; Fadilah et al., 2025; Yadav & Rai, 2017). One might argue that the differences are an inherent tendency of the two generations, but it is possible that the differences are entirely caused by the differences in platforms that the two generations use more often. The difference will be explored further in the next section.

Some considerations suggest that the differences in purchasing patterns reflect the differences in platform design rather than the users' tendencies. As the literature demonstrates, the platforms designed primarily for Millennials serve as a discovery tool that sits alongside the actual act of purchase (Mishra, 2024; Souza, 2024). These platforms do not present the recommendation feeds as a marketplace, so the opportunity to click and purchase is alongside the opportunity to merely scroll (Al-Amin & Scorita, 2025). Contrastingly, the platforms used predominantly by the younger generation do not offer any distinction between the recommendation feeds and the actual marketplace, facilitating the short deliberation period between discovering a product and clicking to purchase. As a result, the two generations, which use different platforms behave differently while interacting with recommendation feeds, limiting the opportunity for purchase for those who use single-purpose platforms.

One argument claims that generational differences demonstrate that younger consumers are significantly more sceptical of the intent of recommendation feeds than their millennial counterparts. However, as the literature cited below does not demonstrate any specific evidence in support of this argument, it cannot be considered more than an unsubstantiated claim. Reviews which discuss the relationship between the exposure to algorithmic curation and consumer trust only describe the overall trends across the tested populations without distinguishing between Millennials and younger generations (Al-Amin & Scorita, 2025; Fadilah et al., 2025; Yadav & Rai, 2017). Thus, there is no evidence to suggest that skepticism is an inherent quality of the younger generation. If anything, the literature cited supports the alternative argument that younger consumers are impulsive, as their deliberation gap is significantly smaller than among Millennials. It is possible to argue that such reduced deliberation is responsible for increased scepticism, however the evidence to demonstrate this assumption is not presented in the reviewed literature either, leaving the claim open for further exploration.

Conclusion & Literature Limitations

Summary of Core Synthesis

Across the reviewed papers, algorithms serve as the initial enabler of product discovery, where choices become narrowed down by the algorithm and the consumer is exposed to faster search results (Theme 1). This exposure cannot be transformed into trust by the algorithm alone, which is why reviews, user opinion, and influencers are important to engross the consumer at the final stages of the decision-making process (Theme 2). The relative importance of these three themes depends on the age group of the respondents where Millennials seem to be more careful and trusting and members of Generation Z engage with feeds as a discovery tool, a proof mechanism and a one click enabler (Theme 3).

The central claim of this review is that algorithmic curation and social proof are better understood as one mechanism than two. The reviewed literature treats them as separate, with the algorithm creating exposure and social proof creating trust. But the recommendation algorithm decides which reviews, influencer videos and comments a consumer sees, so the trust that is meant to check the algorithm is shaped by the algorithm itself (Theme 2). This changes how the other two themes should be read. The limits of recommendation feeds in Theme 1 are limits on what a feed can do on its own, and the generational differences in Theme 3 may partly come from how closely each generation's platforms tie the feed and the proof together.

These themes are consistent with each other while raising questions about the consumer behaviour literature's emphasis on age group differences in the context of the reviewed studies.

All three themes align around the idea of an algorithm-influenced choice system which is then proven by another algorithm in the shape of user reviews. The paper's findings on Gen Z and Millennials may reflect regional variations in technology development rather than inherent age group differences.

Critical Literature Limitations

The review highlighted several limitations of the included papers. First of all, the identified articles failed to consider the role of different algorithm recommendation models in the level of trust that people place in brands and peers which may have significant implications for the trust they have in social algorithm recommendations. Additionally, most of the papers did not take the type of product into consideration, even though people may respond differently to algorithmic suggestions and social proof for hedonic versus utilitarian goods. Zhao et al. (2025) is an exception here, showing that the effect of personalised recommendation on consumers' psychological needs, which in turn shape purchase intention, does vary by product category, however none of the other studies reviewed test whether social proof or generational differences behave the same way across product types. In that experiment, which compared smartphones with shoes, the effect of recommendation diversity on consumers' sense of autonomy was stronger for shoes ($\beta = 0.215$, $p < 0.01$), and so was the effect of recommendation novelty on their sense of relatedness to the platform ($\beta = 0.164$, $p < 0.01$), while the effects of accuracy and novelty on competence did not differ by category ($p = 0.511$ and $p = 0.706$) (Zhao et al., 2025).

Second, the review revealed that the majority of papers utilized the same respondents in terms of geography, as several studies reported their findings on Asian populations, including

Indian and Indonesian participants. The validity of some conclusions in other geographical areas, especially emerging economies, may be compromised as differences in smartphone adoption, online shopping, and digital payment may affect the role of technology in people's consumption and behavior, not their age.

Therefore, although the review concluded that differences in the perception of algorithmic and social recommendations stem from people born at different times, it is unclear to what extent these findings are generalizable to different markets. Third, the review did not show whether any of the papers utilized a longitudinal design, which would enable understanding how consumer behavior and attitude change as they age. The study suggests that the cross-sectional design allowed the authors to identify the characteristics of consumers who engage in certain behaviors or the determinants of their appearance. The included papers examined consumer agencies in terms of their receptivity to recommendations, however, little was said on their abilities, algorithmic literacy and manipulation.

The findings may help marketers engage millennial consumers more responsive to recommendation algorithms than before, taking into account their more liberal views towards shopping in general. However, the review, like the majority of the papers it includes, fails to highlight the mechanisms by which consumer trust in algorithm recommendations emerges. Although the review highlighted the existence of intergenerational differences in people's perception and use of social and algorithmic recommendations stemming from differences in their upbringing, it is unclear if these differences apply to all types of products equally.

Additionally, the review and the papers it includes fail to consider the possibility that consumer fatigue may emerge when interacting with recommendation algorithms, as well as the length of such interaction affecting their reliance on reviews, brand image, and product types. Furthermore, beyond the differences in the receptivity to recommendations, the review barely comments on the differences in consumer behavior. Finally, the review barely elaborates on consumer differences beyond their age, with millennial and centennial consumers being more receptive to algorithmic marketing than previous generations.

Future Research Directions

The future research on the role of technological development on consumer behaviour may build upon the findings in several directions, which may contribute to obtaining a more comprehensive understanding of the phenomenon. First of all, it is essential to conduct a longitudinal analysis of actual spends on products with the use of algorithmic impulses and one-click buying to assess the real impact of such forces on the economy, focusing on young people's purchasing behaviour in particular.

Further, there is a need for research on the role of AI-based shopping assistants and chat-bots in influencing consumers' trust, perception, and buying behaviour through their recommendations and engagement. Another interesting direction of the research may be a comparison of the role of algorithmic recommendations for hedonic and utilitarian products and

services. This would include an analysis of individual preferences and the influence of social feedback loops or word of mouth in different cultures.

The analysis of algorithm literacy and consumer agency is another interesting direction of research, which would help understand how people navigate algorithmic recommendations or use their purchasing power to influence them rather than being passive acceptors of suggestions. Further, the comparison may be made between individualistic and collectivistic cultures to identify whether preferences are culturally dependent, different across generations, or are driven by algorithms. Each of these research directions may contribute to a more comprehensive understanding of the ways algorithms, social influence, design choices, product type, and consumer behaviour intertwine.

Works Cited

Agrawal, D. K. (2022). Determining behavioural differences of Y and Z generational cohorts in online shopping. International Journal of Retail & Distribution Management, 50(7), 880-895. https://doi.org/10.1108/IJRDM-12-2020-0527

Akbar, A., & Dellyana, D. (2023). Maximizing simultaneous social media communication: A study on purchase intention among Millennials and Generation Z in the online language learning industry. International Journal of Current Science Research and Review, 6(6), 3635-3656. https://doi.org/10.47191/IJCSRR/V6-i6-54

Al-Amin, M., & Scorita, K. B. (2025). Digital marketing strategy for Millennials and Gen Z: A literature review. International Journal of Financial Economics (IJEFE), 2(6), 620-629.

Alcántara-Pilar, J. M., Rodriguez-López, M. E., Kalinić, Z., & Liébana-Cabanillas, F. (2024). From likes to loyalty: Exploring the impact of influencer credibility on purchase intentions in TikTok. Journal of Retailing and Consumer Services, 78, 103709. https://doi.org/10.1016/j.jretconser.2024.103709

Fadilah, N., Masarroh, I., & Wildan, M. A. (2025). Effectiveness of recommendation algorithms on impulsive buying in e-commerce platforms: A systematic literature review. Priviet Social Sciences Journal, 5(8), 186-198. https://doi.org/10.55942/pssj.v5i8.485

Mishra, A. (2024). Consumer behaviour in digital age: A comprehensive review. ANALISTA: International Journal of Multidisciplinary Research, 2(1), 1-9.

Ngo, T. T. A., Vuong, B. L., Le, M. D., Nguyen, T. T., Tran, M. M., & Nguyen, Q. K. (2024). The impact of eWOM information in social media on the online purchase intention of Generation Z. Cogent Business & Management, 11(1), 2316933. https://doi.org/10.1080/23311975.2024.2316933

Paliwal, R., Sharma, A., & Khan, M. Z. (2025). Navigating the digital divide: Social media marketing strategies for Millennials and Generation Z. International Journal of Global Multidisciplinary Research and Analytics, 1(2), 88-100. https://doi.org/10.5281/zenodo.17313053

Rahman, J., Hossain, M. A., & Datta, S. (2025). Understanding the impact of social networks on consumer buying behavior: A comprehensive review from a sociological perspective. International Journal of Business, Arts and Scientific Study, 6(1), 74-85. https://doi.org/10.70818/ijbass.2025.v06i01.2508

Souza, M. A. de. (2024). Consumer differences between Generation Y (Millennials) and Generation Z: A bibliometric review. Revista Eletrônica de Administração, 23(1), 104-118.

Yadav, G. P., & Rai, J. (2017). The Generation Z and their social media usage: A review and a research outline. Global Journal of Enterprise Information System, 9(2), 110-116.

Zhao, L., Fu, B., & Bai, S. (2025). Understanding the influence of personalized recommendation on purchase intentions from a self-determination perspective: Contingent upon product categories. Journal of Theoretical and Applied Electronic Commerce Research, 20(1), 32. https://doi.org/10.3390/jtaer20010032

AI-Based Predictive Maintenance With Generative AI Integration for Multi-Sensor Diagnosis and Prediction of Aircraft Gearbox Failures By Vikrant Raj Sathyamurthy

Abstract

Aircraft gearbox failures can compromise system reliability, increase maintenance costs, and cause unexpected downtime. Although artificial intelligence (AI) has strong potential for predicting and diagnosing gearbox faults using multi-sensor data, many systems produce results that are difficult for maintenance personnel to interpret and apply. This review examines how generative artificial intelligence (GenAI) can complement AI-based predictive maintenance for aircraft gearboxes. It synthesizes research on machine learning, deep learning, and multi-sensor data fusion using vibration, temperature, rotational speed, and operational measurements to predict failures in bearings, gears, and shafts. The review then considers how explainable AI and GenAI can translate model outputs into understandable, evidence-based maintenance recommendations. Its primary contribution is a conceptual hybrid framework connecting predictive accuracy with maintenance decision support, specifically addressing the critical challenge of grounding Generative AI and predictive model outputs in aviation certification requirements and approved technical documentation. Key limitations include scarce representative aircraft fault data, model generalization, explanation reliability, and the need to ground GenAI outputs in validated evidence and approved technical documentation.

Keywords Engineering, Mechanical Engineering, Predictive Maintenance, Aircraft Gearboxes, Generative Artificial Intelligence

Introduction

Aircraft gearbox reliability is a safety and operational priority because gears, shafts, bearings, and clutches transfer power through tightly coupled systems in which a localized defect can trigger a rapid failure cascade. The accompanying research presentation highlights three recent gearbox- or drivetrain-related aircraft accidents: the 2023 U.S. Air Force CV-22B accident near Yakushima, Japan, with 8 fatalities; the 2022 U.S. Marine Corps MV-22B accident near Glamis, California, with 5 fatalities; and the 2016 CHC Helikopter Service Flight 241 accident near Turoy, Norway, with 13 fatalities. Across these three cases, 26 lives were lost. The reported initiating mechanisms included fatigue cracking in a high-speed pinion gear, hard-clutch engagement that produced severe drivetrain torque spikes, and a subsurface fatigue crack in a main-gearbox planet gear. These events show why detecting weak degradation signals before they develop into systemlevel failures is an important aviation research problem. [16]

The consequences of gearbox and propulsion-system degradation also extend beyond immediate safety. The business-case evidence summarized in the presentation reports that approximately 15% to 20% of the global Airbus A220 fleet, about 80 to 90 of roughly 450 aircraft, had been parked during the cited period; some engine shop visits faced turnaround times of more than 300 days; and upgrade programs targeted a 40% increase in on-wing engine life

through redesigned hardware. Although these figures represent broader fleet and propulsion-maintenance pressures rather than direct evidence of aircraft gearbox-fault prediction, they illustrate how component reliability, maintenance capacity, spare-part availability, aircraft utilization, and schedule continuity are interconnected. Earlier warnings can therefore create value not only by reducing technical risk, but also by supporting maintenance planning and aircraft availability. [16]

Traditional scheduled inspections and reactive repairs may not reveal a developing defect early enough, especially when initial fault signatures are weak, intermittent, or masked by noise and changing operating conditions. Modern Health and Usage Monitoring Systems (HUMS) provide continuous or repeated measurements such as vibration, temperature, rotational speed, oil condition, acoustic signals, torque, and flight conditions. Vibration is particularly sensitive to bearing wear, gear damage, shaft misalignment, and imbalance, while thermal and operational measurements add information about friction, lubrication, loading, and the state in which a signal occurs. Combining these sensor streams can therefore provide a more complete representation of gearbox health than relying on one measurement alone. [1,3,4,6]

Current AI research provides a strong technical foundation for this transition from reactive to predictive maintenance. Decision Trees, Support Vector Machines, Random Forests, XGBoost, CatBoost, Multi-Layer Perceptrons, convolutional networks, recurrent networks, transfer-learning models, and transformer-based approaches have all been investigated for rotating-machinery diagnosis or aerospace operational prediction. Reported examples include 97.2% accuracy for KNearest Neighbours on a ten-condition rotating-machinery benchmark, an R2 of 0.939 for XGBoost on 742,625 helicopter turboshaft-engine operational records, and 94.90% classification accuracy for a transferred CNN using only five vibration signals per condition. These results demonstrate the potential of AI, while also showing that performance depends on the task, dataset, operating range, and availability of representative labeled data. [3,5,9]

Multi-sensor fusion is especially relevant because individual sensors provide incomplete views of the same mechanical event. In one controlled gearbox experiment covering 12 failure classes, three individual microphone streams achieved 89.58%, 90.28%, and 88.89% classification accuracy. Early fusion of the three streams increased accuracy to 100%, a 9.72-percentage-point improvement over the best individual stream. This result provides quantitative evidence that complementary sensing can improve diagnosis, but it should not be interpreted as guaranteed aircraft performance because laboratory conditions do not reproduce the full variability of flight operations. [14]

Despite these advances, three gaps limit practical adoption in aviation. First, many studies examine one component, one fault type, or one dominant sensor, whereas an aircraft gearbox contains interacting bearings, gears, and shafts operating under changing loads and speeds. Second, strong benchmark accuracy does not ensure generalization to real aircraft, where representative fault data are scarce and imbalanced. Third, many high-performing models function as black boxes, providing a diagnosis without an engineering explanation that maintenance personnel can independently examine. In a safety-critical setting, a useful system

must connect each prediction to sensor evidence, confidence, operating context, and human review rather than provide a fault label alone. [1,7,9,15]

Accordingly, this review investigates a conceptual hybrid Predictive AI + Explainable AI + Generative AI framework for near-real-time aircraft gearbox monitoring. The primary novelty of this work does not lie in developing standalone algorithms, but rather in establishing a structured research pathway for grounding AI and Generative AI outputs in aviation certification requirements, airworthiness standards, and approved technical documentation to support human oversight. Instead, it defines a human-centered research pathway in which qualified engineers retain authority while evaluating whether integrated prediction, traceability, and communication can make AI-based maintenance support more understandable and actionable. [1,2,14,15]

Literature Selection and Review Methodology

To establish a comprehensive and rigorous synthesis of artificial intelligence-based predictive maintenance, multi-sensor data fusion, and generative explainability for aircraft gearboxes, a structured literature review methodology was executed across major scientific repositories, including IEEE Xplore, ScienceDirect, Scopus, Google Scholar, and the AIAA Aerospace Research Central using targeted Boolean combinations of keywords covering aircraft gearbox failures, predictive machine learning, multi-sensor data fusion, and generative AI decision support. Retrieved records were filtered for publication between 2018 and 2026, prioritizing peer-reviewed journal articles, conference proceedings, and official accident investigation reports that provided quantitative diagnostic metrics on rotating machinery or evaluated human-in-the-loop engineering decision support systems, while excluding non-peer-reviewed white papers and articles lacking technical transparency.

Results and Discussion

1. Aircraft Gearbox Health Monitoring and Predictive Maintenance

Aircraft Gearbox System Layout

Aircraft gearbox systems contain interconnected components, including gears, shafts, and bearings, which work together to transmit power. Gears transfer mechanical power, shafts transmit torque, and bearings support the rotating shafts. Because these components interact, a fault in one can affect the measurements observed throughout the gearbox.

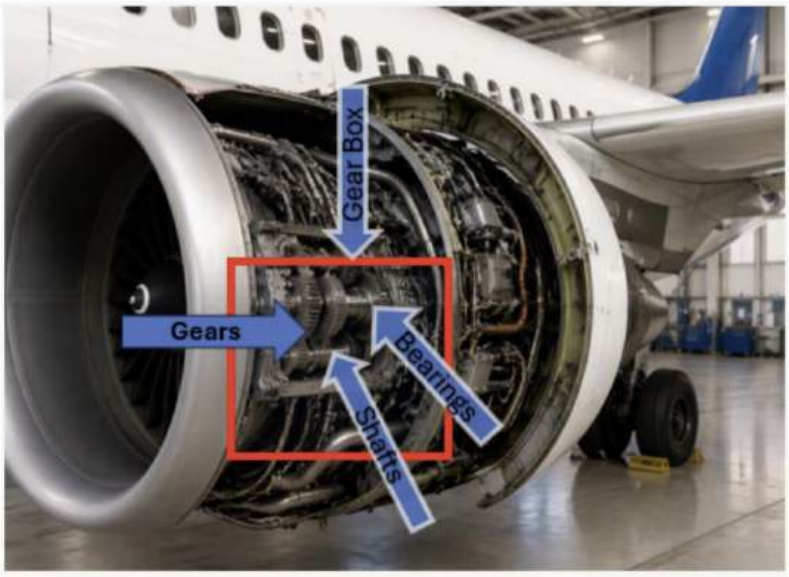

Figure 1. Aircraft gearbox system layout showing the main components.

The gearbox layout provides important context for predictive maintenance. Different faults can affect vibration, temperature, acoustic signals, and other measurements. By combining these sensor measurements, AI-based systems can better identify abnormal gearbox behavior and determine the likely source of a fault. Therefore, understanding the gearbox layout provides a foundation for the multi-sensor AI-based diagnosis discussed in the following sections.

Predictive Maintenance and HUMS for Aircraft Gearbox Monitoring

Predictive maintenance has increasingly shifted aircraft gearbox maintenance from fixed inspection schedules toward continuous condition-based monitoring. Unlike traditional maintenance strategies that rely on predetermined service intervals or corrective repairs after failures occur, predictive maintenance evaluates component health using operational sensor data to identify early signs of mechanical degradation. This proactive approach can reduce unnecessary maintenance, minimize unexpected aircraft downtime, extend component service life, and improve operational safety by allowing developing faults to be addressed before they result in severe damage .1,3,10 As aircraft systems become increasingly complex, predictive maintenance has become an important strategy for maintaining reliability while controlling maintenance costs. Central technology enabling this transition is the widespread use of Health and Usage Monitoring Systems (HUMS), which provide continuous measurements of gearbox operating conditions throughout flight. HUMS can incorporate vibration, temperature, rotational speed, lubrication-related measurements, acoustic signals, and other flight parameters to provide information about the condition of aircraft components under realistic operating conditions .2,1 These measurements allow changes in component behavior to be monitored over time and provide the operational information necessary for identifying abnormal conditions and supporting condition-based maintenance.

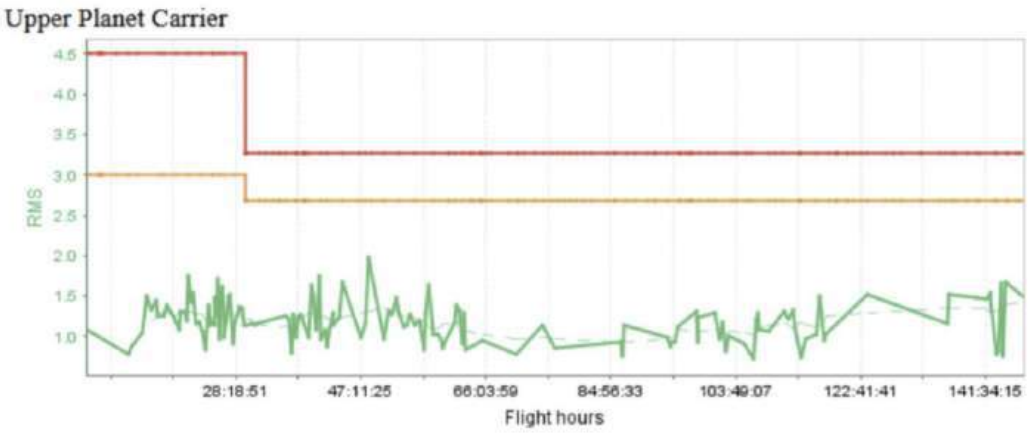

Figure 2. Root-mean-square acceleration (gRMS) of the EC 725 AP upper planetary carrier across flight hours, with amber and red HUMS alert thresholds. Adapted from AIBN (2018), as reproduced in Moreira et al.1.

Figure 2. illustrates an operational HUMS condition indicator by showing the root-mean-square acceleration of the EC 725 AP upper planetary carrier across flight hours. The graph includes amber and red alert thresholds that represent increasing levels of maintenance concern. Rather than evaluating the vibration signal only as an isolated measurement, HUMS uses changes in the condition indicator and its relationship to established thresholds to support maintenance decisions. Although Figure 1 illustrates a vibration-based condition indicator, modern monitoring systems can incorporate multiple complementary measurements to provide a more complete assessment of gearbox health. Research has shown that combining different sensor streams can provide information that may not be captured through vibration measurements alone, improving the ability to identify changes in component condition .6,11 Consequently, multi-sensor monitoring provides an important foundation for modern aircraft health management and predictive maintenance systems.

The growing availability of benchmark datasets has further supported advances in predictive maintenance research by allowing researchers to evaluate diagnostic algorithms under defined operating conditions. Publicly available datasets such as the NASA Bearing Dataset, MaFaulDa, PHM09 Gearbox Dataset, and helicopter HUMS datasets have been used to investigate bearing and gearbox fault diagnosis and condition monitoring .7,9,1 However, existing studies also recognize that many laboratory datasets differ substantially from real aircraft environments because they may involve controlled operating conditions, simplified loading profiles, or limited representations of real-world variability. Consequently, incorporating operational flight data and multiple sensor sources is important for improving the robustness and generalization of predictive maintenance models when they are applied to practical aviation environments .[2,6,3]

Despite these advances, the large volume and complexity of data generated by modern monitoring systems create analytical challenges that conventional diagnostic techniques may not efficiently address. Extracting meaningful information from continuously changing, highdimensional sensor data requires computational methods capable of identifying complex

relationships and subtle fault patterns across different operating conditions .3,1 These challenges have contributed to the increasing use of Artificial Intelligence techniques, particularly machine learning and deep learning, for automated fault diagnosis and predictive maintenance .8,9,5 The next section therefore examines how Artificial Intelligence models can analyze aircraft gearbox sensor data to diagnose and predict mechanical faults, with particular attention to traditional machine learning and deep learning approaches.

2. AI-Based Fault Diagnosis and Prediction Models

Artificial Intelligence has become increasingly important in predictive maintenance because machine learning and deep learning models can identify complex patterns within vibration, temperature, and operational sensor data that may be difficult to capture using conventional diagnostic methods .9 Public datasets, including the PHM09 Gearbox Dataset, NASA Bearing Dataset, MaFaulDa, and helicopter HUMS datasets, have provided standardized environments for evaluating these approaches in rotating-machinery fault diagnosis and condition monitoring .12,11,2 However, the performance reported in these datasets does not necessarily translate directly to real aircraft applications because differences in sensor configurations, operating conditions, and the limited availability of representative failure data can restrict model generalization .1,10 Consequently, evaluating both traditional machine learning and deep learning approaches across controlled datasets and operational data is important for determining their practical suitability for aviation predictive maintenance.

Traditional Machine Learning Approaches

Classical machine learning approaches typically rely on engineered features extracted from sensor measurements in the time or frequency domain and have established strong baselines for rotating-machinery fault classification .5 On a ten-fault-condition benchmark, K-Nearest Neighbours achieved 97.2% classification accuracy, while Support Vector Machines and a wide neural-network classifier achieved 96.2% and 94.1%, respectively, using 80% of the available data for training .5 These results demonstrate that relatively conventional models can achieve high diagnostic accuracy when informative features and representative training data are available.

Ensemble learning methods have also demonstrated strong performance when applied to operational aerospace data. In a large operational dataset containing 742,625 records from helicopter turboshaft engine operation, machine-learning models were evaluated for predicting Engine Torque Fraction (ETF) under different operating conditions .10,3 XGBoost, CatBoost, and Random Forest substantially outperformed a linear-regression baseline, with XGBoost achieving an RMSE of 0.006 and an R2 of 0.939 over the ETF 0.97–1.05 range, while CatBoost and Random Forest achieved R2 values of 0.905 and 0.906, respectively .3 These results are important because they demonstrate that traditional machine-learning models can capture nonlinear relationships in large operational aerospace datasets, although the task represents operational parameter prediction rather than direct gearbox-fault classification.

Table 1. Comparison of tree-based ensemble models for Engine Torque Fraction prediction on operational helicopter turboshaft-engine data. RMSE & R2 are reported for the ETF operating ranges 0.97–1.05 & 0.90–0.99 .3.

Model	RMSE (ETF 0.97–1.05)	R^2 (ETF 0.97–1.05)	RMSE (ETF 0.90–0.99)	R^2 (ETF 0.90–0.99)
XGBoost	0.006	0.939	0.009	0.898
CatBoost	0.007	0.905	0.008	0.903
Random Forest	0.007	0.906	0.011	0.846
Linear regression (ceiling)	0.019	0.458	0.024	0.382

The results in Table 1 further illustrate the advantage of nonlinear ensemble models over the linear-regression baseline. Across the ETF 0.97–1.05 range, XGBoost produced the lowest RMSE (0.006) and highest R2 (0.939), while the linear-regression model remained below R2=0.46 . 3 The ensemble results therefore indicate that traditional machine learning can remain highly effective for complex operational aerospace data, particularly when structured sensor features are available.

Deep Learning Approaches

Deep-learning approaches extend traditional machine learning by learning feature representations directly from complex data and can therefore reduce dependence on manually engineered diagnostic features. Recurrent architectures such as Long Short-Term Memory (LSTM) networks are particularly suited to sequential monitoring data because they can model relationships between measurements across time. Figure 3 illustrates two LSTM-based architectures proposed for aircraft-gearbox health classification. The architecture uses recurrent layers followed by fully connected layers to map sequential monitoring signals to a binary health classification, representing good and bad gearbox conditions (Mallikarjuna et al., 2021).

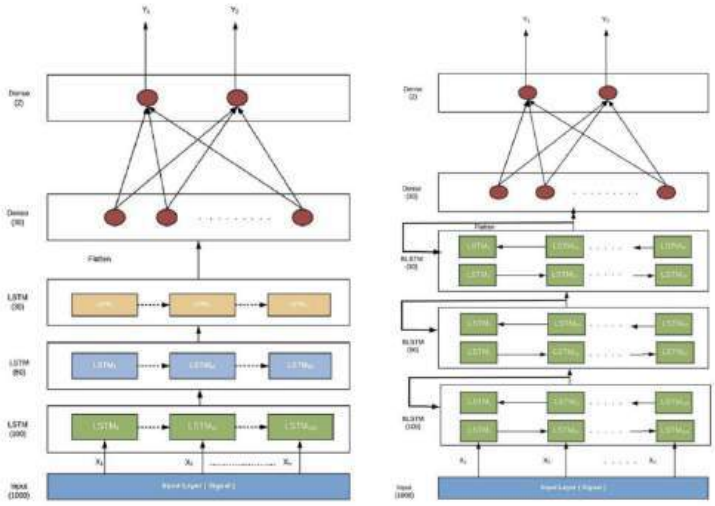

Figure 3. Proposed deep-learning architectures for two-class classification of aircraft-gearbox health conditions. (a)LSTM architecture consisting of three recurrent layers followed by fully connected classification layers. (b) Bidirectional LSTM architecture that processes the input sequence in both forward and backward directions. Adapted from Mallikarjuna et al. (2021).

Convolutional Neural Networks (CNNs) provide another important deep-learning approach because they can automatically learn local patterns from vibration signals and other structured inputs. CNN-based architectures have demonstrated strong performance on industrial gearbox datasets and can reduce reliance on manually extracted features (Li & Wang, 2020; Zhang et al., 2020). Transfer learning is particularly relevant to aviation because collecting large numbers of labeled failure examples from aircraft is difficult. Cao et al. (2018), for example, reported 94.90% classification accuracy when only 5 vibration signals per condition were available for fine tuning a transferred CNN model, while accuracy increased to approximately 99–100% with larger training sets. However, performance declined substantially when the available training data became extremely limited, reaching 72.22% when only 2% of the training data were used (Cao et al., 2018). These results demonstrate both the potential of transfer learning to reduce data requirements and the continuing sensitivity of deep-learning models to limited labeled data.

More recent architectures have attempted to improve this trade-off between model complexity and data requirements. Lightweight channel-attention CNNs have demonstrated transfer-learning capabilities while reducing the number of model parameters. .[13] Transformer-based tabular models have also been investigated for rotating-machinery datasets, providing another approach for classification when structured sensor features rather than raw signals are used . 9 Together, these approaches demonstrate that deep learning can provide powerful automated feature extraction, but their effectiveness remains dependent on the quality, diversity, and representativeness of the training data.

Cross-Method Evaluation and Aviation Suitability

The comparison of AI approaches should extend beyond classification accuracy because aviation applications also require consideration of data requirements, computational demands, generalization, and operational reliability. On rotating-machinery benchmarks, Magadán et al.9 reported strong performance from a tabular-transformer approach across accuracy, precision, recall, and F1-score. Similarly, transfer learning allowed a CNN-based model to achieve 94.90% accuracy with only five training samples per condition and approximately 99–100% accuracy with larger training sets (Cao et al., 2018). These findings suggest that both transformer-based and transfer-learning approaches may be valuable when labeled fault data are limited.

Traditional ensemble methods remain particularly competitive when structured operational data is available. The XGBoost, CatBoost, and Random Forest results in Table 1 demonstrate that these models can capture nonlinear relationships in large operational aerospace datasets and substantially outperform a linear-regression baseline for Engine Torque Fraction prediction .3 This finding is relevant to aircraft predictive maintenance because operational aircraft data can contain complex relationships between sensor measurements and system behavior. However, the comparison also highlights an important distinction: strong performance on an operational parameter-prediction task does not by itself demonstrate reliable aircraft gearbox-fault diagnosis.

Overall, the literature indicates that no single AI approach is universally optimal. Traditional machine-learning models can provide strong performance with structured features and comparatively straightforward computational requirements, while deep-learning approaches offer automated feature extraction and greater capacity to model complex sequential or signal-based patterns. Transfer learning can partially address the limited availability of labeled fault data, but performance may still decrease when training data are scarce or when models are transferred across substantially different operating conditions .1 These limitations are particularly important in aviation, where models must generalize beyond controlled datasets and operate reliably across changing flight conditions.

However, the AI approaches discussed in this section largely depend on the information available from individual sensor streams or structured feature sets. The next section therefore examines multi-sensor data fusion and how combining vibration, temperature, rotational speed, acoustic, and operational flight data can provide a more comprehensive representation of aircraft gearbox condition than single sensor approaches alone.

3. Multi-Sensor Data Fusion for Aircraft Gearbox Diagnosis

Section 2 established that AI models can successfully diagnose mechanical faults from individual sensor signals when appropriate features and training data are used. However, relying on a single sensor limits the amount of information available to the diagnostic system. Aircraft gearboxes operate under changing loads, speeds, temperatures, and mechanical interactions, meaning that a fault can produce several physical symptoms at the same time. Vibration may capture impacts or changes in gear-mesh behavior, temperature may reveal friction and lubrication problems, acoustic measurements may detect changes in emitted noise, and operational variables such as rotational speed and torque provide information about the conditions under which the gearbox is operating .[14] Multi-sensor data fusion addresses this limitation by combining information from several sensor streams so that the AI model can use complementary evidence rather than depending on one signal alone .[14]

Limitations of Single-Sensor Gearbox Monitoring

Table 2. summarizes the main limitations of single-sensor monitoring compared with multisensor monitoring for gearbox fault diagnosis. Single-sensor systems provide useful information but are more dependent on sensor location and may fail to capture certain fault types or distinguish between simultaneous faults. In contrast, multi-sensor monitoring combines complementary measurements, providing a more complete representation of gearbox condition. This allows AI-based models to better identify different failure mechanisms, reduce uncertainty, and improve the reliability of fault diagnosis.

Aspect	Single-Sensor Monitoring	Multi-Sensor Monitoring
Information captured	Captures only one type of signal	Captures multiple types of signals
Dependence on location	Highly dependent on sensor location and orientation	Less dependent; different sensors compensate for location effects
Detection of different fault types	Some faults may not be clearly detected	Different sensors detect different fault mechanisms
Multiple simultaneous faults	Signals can overlap and cause interference	Easier to separate and identify individual faults
Overall reliability	Lower reliability in complex conditions	Higher reliability and more consistent results
Overall assessment	Partial view of gearbox condition	More complete and accurate assessment

Complementarity of Multiple Sensor Streams

The main advantage of multi-sensor monitoring is sensor complementarity. Different sensors respond to different physical consequences of the same mechanical condition. Instead of asking one sensor to capture every aspect of a fault, a diagnostic system can combine several measurements and allow the AI model to determine which information is most useful for distinguishing between fault classes .[14]

Sensor modality	Direct evidence it captures	Failure mode it cannot see
Vibration (accelerometer)	Bearing wear, gear-tooth cracks, imbalance, shaft misalignment, impact events	Slow thermal drift, long-timescale lubrication loss
Temperature (thermocouple / RTD)	Friction, lubrication loss, thermal loading	Impact-event timing, tooth-mesh condition
Rotational speed (tachometer)	Speed variation, torque ripple, second-order gear harmonics	Fault-severity quantification, fault-type classification
Acoustic emission (microphone)	Broadband noise from crack propagation, bearing-race spalling	Slow thermal evolution, mounting-position bias
Operational flight data	Load, torque, flight regime, engine speed	Internal mechanical signatures hidden in sensor envelope

Figure 4. Complementarity of Multiple Sensor Streams for Aircraft Gearbox Fault Diagnosis.

The important point is that these streams should not be viewed as independent measurements. Their combination provides contextual information about the same physical system. For example, an increase in vibration may have a different diagnostic meaning under

different rotational speeds or thermal conditions. Combining sensor information therefore allows an AI model to evaluate the fault signal together with the operating state in which that signal occurred .[14]

This principle is particularly relevant to aircraft Health and Usage Monitoring Systems (HUMS). Aircraft gearboxes operate under changing mechanical and environmental conditions, so interpreting one measurement without considering the other available measurements can reduce diagnostic reliability. A multi-sensor architecture can instead use several measurements simultaneously, potentially reducing the uncertainty associated with any individual sensor .[14]

Strategies for Multi-Sensor Data Fusion

Once several sensor streams are available, the information can be combined at different stages of the AI pipeline. Three commonly discussed approaches are early fusion, feature fusion, and decision fusion. The main difference between them is the point at which information from the sensors is combined .[14]

Early fusion: Combines features from multiple sensors into one input before classification so the model can learn relationships between sensor data.

Feature fusion: Extracts features from each sensor separately and combines the learned representations before classification.

Decision fusion: Combines the predictions of separate sensor-specific models to produce one final diagnosis.

For the aircraft-gearbox framework considered in this review, early fusion is the most directly supported approach by the experimental evidence discussed below. Habbouche et al.14 specifically demonstrated that concatenating features from multiple sensor channels before classification can substantially improve gearbox fault-diagnosis accuracy. Early fusion is therefore used here as the main example of how multi-sensor information can improve AI-based gearbox diagnosis.

Experimental Evidence for Early Fusion

One of the clearest experimental demonstrations of the value of multi-sensor fusion comes from Habbouche et al.,14 who investigated gearbox failure diagnosis using multiple microphones. Their experimental system used an HTM90 gearbox-and-bearing test bench with three microphones positioned in different orientations. The three microphones represented different measurement perspectives of the same gearbox, allowing the researchers to investigate whether combining their information could improve fault classification .[14]

The experiment considered 12 failure classes produced by combinations of gearbox and bearing conditions. The signals were sampled at 25.6 kHz, and the recordings were divided into shorter segments for classification .14 Each microphone was processed independently using linear predictive coding (LPC), with the first 15 LPC coefficients used as features. These features were then supplied to an LSTM-based classifier consisting of a 100-node LSTM layer, a 10-node fully connected layer, and a softmax classification layer .[14]

Method	1st microphone	2nd microphone	3rd microphone	Early fusion (3 fused)
LPC–LSTM (per stream vs. fused)	89.58 %	90.28 %	88.89 %	**100 %**
MFCC–CNN–LSTM (per-stream alt.)	97.9 %	98.3 %	100 %	— (single-stream)

Figure 5. Classification Accuracy of Individual and Early-Fused Microphone Streams.

The individual microphones achieved 89.58%, 90.28%, and 88.89% accuracy, showing that each sensor provided useful but incomplete diagnostic information .[14] Because the microphones had different orientations, they did not make the same classification errors, meaning that information unclear in one sensor could be clearer in another. To exploit this complementarity, the researchers combined the 15 LPC coefficients from each microphone into a single 45- coefficient input vector for the same LPC-LSTM classifier. This early-fusion approach achieved 100% accuracy across all 12 failure classes, compared with less than 91% for the individual microphones .[14] This represents a 9.72-percentage-point improvement over the best individual microphone and demonstrates that combining complementary sensor information can significantly improve fault diagnosis.

Implications for Aircraft Gearbox HUMS

Multi-sensor fusion improves fault classification by combining different measurements of the same mechanical event. Habbouche et al.[14] showed that sensor diversity can improve diagnostic performance, supporting its use in aircraft gearbox monitoring. However, their 100% accuracy was achieved in a controlled test bench and should not be assumed for real aircraft conditions. Real aircraft systems face varying speed, load, temperature, vibration, and configuration, making representative training data essential .9 multi-sensor fusion provides more physical information while reducing dependence on any single sensor affected by noise, position, or operating conditions. Nevertheless, additional sensors increase data complexity, requiring a balance between accuracy, computational cost, reliability, synchronization, and training data. Section 4 therefore examines explainable AI (XAI), focusing on how the reasoning of AI-based gearbox diagnosis can be made understandable and trustworthy to human users.

4. Explainability Challenges in AI-Based Predictive Maintenance

Sections 2 and 3 demonstrated that AI can effectively learn gearbox fault patterns and fuse information from multiple sensors. However, high classification accuracy is not enough for aircraft maintenance. In safety-critical aviation environments, engineers and regulators must also understand why an AI system predicted a particular fault. AI models are often described as black

boxes because their internal decision-making is difficult for humans to interpret, which can reduce trust in their predictions .[15] Therefore, explainable artificial intelligence (XAI) is necessary to connect accurate AI predictions with understandable engineering evidence.

The XAI Taxonomy

XAI methods can be organized along three main dimensions: global versus local explanations, model-specific versus model-agnostic methods, and intrinsic versus post-hoc explanations .[15] Global explanations describe overall model behavior, while local explanations focus on why a specific prediction was made. Model-specific methods are designed for particular model architectures, whereas model-agnostic methods can be applied across different models. Intrinsic methods are interpretable by design, while post-hoc methods explain predictions after the model have produced them .[15]

Table 3. Explainability categories and limitations.

Axis	Categories	What it discriminates	Limitations affecting explainability
Source of explainability	Intrinsic interpretability vs post-hoc explanation	Models built to be interpretable on their own (e.g. linear regression, decision trees) versus models explained by an external interpretive layer applied to a black-box prediction	• Class imbalance due to rare gearbox failures can bias explanations toward majority (healthy) classes (Ali et al., 2019). • Data balancing methods (e.g. SMOTE, TabPFN) are needed but not always sufficient for reliable, interpretable insights.
Scope of explanation	Global vs local	Global explanation describes the model's overall behaviour averaged across many forecasts; local explanation describes the model's behaviour on a single forecast	• Laboratory-to-fleet transfer gap: models trained on controlled lab data may not generalize to varying real-world operating conditions (Magadán et al., 2023). • Explanations may lose validity when the model encounters conditions outside the training distribution.
Dependency on the forecasting model	Model-agnostic vs model-specific	Model-agnostic techniques can be applied to any AI model; model-specific techniques can only be applied to models built with a particular algorithm	• Explanations are only as reliable as the underlying model and data quality. • An explanation of an unreliable or out-of-distribution prediction cannot automatically be considered a physically meaningful explanation (Magadán et al., 2023).

For aircraft gearbox diagnosis, these approaches can help engineers identify which sensor measurements or signal characteristics influenced a predicted fault. For example, LIME can provide a local approximation of a model's decision, while SHAP can quantify the contribution of individual features to a prediction .15 In a multi-sensor system, such explanations could show whether vibration, acoustic, temperature, or other measurements contributed most strongly to a diagnosis.

Explainability Requirements in Aviation

These limitations make explainability particularly important for the three main users of aviation predictive-maintenance systems: maintenance engineers, regulators, and aircraft operators. Engineers need explanations to support diagnosis, regulators require traceable evidence behind maintenance decisions, and operators need confidence that predictions are supported by reliable sensor information (Kabashkin et al., 2025).

An effective XAI system should therefore provide more than a fault label. It should connect the predicted fault with its confidence, the sensor evidence influencing the prediction, and the relevant features or signal patterns. This creates a traceable process from sensor data to AI prediction and, through XAI, to the engineering decision. AI should support rather than replace maintenance personnel by helping engineers verify diagnoses and determine appropriate corrective actions (Kabashkin et al., 2025).

From XAI to Generative AI

Although XAI explains why a model produced a prediction, it does not necessarily turn that explanation into a complete engineering recommendation. For example, a classifier may identify a bearing fault while XAI shows that vibration features strongly influenced the prediction, but an engineer must still interpret what this evidence means for maintenance. This creates the gap addressed in Section 5. The AI system developed across Sections 2–4 can produce a fault classification, confidence value, and sensor-level explanation, but these outputs may remain technical and fragmented. Generative AI can build on these results by converting structured predictions and XAI outputs into a human-readable engineering narrative.

5. Generative AI Integration for Predictive Maintenance Decision Support

Previous sections established that AI could diagnose aircraft gearbox faults and that multisensor fusion can improve diagnosis by combining complementary sensing information. XAI methods such as SHAP, LIME, and Layer-wise Relevance Propagation (LRP) can identify features contributing to a prediction, but these outputs still require human interpretation .[15] GenAI can therefore act as a communication layer that translates these technical results into understandable engineering explanations.

For example, if an AI classifier predicts a bearing fault and SHAP identifies vibration features as major contributors, GenAI can combine this evidence with sensor measurements and operating conditions to explain the prediction in an engineering context. It can also combine predictive results, confidence levels, XAI outputs, sensor trends, and maintenance records into a structured technical summary, helping engineers review the suspected fault, supporting evidence, and areas requiring further inspection .[15]

However, GenAI must operate under strict reliability and safety constraints because it can generate information that is not supported by evidence. The system should therefore be grounded in validated AI predictions, XAI results, sensor data, maintenance history, and approved technical documentation. GenAI should communicate and contextualise evidence rather than independently determine maintenance actions, with final decisions remaining with qualified maintenance personnel (Kabashkin et al., 2025).

Overall, the proposed framework gives each technology a complementary role: predictive AI identifies the likely fault, XAI identifies the supporting evidence, and GenAI explains and contextualises that evidence for maintenance personnel. This provides the foundation for the hybrid AI–GenAI framework presented in Section 6.

6. Proposed Hybrid AI–Generative AI Framework for Aircraft Gearbox Predictive Maintenance

The literature reviewed in the previous sections demonstrates that AI-based predictive maintenance can improve aircraft gearbox fault diagnosis, while multi-sensor fusion provides complementary information that may improve the reliability of these predictions .[14,1] However, accurate prediction alone is not sufficient for practical aviation maintenance because engineers must also understand the evidence supporting an AI prediction and determine the appropriate maintenance response .[15] Based on these findings, this review proposes a hybrid Predictive AI + Generative AI framework that integrates multi-sensor monitoring, predictive AI, explainability, and human decision support.

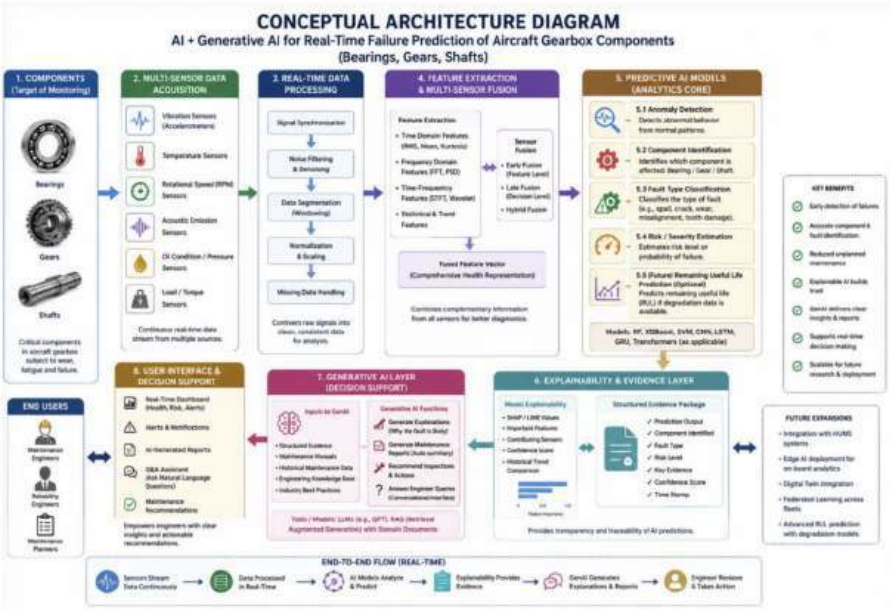

Figure 6. Proposed Architecture Diagram.

6.1 Proposed System Architecture The proposed architecture begins with continuous monitoring of critical aircraft gearbox components, specifically bearings, gears, and shafts. Multiple sensor types, including vibration, temperature, rotational speed, acoustic emission, oil condition, pressure, torque, and load sensors, provide complementary information about component condition and operating state .[2,1]

The collected signals are then synchronized, filtered, segmented, normalized, and prepared for analysis. The processed data are transformed into time-domain, frequency-domain, time-frequency, and trend-based features before being combined through a multi-sensor fusion stage. This follows the evidence presented in Section 3, where combining information from multiple sensors improved gearbox fault classification compared with individual sensor measurements .[14] The resulting fused representation provides the Predictive AI models with a more comprehensive description of gearbox condition.

Explainability and Generative AI Integration

After the Predictive AI models generate a diagnosis, the proposed framework introduces an Explainability and Evidence layer. XAI methods such as SHAP and LIME can identify influential features, contributing sensors, confidence values, and other evidence associated with prediction . [15] This creates a traceable connection between the original sensor data and the predicted component or fault.

Generative AI is then positioned as a communication and decision-support layer rather than as the primary diagnostic model. It can use structured prediction, XAI evidence, historical maintenance information, and approved technical knowledge to generate understandable explanations and maintenance reports .15 This allows engineers to interact with the system using natural language while keeping the underlying diagnosis based on the validated Predictive AI models.

Real-Time Operating Cycle and Human Oversight

The complete system is designed to operate as a near-real-time cycle: sensor data are collected, processed, fused, analyzed by Predictive AI, interpreted through the explainability layer, and communicated through Generative AI. The resulting information is presented to qualified maintenance personnel through alerts, dashboards, reports, and decision-support tools .2

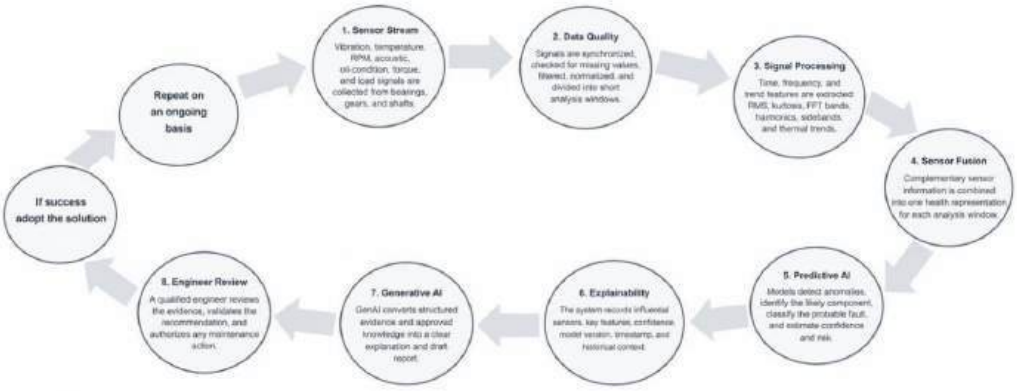

Figure 7. Near-Real-Time AI–GenAI Operating Cycle with Human Oversight. Figure 7 shows that AI provides evidence and recommendations that must be reviewed and validated by qualified personnel. This human-in-the-loop approach is particularly important for safety-critical aviation applications (Kabashkin et al., 2025).

Research Contribution and Limitations

The main contribution of the proposed framework shifts the focus of AI-assisted maintenance toward certification grounding: integrating multi-sensor fusion and Predictive AI with an Explainability and Generative AI layer designed to respect strict aviation compliance, airworthiness standards, and approved maintenance documentation. Multi-sensor fusion provides complementary physical information, Predictive AI identifies probable faults, XAI provides

evidence for those predictions, and GenAI converts the resulting information into understandable decision support .[14,15]

However, the framework remains conceptual and has not yet been experimentally validated as an integrated aircraft system. The strong results reported in controlled gearbox experiments cannot automatically be assumed to represent real aircraft performance, where operating conditions, sensor behavior, and available fault data can differ substantially .14,1 Future research should therefore evaluate the proposed architecture using representative multi-sensor datasets and assess diagnostic accuracy, generalization, explanation reliability, GenAI grounding, computational requirements, and engineer decision quality.

Overall, the proposed framework provides a pathway from continuous sensor monitoring to AI-based diagnosis, explainable evidence, Generative AI communication, and human maintenance decision-making. It therefore serves as a conceptual foundation for future research into scalable AI-assisted predictive maintenance for aircraft gearbox bearings, gears, and shafts.

Conclusion

AI-based predictive maintenance can improve aircraft gearbox monitoring by analyzing sensor data to detect faults before they become severe. The literature shows that machine-learning and deep-learning models can successfully diagnose mechanical faults, but their performance depends on representative training data and may decrease when moving from controlled laboratory datasets to real aircraft conditions .[9,1]

Multi-sensor fusion can improve diagnosis by combining complementary information from vibration, temperature, rotational speed, acoustic, and operational measurements. The reviewed evidence shows that combining sensors can provide better fault classification than relying on a single sensor .[14] However, accurate predictions must also be understandable to maintenance personnel. XAI can provide evidence for AI predictions, while Generative AI can translate that evidence into clear maintenance explanations and reports.[15]

Overall, this paper proposes a conceptual hybrid architecture for monitoring aircraft gearbox bearings, gears, and shafts. The framework has not yet been experimentally validated as a complete aircraft system, so future research should test it using representative multi-sensor datasets and evaluate its accuracy, generalization, explainability, and reliability. The proposed approach provides a foundation for future AI-assisted, human-centered predictive maintenance in aviation.

Acknowledgements

I would like to express my sincere gratitude to my Indigo Research professors, Dr. Siddharth Krishnan and Dr. Samuel Lefcourt, for their invaluable guidance, support, and expertise throughout the development of this research. I would also like to extend my appreciation to my peers for their constructive feedback, thoughtful suggestions, and encouragement, which contributed meaningfully to the improvement of this work.

Code and Data

Availability This review did not develop or execute original software code and did not generate a new dataset. The studies and public datasets discussed in the review are identified in the cited sources. Any future experimental implementation of the proposed framework should provide persistent links to the code, trained models, configuration files, and data used for evaluation.

AI Acknowledgement

During the preparation of this work, the author used Microsoft Copilot to generate the Figure 6 - Proposed Architecture Diagram

Works Cited

Ali, M. A., & Al-Saadi, A. (2019). The influence of handling imbalance classes on the classification of mechanical faults using neural networks. 2019 8th International Conference on Modeling Simulation and Applied Optimization (ICMSAO), 1–5. https://doi.org/10.1109/ICMSAO.2019.8880437

Accident Investigation Board Norway (NSIA). (n.d.). Report on the air accident near Turøy, Øygarden municipality, Hordaland county, Norway 29 April 2016 with Airbus Helicopters EC 225 LP, LN-OJF, operated by CHC Helikopter Service AS. https://nsia.no/Aviation/Aviation/Published-reports/2018-04

Cao, P., Zhang, S., & Tang, J. (2017). Pre-processing-free gear fault diagnosis using small datasets with deep convolutional neural network-based transfer learning. arXiv. https://arxiv.org/abs/1710.08904

Cheng, X., Dou, S., Du, Y., et al. (2024). Gearbox fault diagnosis method based on lightweight channel attention mechanism and transfer learning. Scientific Reports, 14, Article 743. https://doi.org/10.1038/s41598-023-50826-6

Dabhi, L., et al. (2026). Data-driven estimation of helicopter engine power using regular flight data: A machine learning approach. Electronics, 15(1), Article 141. https://doi.org/10.3390/electronics15010141

Habbouche, H., Benkedjouh, T., Amirat, Y., & Benbouzid, M. (2021). Gearbox failure diagnosis using a multisensor data-fusion machine-learning-based approach. Entropy, 23(6), Article 697. https://doi.org/10.3390/e23060697

Hünemohr, D., et al. (2022). Usage monitoring of helicopter gearboxes with ADS-B flight data. Aerospace, 9(11), Article 647. https://doi.org/10.3390/aerospace9110647

Ignjatovska, A., et al. (2023). Classification of present faults in rotating machinery based on time and frequency domain feature extraction. Vibroengineering Procedia, 51, 22–28. https://doi.org/10.21595/vp.2023.23667

Kabashkin, I., et al. (2025). AI-driven fault detection and maintenance optimization for aviation technical support systems. Processes, 13(3), Article 666. https://doi.org/10.3390/pr13030666

Li, Y., & Wang, K. (2020). Modified convolutional neural network with global average pooling for intelligent fault diagnosis of industrial gearbox. Eksploatacja i Niezawodność, 22(1), 63–72. https://doi.org/10.17531/ein.2020.1.8

Loewenson, I. (2023, July 21). Clutch problem in Osprey led to death of 5 Marines in 2022, Corps says. Marine Corps Times. https://www.marinecorpstimes.com/news/your-marine-corps/2023/07/21/clutchproblem-in-osprey-led-to-death-of-5-marines-in-2022-corps-says/

Magadán, C., et al. (2023). Early fault classification in rotating machinery with limited data using TabPFN. IEEE Sensors Journal, 23(24), 30960–30970. https://doi.org/10.1109/JSEN.2023.3331100

Mallikarjuna, P. B., Sreenatha, M., Manjunath, M., & Kundur, N. C. (2020). Aircraft gearbox fault diagnosis system: An approach based on deep learning techniques. Journal of Intelligent Systems, 30(1), 258–272. https://doi.org/10.1515/jisys-2019-0237

Molano, C. C. (2019). Analysis of NASA bearing dataset of the University of Cincinnati by means of Hjorth's parameters. In the International Conference on Engineering Dynamics (ICEDyn). https://iris.unimore.it/bitstream/11380/1203704/1/Final%20paper%20ICEDyn%2 02019.pdf

Moreira, G., et al. (2024). Addressing gearbox health monitoring challenges for helicopters: A machine learning approach. Anais da Academia Brasileira de Ciências, 96, e20240404. https://doi.org/10.1590/0001-3765202420240404

Sofianidis, G., Rožanec, J. M., Mladenić, D., & Kyriazis, D. (2021). A review of explainable artificial intelligence in manufacturing [Preprint]. arXiv. https://arxiv.org/abs/2107.02295

Spadini, T., et al. (2024). Intelligent fault diagnosis of type and severity in low-frequency, low bit-depth signals [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2411.06299

Swati, K. (2026, July 18). Airbus sees the end of the A220 engine crisis by late 2026 as Pratt improves GTF reliability. MRO Business Today. https://mrobusinesstoday.com/airbus -sees-end -of-a220-engine-crisis-by-late2026-as-pratt-improves-gtf-reliability/

The Associated Press. (2024, August 2). The Osprey crash was caused by the pilot's decision to keep flying with cracks in a gear. NPR. https://www.npr.org/2024/08/02/nx-s1-5061194/ osprey-crash-was-caused-bypilots-decision-to-keep-flying-with-cracks-in-a-gear

Wu, Z., et al. (2024). Robust health condition prediction of helicopter turboshaft engines using ensemble machine learning models. Annual Conference of the PHM Society, 16(1). https://doi.org/10.36001/phmconf.2024.v16i1.4195

Xu, K., et al. (2024). Research on temperature rise characteristics prediction of main shaft dualrotor rolling bearings in aircraft engines. Lubricants, 12(6), Article 210. https://doi.org/10.3390/lubricants12060210

Zhang, S., Zhang, S., Wang, B., & Habetler, T. G. (2019). Machine learning and deep learning algorithms for bearing fault diagnostics—A comprehensive review. arXiv. https://arxiv.org/abs/1901.08247

Zhao, Z., et al. (2023). A fault diagnosis method of rotor systems based on parallel convolutional neural network architecture with attention mechanism. Journal of Signal Processing Systems, 95(8), 965–977. https://doi.org/10.1007/s11265-023-01846-y

Zhou, L., et al. (2018). Using frequency domain analysis techniques for diagnosis of planetary bearing defects in a CH-46E helicopter aft gearbox. Engineering Failure Analysis, 92, 71–83. https://doi.org/10.1016/j.engfailanal.2018.04.051

Author Vikrant Raj Sathyamurthy is a rising senior at PrepaTec (Tecnológico de Monterrey) Multicultural High School in Guadalajara, Mexico, conducting research in Mechanical Engineering and Artificial Intelligence under the mentorship of Dr. Siddharth Krishnan and Dr. Samuel Lefcourt through the Indigo Research Program. His project explores a hybrid AI and physics-based approach for predictive maintenance of aircraft gearboxes. He was drawn to this topic by his passion for AI, engineering, and aviation safety. What excited him most during the research process was seeing how recent advances in Generative AI could be leveraged alongside traditional AI and physics-based models to enhance predictive maintenance, opening new possibilities for smarter and more reliable aerospace systems.

The Architecture of Exception: What Legal Exclusions in Singapore Reveal About Global Migration Governance By Ava Rose Attari

Research Question: What does the exclusion of domestic and maritime labour in Singapore's Employment Act tell us about the structural power asymmetries of global migration governance?

Abstract

Expatriates contribute nearly 40% of Singapore's population, forming the economic backbone of its development. However, two key sectors remain governed by frameworks that exclude them from key legal protections. This paper investigates a critical question within global migration governance: what does the exclusion of domestic and maritime labour in Singapore's Employment Act reveal about structural power asymmetries? I evaluate how statutory exclusions are not accidental policy oversights, but a calculated strategy that institutionalised state control to maximise capital output. By evaluating live-in domestic workers and maritime crew, I scrutinise labour legislation alongside international standards. Applying theoretical frameworks of isolated precarity, state exception, and fragmented global governance, and calibrating them to a Singaporean context, this paper showcases how marginalisation and overwork are produced. Overall, through qualitative textual analysis, the paper proposes the consequences of statutory exclusion, while exploring how origin countries are burdened by remittance dependencies – lacking the diplomatic leverage to advocate for their citizens abroad.

Keywords Employment Act, Migrant Domestic Workers, Maritime Labour, Spatial Control, State Power Asymmetries

Introduction

Singapore remains one of the world's most affluent countries, accumulating a gross domestic product exceeding $650 billion in 2026 (Singapore GDP (2026)). However, the country is facing a declining fertility rate, which 'Worldometer' reports has dropped from 1.78 in the 1980s to 1.26 in 2026. As a result of this, alongside Singapore's rapid economic expansion, the state has become fundamentally reliant on importing an immense number of low-wage foreign labour. The transient workforce is primarily drawn from low-income countries (LICs) across South/Southeast Asia, primarily Indonesia, the Philippines, Myanmar, Bangladesh, and India. (Chin, 2019). It is evident that Singapore hosts over 1.2 million foreign workers on temporary permits. Strikingly, there are more than 350,000 workers who are concentrated in two essential sectors: modern domestic workers (MDWs), contributing towards the nation's social reproduction, and maritime labourers, driving global trade (Foreign Workforce Numbers, Ministry of Manpower).

The management of labour within Singapore's population is underpinned by a strict legal architecture, tailored to ensure that foreign workers remain strictly temporary. The Work Permit system ensures that low-skilled migrant workers are bound to their employer, barring them from

bringing family members in, marrying any local citizen, and even denying them permanent residency. The policy regime deliberately constructs foreign labour as an economic input rather than members of society. However, the most severe factor of vulnerability lies in Singapore's primary labour statute: The Employment Act. While this act does establish a baseline protection for most employees, domestic and maritime labourers are explicitly excluded from its main provisions. Despite the labourers' vital role in sustaining Singapore's economy, they are systematically denied fundamental labour protections; both MDWs and maritime labourers are excluded from protections regarding maximum working hours, mandatory rest days, dispute resolution mechanisms, and overtime pay. (Chin, 2019)

This article focuses on Singapore as a case study because it is exceptionally revealing for examining the disempowerment of migrant labour. Its contemporary labour regime is deeply rooted in British colonial legacies. Established first in 1819 as an "imperial entrepot", Singapore's economic model was built up through importing vast numbers of manual labourers from LIC's such as China, India and the Malay Archipelago to service trade ports - without giving them the right of residency. After gaining its independence, Singapore inherited this imperial logic through modern legal mechanisms such as transforming colonial labour practices into a contemporary work permit structure (Kaur, 2007). The current exclusion of both domestic and maritime workers from statutory protections reflects a direct continuation of this colonial legacy. Specific classes of labourers are governed not through a standard measure of employment law, but through administrative controls specifically designed to maximise economic growth. By evaluating Singapore as a case study, I reveal how a postcolonial buildup can retain imperial governance mechanisms to maintain structural power over a migrant workforce in a modern global economy.

This paper argues that the inherent exclusion of domestic and maritime labour from the Employment Act is not an accidental oversight, but rather a deliberate strategy implemented by the state. It is evident that the explicit exclusion of these two labour branches is due to the Singaporean state's inherent effort to preserve its maximum legislative flexibility, minimising regulatory costs for family homes and maritime firms. This also ensures a continuous supply of inexpensive, exploitable workers. Driven by geopolitical power imbalances, source countries of economic migrants are trapped by their reliance on foreign remittances, leaving them lacking the structural power that's required to protect their citizens abroad. Overall, this constructs a localised regime which is defined by social isolation and severe overwork - directly mirroring the unequal interactions between source and host states on the international stage. As a result, these workers completely lack "adequate labour protection", creating a localised power imbalance favouring employers. Overall, to adverse working and living conditions for MDWs and maritime labour. (Parreñas et al., 2021).

To unpack this argument, this paper is structured into three interconnected sections: Section 1 aims to highlight the theoretical background for the paper by evaluating mainstream migration networks, specifically Castles' (2010) model of labour commodification, Ruhs (2013) analysis of visa precarity, and the concept of "spaces of exception", cultivated by Agamben in

1998. These arguments have historically been analysed through a Euro-American lens. Therefore, this section recalibrates these theories towards a Singaporean context. Additionally, section 2 strives to broaden the macro-geopolitical scale, locating Singapore's domestic policies within the wider context of global migration governance. Drawing on theories of fragmented governance and structural power, it also takes a deep dive into the colonial legacies of Singapore. Section 3 moves away from the macro scale towards an evaluation of how Singapore's statutory exclusions under the Employment Act operate in practice. It emphasises how the state works alongside employment recruitment agencies and private employers to form a mutually reinforcing network of control. Finally, section 4 aims to examine the true lived realities of migrant workers on the ground and at sea. Through analysing their physical environment - specifically the private domestic household and the isolated vessel. This exemplifies how the exclusion of these workers from the Employment Act can not only restrict them from day-to-day activities but also have a direct effect on their mental health.

Literature Review

Literature on migrant labour in Singapore distinctly investigates the daily lived realities of low-wage foreign workers within both the private household and maritime domains. Existing scholars demonstrate how structural isolation is operationalised on a micro level, particularly through physical isolation and psychological control. Firstly, Anjara et al. (2017) scrutinise the severe mental health issues that affect MDWs, showcasing that their prolonged exposure to severe employer surveillance and overwork directly intensifies rates of depression and distress. To expand on this idea, Goh (2019) is utilised to examine how the private household can function as a site of compulsory social isolation, with its lack of privacy also functioning to strip domestic workers of personal autonomy. Although these studies highlight terrestrial care labour, Yea (2022) extends this analysis towards the maritime domain. It exemplifies how fishing vessels operate as extremely isolated "spaces of exception" - seafarers endure extreme physical confinement, withholding of wages, and abuse -- which are far from regulatory oversight. For this paper, these specific papers contribute useful evidence for how intense isolation and overwork affect workers on the ground. However, many pieces of localised literature frame the conditions as personal employer-to-employee dynamics. This piece of research aims to bridge the analytical gap by linking the micro-level experiences directly to the state-level exclusions under the Employment Act, showcasing that daily control is not only a private workplace dynamic, but an intended output of a state legal regime.

To further expand, global migration has been analysed to highlight structural asymmetry and the power of the state. These include Castle's (2010) commodification model and Ruh's (2010) analysis of how host states trade off workers for a high import volume. It can also be noted that Agamben's (1998) concept of "spaces of exception" has investigated how these workers are stripped of their political agency and mobility. This is due to the "spaces" being private and isolated - authorities are unable to conduct routine inspections of workers' well-being and treatment. Many scholars' studying international political economy have investigated how

host states utilize restrictive visas, e.g., single employer sponsorship. This constructs a guest worker regime, stripping foreign labour of social and political autonomy. However, a significant limitation within popular global migration studies is the persistent Western (Euro-American) bias; theories regarding borders, commodification of labour, and state sovereignty remain focused on HICs - such as the United Kingdom and the United States. Because of this, traditional scholarship regarding migration constantly overlooks the unique governance of economic powerhouses in the Global South, especially Singapore. Crucially, these localized conditions do not exist in a historical vacuum; they reflect the enduring imprint of British colonial labour management, which historically isolated unfree migrant workforces to maximize economic extraction.

Although authors such as Chin (1998) discuss the early analysis of labour regimes within Asia - exposing how post-colonial countries create "complicit networks" with private employment agencies and isolating female MDWs as temporary commodities to fill state reproduction goals - mainstream global governance literature has largely overlooked these regional dynamics. To expand on this, Kainz & Betts (2021) highlight why our global system remains inherently fractured, arguing that economically powerful labour host states deliberately reject universal frameworks implemented by the United Nations. Instead, they opt for non-binding agreements - receiving states maintain their total sovereignty over migrant workers. To complement this, Likić-Brborić (2018) scrutinises the political economy, emphasising that the current system creates an "asymmetric dual regime."

Within this structure, wealthy states prioritise cheap labour flexibility, marginalising rights-based frameworks and international labour conventions. It could also be highlighted that Wickramesekara (2015) examines the bilateral mechanisms that execute this inequality on the ground; his work highlights that source countries are pressured into signing weak, non-binding 'Memoranda of Understanding' rather than actual enforceable Bilateral Labour Agreements. Due to these source countries being economically dependent on worker remittances, such as India, where 3.8% of its GDP is made up of remittances (World Bank, n.d.). Host states can dictate the terms of these agreements, ensuring that standard labour protections (such as minimum working hours) are stripped before the work even commences. Through analysing these theories together, this paper recalibrates the theoretical view. Foundational global migration theories will be applied and recalibrated towards a Singaporean context, allowing their findings to fit within an economic hub in the global south. By taking into consideration the legislation and social norms, this allows the thematic analysis to be as precise as possible for the case study. These theories will be paired with micro-scale realities of migrant workers, showcasing how the lived realities of both MDWs and maritime crew act as a direct manifestation of structural power asymmetries embedded within global migration governance. Overall, this paper demonstrates that the exclusion of domestic and maritime labour within Singapore's primary labour statute is not an isolated local policy, but rather a direct evolution of its British colonial heritage, modernising its imperial mechanisms of labour in an attempt to exploit contemporary global power asymmetries between host and source states on the international stage.

Ultimately, this paper fills a major gap within existing literature on global migration governance - synthesising analytical studies into a cohesive framework. These studies do highlight the structure of the global political economy - how it functions, the inequalities endured within high-income and low-income countries, and the micro-living realities. Very few studies explicitly connect both sides of the spectrum. By examining how the spatial isolation and legal exclusions under Singapore's Employment Act are an input of wider geopolitical statecraft, this study exemplifies a deep analysis that examines how daily worker exploitation can be embedded within the architecture of global migration governance.

Section 1: Theoretical Background: Recalibrating Modern Migration Theories Towards Singapore

To emphasise how host countries completely disempower migrant workers, this article employs a theoretical framework that is grounded in political economy and spatial control. It can be highlighted that the structural vulnerability of foreign migrant workers is not a byproduct of migration, but the deliberate result of state-driven labour commodification. A method utilised by host countries to completely strip human beings of their civil, political and social personhood, reducing them to a mere economic input (Castles 2010). Throughout this process, human rights ensure that foreign workers are granted conditional access to territory to fulfil economic demands but are completely insulated from accessing permanent residency or political representation. As a result of this structure, host countries treat these labourers as an economic input/disposable commodity which can be discarded during economic downturns without the state worrying about long-term settlement (Castles, 2010, 2013).

It is evident that host states execute a calculated 'trade-off' between labour volume and human rights to operationalise this commodification without triggering any domestic political backlash or wage inflation (Ruhs, 2013). This strategy relies immensely on employer-sponsored visa architectures. The worker is essentially tied to their employer. By legally binding a worker's right to remain in a country to the sponsorship of a single employer, the state effectively delegates a sovereign disciplinary power to private actors. This specific mechanism completely marginalises workers' labour market mobility and choice - leaving an abusive employer can result in immediate visa cancellation and compulsory deportation. Overall, this allows the threat of deportation to remain a constant instrument of workplace exploitation. It can also be highlighted that this visa precarity creates hyper-dependence: the employer is granted complete administrative control over the worker, transforming the threat of unilateral contract termination/deportation into a tool of workplace discipline.

In recent years, advocates of Singapore's model argue that statutory exclusions and strict visa controls act as a pragmatic policy trade-off- maintaining high labour volumes and driving remittance flows towards source countries (Ruhs, 2013). This perspective could also highlight that denying labour rights is a necessary compromise for national economic competitiveness. However, this logic conceals a dramatic power asymmetry: using economic efficiency to legitimize structural devaluation and hyper-exploitation of these foreign workers. In essence,

these policies can be highlighted as ‘jurisdictional spaces of exception’ (Agamben, 1998; Yea, 2022). As private households and maritime vessels have been deliberately carved out from standard oversight, the state has created a physical ‘micro zone’ where standard labour protections are suspended. Throughout these isolated domains, authority is outsourced directly to private employers’/vessel captains, transforming these private spaces into a site of unfree labour and complete institutional control. This encourages exploitation of workers, resulting in long working hours, unfair living conditions, and deteriorating mental health. Through expanding on Agamben’s framework, this spatial exception reduces the migrant worker to the status of a ‘homo sacer’-a person who is legal but situated in a site of double exclusion, existing in the physical boundaries of a nationstate but abandoned by its legal protections (Agamben, 1998). Because the state suspends labour laws in the private home and on an isolated vessel, the worker is exposed to harm and systemic marginalisation with legal impunity.

Moreover, the Singaporean regime completely strips the workers' existence to “bare life” - otherwise known as (zoē). Rather than these workers being recognised as residents with political and social rights, or even personal autonomy, they are treated purely as biological units of labour, which need to be isolated and monitored. Additionally, the reduction of these workers to bare life in unmonitored spaces institutionalized violence, leaving workers in long and unbearable working hours and degrading living conditions - overall contributing to the deterioration of their psychological health. While many concepts of labour commodification and spaces of exception do provide crucial analytical tools, their development suffers from a complete Western bias. Formulated through studying democracies in North America and Europe, these frameworks model state power as an incomplete process. In a Western context, if the state attempts to enforce labour commodification, these pieces of legislation are scrutinised and disrupted by independent judicial review and constitutional rights protections such as the European Convention on Human Rights (ECHR), utilised in 46 member states. (Home et al. n.d.) Therefore, applying these theories to Singapore requires complete theoretical recalibration across its core dimensions - how this city-state transforms Western concepts into its own totalised administrative mechanisms:

Singaporean Recalibrations

1. Contested Labour Commodification to Permanently Temporary Status:

It is highlighted that Castles (2010) showcases commodification as a struggle against legal integration; Singapore operationalises complete commodification. The state utilises this by embedding control into law, such as mandatory semi-annual pregnancy and HIV tests (PreEmployment Medical Examination for MDWs, Ministry of Manpower). Overall, commodification is completely recalibrated from a simple market dynamic to a legal and biological boundary.

2. Employer Power to Delegated State Control

Ruh (2013) treats tied visas as a policy 'trade-off' where employers gain leverage. Singapore converts this trade-off into a complete delegation of sovereign power. Therefore, through employer-funded $5,000 security bonds (Security Bond Requirements for MDWs, Ministry of Manpower), private employers are incentivised to act as border wardens - outsourcing immigration enforcement directly into the hands of the private employer.

3. Temporary Border Exception to Normalized Daily Control

Agamben (1998) exemplifies "spaces of exception" that are created as a temporary, peripheral zone during state crises, such as border detention centres. Utilising Agamben's argument for Singapore as a case study highlights that the space of exception is not peripheral or temporary; it is a permanently normalised area that is embedded into the ordinary geography of the country - working routinely inside high-rise residential apartments and commercial maritime vessels.

Western Concept	**Operational Environment in the West**	**Singaporean Recalibration**
Commodification (Castles, 2010)	• Scrutinised by courts, judiciary, unions and naturalisation pathways	• Absolute • Biological (pregnancy bans, mandatory deportation, zero citizenship route)
Tied Visa Precarity (Ruhs, 2013)	• Employer leverage because of labour laws	• Outsourced Sovereignty (Employer security bonds mimic privatised border enforcement)
Spaces of Exception (Agamben, 1998)	• Temporary emergency zones	• Normalised and extremely institutionalised for workers excluded under the Employment Act • Embedded into private homes and isolated maritime ships

Consequently, this paper emphasises Singapore's contemporary guest worker regime as not existing in an institutional vacuum, but a complete continuation of colonial governance. As post-colonial state theory showcases, newly independent states - such as Singapore, which only gained independence in 1965 (Office of the Historian, n.d.) - increasingly function as "complicit networks". Meaning that they actively preserve and modernise imperial governance strategies to service modern capitalist accumulation (Chin, 1998).

The architecture embedded within Singapore of tied work permits, security bonds and labour exemptions from its primary statute represents a direct evolution of British colonial indentured labour systems. These were designed to manage and exploit unfree Chinese and Indian labourers. Through utilising both Western political economy theories and post-colonial state theory, this article aims to establish a multi-scalar framework: the exclusion of domestic and maritime labour under the Employment Act is not an isolated decision, but the downstream

execution of an inherited imperial governance model functioning to increase inequality within an asymmetric global political economy.

Section 2: The Global Matrix: Geopolitical Power Asymmetries and Colonial Inheritance

The vulnerability and exclusions that define the day-to-day lives of foreign MDWs and maritime labourers in Singapore are not isolated anomalies. Instead, the macrorealities of the workers are a complete expression of an asymmetric global migration architecture. To understand the systemic disempowerment of workers within Singapore, its frameworks must be compared to the larger international structures of global migration governance, postcolonial state formation, and the global political economy.

Singapore's modern framework of non-resident labour management is deeply rooted in colonial-era empire governance strategies (Amrith, 2013). Under the rule of the British, Singapore functioned as a free port where infrastructural and commercial expansion was reliant on low-wage indentured labour filled by migrants across Asia. (Huff, 1994). This political economy has established a clear parallel: labour power was imported as a temporary tool of production, insulated from political rights and long-term residency. Rather than dismantling these imperial mechanisms, Singapore's postcolonial state institutionalised them (Chin, 1998). This contemporary framework modernised colonial labour management by substituting imperial administration with an efficient state apparatus, designed to maximise economic output while rigidly enforcing these laws.

Colonial Legacy

It is argued that the institutional continuity between British colonial imperialism and the modern Singaporean state is visible in how policy continues to enforce worker disposability. In the colonial era, migrant networks were categorised as external economic inputs rather than genuine members of society, creating immense legal and social barriers to prevent integration of these migrants (Amrith, 2013).

Instruments utilised in a contemporary context - including the non-transferable work permit system, employer-tied visas and exclusion from the primary employment statute - are a direct parallel to the instruments of population control within the colonial era. The employer dependency allows the state to preserve the foundational colonial divide between a sovereign citizen and a disposable foreign labour force. This is crucial to this paper's argument, as it showcases how the Singaporean state maintained these rigid administrative barriers between membership and non-citizen labour; the postcolonial state extracts minimum economic value from regional source states by insulating domestic and maritime workers from welfare obligations and healthcare liabilities. (Chin, 1998; Castles, 2010). Isolation of these workers is not a policy oversight, but an institutionalised legacy of British colonialism which is designed to treat foreign labour from LICs as temporary, disposable factors of production.

Impact On the Global Stage

This domestic architecture can also be highlighted as reinforcing the structural inequalities that operate on an international level. The complete powerlessness of an individual foreign worker in the Singaporean household or an offshore vessel is a direct parallel to the structural weakness of the source state on the international geopolitical stage. Primary source countries of labour - such as Indonesia, Myanmar and the Philippines - operate under economic vulnerabilities. It can also be highlighted that citizens of Myanmar have extreme difficulty returning home - if deported - due to the civil war, making their dependence on Singapore as a haven even larger. Characterised by persistent fiscal constraints, the structural pressures trap origin states in a state of "remittance dependency" - flows of capital sent home by migrant workers become vital for national economic survival and poverty mitigation (Wickramasekara, 2012; 2015).

Additionally, this dependency exemplifies the diplomatic leverage of receiving governments (Wickramasekara, 2015). It can be highlighted that origin states are effectively forced into a "race to the bottom" against competing labour-exporting countries, securing the host state's employment quotas. If a country attempts to assert its protection over overseas citizens, e.g., enforcing standard employment contracts, wealthy host states can redirect recruitment pipelines towards more compliant states. The individual migrant workers' constant fear of deportation (because of contract termination) is the micro-scale equivalent of their home governments' fear of losing access to extremely useful foreign exchange markers to generate capital.

Within global governance, wealthy host countries continuously leverage their high political and economic rankings to preserve a fragmented migration regime. Unlike global trade or financial schemes, which are constantly regulated by outsider multilateral organisations such as the WTO, global migration governance is decentralised (Kainz and Betts, 2021). Receiving states - Singapore - exploit the institutional fragmentation to bypass international labour conventions, such as the UN International Convention on the Protection of the Rights of All Migrant Workers and Members of their Families (United Nations, 1990). Overall, the central absence of a sovereign governing body allows wealthy host countries to treat internationally recognised labour standards as an optional guideline, rather than a legally binding obligation. The Singaporean State insulates its domestic legal apparatus from external legal scrutiny and global accountability. This includes the exclusion of domestic and maritime labour from the Employment Act. Many Singaporean officials argue that the use of bilateral agreements and non-binding MOUs provides flexible arrangements, securing labour pipelines without compromising their national sovereignty (Wickramasekara, 2015). However, this defence ignores the core power imbalance: without genuine binding international framework benchmarks, bilateral negotiation allows economically dominant host states to dominate terms. This strips origin states of collective bargaining power and further reinforces global inequalities.

Overall, this asymmetric power dynamic within the region is reinforced by host states by favouring bilateral labour agreements and informal Memoranda of Understanding (MOU) over rights-based international frameworks (Wickrameskara, 2015). Bilateral agreements let the host

and source countries negotiate, allowing the wealthy host state to act from a position of overwhelming economic dominance - they dictate employment terms on an individualised basis. Consequently, this approach, which favours bilateral agreements, sends source states into a competitive "race to the bottom". Origin states undercut one another on degrading labour standards to secure their market share to ensure they gain the remittances needed (Wickrameskara, 2012). Due to MOUs being non-binding, they lack the mechanisms needed to hold the host states accountable for exploiting these workers. In addition, these mechanisms externalise social reproduction costs and healthcare liabilities which are associated with foreign labour back towards the origin state, creating a large supply of disposable labour/economic inputs (Castles 2010, 2013). This systemic precarity, which is highlighted through legal exclusions, post-colonial imperialism continuity, and geopolitical power asymmetries, operates to exemplify and maintain the dominance of rich host states.

Section 3: The Legal Side: How Do State Exclusions Work?

The true isolation and distress inflicted on foreign MDWs and maritime crew in Singapore are not accidental policy oversights, but the direct result of a carefully crafted nationwide statute to construct low-wage workers as temporary economic inputs. This demonstrates why they are legally unprotected. Throughout these statutory omissions, Singapore normalises the precarity of these workers. This strategy, enforced by the Singaporean government, demonstrates that worker vulnerability is a calculated structure which insulates private employers from state regulation. Commodifying this legislative silence allows the Singaporean state to normalise worker precarity, transforming disenfranchisement into a routine operational condition of its economy.

<u>Fundamentals of the Regime</u>

At its core, the institutional design is an explicit exclusion of foreign MDWs and seafarers from key provisions of the labour law. Part IV of the Employment Act (1968) - mandating rest days, overtime work pays and leave - excludes these provisions, allowing their working conditions to be set and determined by their employers. This is a calculated strategy by the state; keeping this sector outside standard employment legislation allows private households to stay insulated. Rather than asserting its direct administrative oversight throughout private households and maritime domains, the government withdraws its standard apparatus to insulate these areas from state interference. Embedded throughout the 'Employment of Foreign Manpower Act' (EFMA), the state passes the burden of policing to the employers. Both conditions 4 and 12 of the EFMA Work Permit Conditions illustrate that employers are legally responsible for the "upkeep, maintenance and medical treatment" of their domestic workers. (Ministry of Manpower, n.d.) This sole discretion, which is granted to employers, gives them the authority to abuse their workers freely without any governmental intervention. Additionally, these state guidelines also exempt them from workplace safety and labour inspections under 'The Workplace Safety and Health Act' (WSHA), effectively positioning a private home as a

completely unmonitored zone. Maritime workers also fall into regulatory grey zones where standard protections don't apply. Overall, this system prioritises employer flexibility and low operational costs over basic human rights.

For this regime to work, the state operates within a "complicit network" (Chin, 1998), creating a partnership with both individual employers and employment companies. This network relies on specific legal mechanisms such as the Work Permit system and a mandatory $5000 security bond posted by the employer. Due to the employer risking forfeiting the bond - if a worker absconds - the state effectively gives power to the employer, as they are responsible for policing their workers' behaviour and location. Moreover, recruitment debts (which are very high) are managed by employment brokers who work specifically to immensely restrict the workers' financial ability to leave an environment - even if it is abusive. Because of these mechanisms, the state delegates its surveillance function to private actors, ensuring worker compliance.

This administrative framework is sustained through political logic that characterises low-skilled migrant workers as non-citizen outsiders (Neo, 2015). This has created a legal boundary between citizens - who possess social and political rights - and guest workers - who are denied both. Under this framework, the Singaporean State's economy is explicitly decoupled from long-term integration. Foreign workers are further forbidden from marrying Singaporean citizens or even bringing dependents into the country under their visa. This overall reinforces their social isolation. By legally defining these workers as an economic commodity, the withholding of their protection is legitimised - many believe it protects national sovereignty and preserves economic competitiveness.

To prevent any workers from making permanent social or spatial claims, the legal system enforces strict control over workers' personal lives. These regulations prohibit work permit holders from marrying Singaporean citizens or permanent residency holders. These restrictions are enforced through mandatory biannual medical examinations, which infringe on workers' personal lives, carrying an immediate threat of visa cancellation and deportation. Studies conducted by civil society organisations highlight that these mandatory medical screenings/pregnancy bans incentivise sovereign employers to act as moral custodians. (HOME, 2015, 2019) This leads to routine bodily invasions such as restricted phone access and random bedroom searches to prevent "unauthorised" social contact.

Overall, reducing foreign workers to economic commodities allows the state to legitimize the withholding of statutory protections; structural disenfranchisement is not only framed as an ethical failing but as a necessary and rational policy trade-off. This is required to safeguard national sovereignty and maintain social cohesion. Ultimately, it can be viewed that the Singaporean legal framework serves as the primary mechanism in which worker marginalisation is produced and maintained. By creating these deliberate exceptions and deputising employers, the state constructs loopholes that normalise isolation on the ground. Domestic labour law fails to protect foreign workers, and it intends to maintain a manageable, disposable labour force.

Section 4: Lived Realities: A Regime of Isolation, Overwork, and Spatial Control

The extreme vulnerability and isolation endured by foreign MDWs and maritime crew in Singapore is maintained through the micro-level control of physical space. Throughout private homes and offshore vessels, isolation serves as a tool for extracting labour power while confining worker autonomy. Instead of serving as a site of employment, these environments are legal grey zones where there is no regulatory oversight (Goh, 2019). Through this confinement, spatial control is the mechanism in which foreign labour is commodified and exploited on a daily basis.

It is evident that spatial control is the purpose of the private household and offshore vessel due to sites of jurisdictional exceptionalism (Goh, 2019). Because private homes are legally protected as domestic spaces, they stay immune to mandatory state labour exceptions. This transforms the home into an unregulated place of work - the power of the employer is absolute and sovereign. This parallel dynamic governing maritime workers on Singapore's coast creates a legal vacuum - separating daily labour practices from public and regulatory view.

Within both spheres, physical isolation creates a state of exception where standard employment norms do not apply. This leaves workers reliant on the goodwill of employers. Within these areas, labourers are forced into a regime of overwork and surveillance, dismantling their personal agency. For those in the domestic sector, the lack of defined working hours under a contract translates into endless working days. In addition, living at a place of employment blurs the boundary between work and rest. This exploitation is showcased through physical and technological surveillance, including closed-circuit cameras and communication devices. Similarly, maritime crew endure confinement at sea, without sufficient rest intervals or social support networks. By restricting mobility and monitoring their daily behaviour, the employer completely transforms the workspace into an institution that only accepts utter compliance.

<u>Mental and Physical Effects</u>

As a result, there is a severe psychological and physiological toll on the workers. This was empirically documented by Anjara et al. (2017). Their cross-sectional study demonstrates that constant physical confinement and lack of rest correlate with elevated rates of anxiety, distress, and decreased quality of life among female MDWs in Singapore. To survive hyper-surveillance, workers are forced to accept internal coping strategies. Lai & Yeoh (2026) further exemplify this view through a "reflective thematic analysis" - migrant domestic workers rely on informal peer networks as a mechanism of preserving emotional resilience. In an oppressive environment, psychological agency acts as a tactic of self-preservation enacted under conditions of forced isolation.

Crucially, the micro-lived experiences of psychological and physiological distress endured by workers are not isolated workplace tragedies; rather, they serve as the direct manifestation of structural power asymmetries embedded within global migration governance. The exclusion of domestic and maritime labour from the Employment Act allows the host state to operationalise a global migration model that holds capital accumulation as the utmost priority. This severe mental and physical toll endured by workers reflects how decentralised global

governance is - allowing wealthy receiving states to enforce isolation and strip workers of legal protections, shifting the costs of economic integration onto the vulnerable workers.

Taking all into account, both physical isolation and immense control are not only secondary consequences of low-wage employment, but mechanisms that make labour exploitation possible. Trapping workers within a space isolated from public scrutiny allows the host country to diminish the workers' ability to organise and escape abusive conditions. Daily lived realities within offshore waters and private homes in Singapore highlight how confinement exacerbates mental health issues. This overall ensures that foreign labour remains productive and legally disempowered.

Conclusion

This paper has demonstrated that the constant isolation, legal disenfranchisement, and exploitation experienced by both foreign modern domestic workers (MDWs) and maritime workers in Singapore are not governance failures or isolated workplace abuses. Rather, as this multi-scalar analysis showcases, the micro-level realities are a deliberate outcome of an integrated socio-legal system: operating across individual, state, and global scales.

On the individual scale, spatial isolation within both private homes and offshore ships functions as a primary mechanism of labour extraction (Goh, 2019). These workers are trapped within environments of jurisdictional exceptionalism; they are trapped in a spatialised “state of exception” (Agamben, 1998; Yea, 2022). Through suspending routine workplace checks of the home or maritime vessel, the state reduces migrant workers to “bare life”, positioning them as a “homo sacer” (Anjara et al., 2017). In response to this, workers are forced to utilise micro tactics (such as emotional self-compassion) to maintain limited amounts of autonomy in a condition of forced isolation (Lai & Yeoh, 2026).

On the national scale, these regimes are normalised through Singapore’s statutory architecture, excluding both MDWs and maritime crew from Part IV of the Employment Act (1968). The state utilises Ruh’s (2013) framework of “numbers versus rights”, deliberately restricting labour rights to preserve employer flexibility and low costs. This “complicit network” (Chin, 1998) of employer-tied visas and $5000 security bonds - the state gives private employers complete authority. The privatised surveillance regime is sustained by a logic constructing low-skilled non-citizens as temporary economic inputs, excluded from any membership in the country (Neo, 2015).

On a macro scale, Singapore’s modern guest-worker regime can be highlighted as directly inheriting and refining imperial policies established during British colonial rule (Huff, 1994). The colonial settlements relied on vast networks of low-wage indentured labour to fuel mercantile expansion. This established an administrative blueprint: treating foreign workers as a factor of production who were insulated from rights of permanent settlement. Instead of the modern Singaporean state dismantling imperial mechanisms, it modernised them. By substituting colonial administration with a hyperefficient Work Permit System (Castles, 2010; Chin, 1998), the Singaporean state leverages postcolonial policies to bypass binding multilateral conventions

in favour of bilateral agreements. Furthermore, the source countries are trapped in "remittance dependency" - lacking diplomatic leverage to demand basic human rights for their citizens.

Overall, ending the isolation and confinement of domestic and maritime labourers in Singapore requires going beyond policy tweaks or isolated calls for empathy. Reform requires addressing the entire system of disempowerment - extending protections for these workers under domestic law, dismantling employer-tied visas (which were inherited from colonial networks), and establishing a binding multilateral standard to rebalance geopolitical power between host and source countries.

Building on the findings of this paper, future research could investigate how digital networks and cross-border civil society coalitions could disrupt the entrenched state-sanctioned asymmetries. Other researchers could examine whether transnational borders can leverage digital channels to challenge non-binding bilateral frameworks - forcing host states towards binding multilateral compliance. By exploring these different aspects of resistance, it will be vital to conclude whether the precarity of guest worker regimes can be dismantled within an increasingly digitalised global economy.

Until migrant labour is broken away from institutionalised state exclusion, the economic developmental model in Singapore will remain dependent on commodified labour of a legally disempowered workforce.

Acknowledgements

I would like to express my sincere gratitude to my mentors, Dr. Ida Danewid and Dr. Tanja Weis. Their invaluable guidance, constant prompt feedback, and generous support were vital throughout every stage of writing this paper. Their immense expertise in global governance and insightful mentorship constantly challenged me to expand my ideas and sharpen my analytical frameworks. Overall, the support of my mentors pushed the depth of my intellectual inquiry.

In addition, I would like to thank my parents; their constant encouragement provided the foundation and background support needed to write this paper. To my mother, thank you for your unwavering belief in my ability to write - your voice has given me the confidence to highlight my ambitious ideas on a scholarly level. To my father, thank you for always taking the time to help me clarify complex concepts, keep my paper clear, and ensure it remains valid and precise.

Finally, I would like to extend my appreciation to my A-level Human Geography teachers, Ms. Zarina Long and Mrs. Claire Huetson. Their exciting lessons and teaching cultivated my deep fascination for global migration governance, Additionally, their combined intellectual spark directly inspired the research that was undertaken in this paper.

Works Cited

Agamben, G. (1998). Homo Sacer: Sovereign Power and Bare Life (D. Heller-Roazen, Trans.). Stanford University Press. https://www.sup.org/books/religious-studies/homo-sacer

Amrith, S. S. (2013). Migration and diaspora in modern Asia. Bijdragen Tot de Taal-, Land- En Volkenkunde / Journal of the Humanities and Social Sciences of Southeast Asia, 169(4), 495–512. https://doi.org/10.1163/22134379-12340068

Anjara, S. G., Nellums, L. B., Bonetto, C., & Van Bortel, T. (2017). Stress, health and quality of life of female migrant domestic workers in Singapore: A cross-sectional study. BMC Women's Health, 17(1), 98. https://doi.org/10.1186/s12905-017-0442-7

Castles, S. (2010). Understanding global migration: A social transformation perspective. Journal of Ethnic and Migration Studies, 36(10), 1565–1586. https://doi.org/10.1080/1369183X.2010.489381

Castles, S. (2013). The forces driving global migration. Journal of Intercultural Studies, 34(2), 122– 140. https://doi.org/10.1080/07256868.2013.781916

Chin, C. (2019). Precarious work and its complicit network: Migrant labour in Singapore. Journal of Contemporary Asia, 49(4), 528–551. https://doi.org/10.1080/00472336.2019.1572209

Chin, C. B. N. (1998). In service and servitude: Foreign female domestic workers and the Malaysian "modernity" project. Columbia University Press.

European Implementation Network. (n.d.). Countries overview. Retrieved August 3, 2026, from https://www.einnetwork.org/countries-overview

Goh, D. P. S. (2019). Super-diversity and the bio-politics of migrant worker exclusion in Singapore. Identities, 26(3), 356–373. https://doi.org/10.1080/1070289X.2018.1530899

Huff, W. G. (1994). The economic growth of Singapore: Trade and development in the twentieth century. Cambridge University Press. https://doi.org/10.1017/CBO9780511470714

International Labour Organization. (2015). Bilateral agreements and Memoranda of Understanding on migration of low-skilled workers: A review (International Migration Papers No. 120). https://www.ilo.org/publications/bilateral-agreements-and-memoranda-understanding-migration-lowskilled

Kainz, L., & Betts, A. (2021). Power and proliferation: Explaining the fragmentation of global migration governance. Migration Studies, 9(1), 65–89. https://doi.org/10.1093/migration/mnaa015

Kaur, A. (2007). International labour migration in Southeast Asia: Governance of migration and women domestic workers. Australian National University. https://doi.org/10.25911/7E1R-EG94

Lai, R. L. L., & Yeoh, B. S. A. (2026). Self-compassion in context: A reflexive thematic analysis of migrant domestic workers' experiences in Singapore. International Journal of Qualitative Studies on Health and Well-Being, 21(1), 2617449. https://doi.org/10.1080/17482631.2026.2617449

Lemke, T., Casper, M. J., & Moore, L. J. (2020). Sovereign power and bare life: Giorgio Agamben. In Biopolitics (pp. 53–64). New York University Press. https://doi.org/10.18574/nyu/9780814753378.003.0009

Likić-Brborić, B. (2018). Global migration governance, civil society and the paradoxes of sustainability. Globalizations, 15(6), 762–778. https://doi.org/10.1080/14747731.2018.1503841

Macrotrends. (2026). Singapore fertility rate 1950–2026. Retrieved July 23, 2026, from https://www.macrotrends.net/global-metrics/countries/sgp/singapore/fertility-rate

Ministry of Manpower, Singapore. (n.d.-a). Foreign workforce numbers. Retrieved July 23, 2026, from https://www.mom.gov.sg/foreign-workforce-numbers

Ministry of Manpower, Singapore. (n.d.-b). Pre-employment medical examination for MDWs. Retrieved August 3, 2026, from https://www.mom.gov.sg/passes-and-permits/work-permit-forforeign-domestic-worker/eligibility-and-requirements/medical-examination Ministry of Manpower, Singapore. (n.d.-c). Security bond requirements for MDWs. Retrieved August 4, 2026, from https://www.mom.gov.sg/passes-and-permits/work-permit-for-foreign-domesticworker/eligibility-and-requirements/security-bond

Neo, J. L. (2015). Riots and rights: Law and exclusion in Singapore's migrant worker regime. Asian Journal of Law and Society, 2(1), 137–168. https://doi.org/10.1017/als.2015.1

Office of the Historian. (n.d.). A guide to the United States' history of recognition, diplomatic, and consular relations, by country, since 1776: Singapore. U.S. Department of State. Retrieved August 4, 2026, from https://history.state.gov/countries/singapore

Parreñas, R. S., Kantachote, K., & Silvey, R. (2021). Soft violence: Migrant domestic worker precarity and the management of unfree labour in Singapore. Journal of Ethnic and Migration Studies, 47(20), 4671–4687. https://doi.org/10.1080/1369183X.2020.1732614

Ruhs, M. (2013). The price of rights: Regulating international labor migration. Princeton University Press.

United Nations. (1990). International Convention on the Protection of the Rights of All Migrant Workers and Members of Their Families (A/RES/45/158). OHCHR. https://www.ohchr.org/en/instruments-mechanisms/instruments/international-convention-protectionrights-all-migrant-workers

Wickramasekara, P. (2008). Globalisation, international labour migration and the rights of migrant workers. Third World Quarterly, 29(7), 1247–1264. https://doi.org/10.1080/01436590802386478

Wickramasekara, P. (2012). Something is Better than Nothing? The Role of Bilateral Labour Agreements in the Protection of Migrant Workers (Global Migration Research Paper No. 7). Graduate Institute of International and Development Studies.

Wickramasekara, P. (2015). Bilateral agreements and Memoranda of Understanding on migration of low-skilled workers: A review (International Migration Papers No. 120). International Labour Organization.

https://www.ilo.org/publications/bilateral-agreementsand-memoranda-understanding-migration-low-skilled World Bank. (2026).
World Bank open data. Retrieved August 2, 2026, from https://data.worldbank.org Worldometer. (2026). Singapore GDP. Retrieved July 23, 2026, from https://www.worldometers.info/gdp/singapore-gdp/
Yea, S. (2022). Human trafficking and jurisdictional exceptionalism in the global fishing industry: A case study of Singapore. Geopolitics, 27(1), 238–259. https://doi.org/10.1080/14650045.2020.1856338

Protection Without Borders: Glocalization and the Distribution of Safeguarding in the NBA's Markets By Armaan Batra

Abstract

In 2014, the National Basketball Association (NBA) banned Donald Sterling, owner of the Los Angeles Clippers, for life within four days of a recording of his racist remarks becoming public. Six years later, when American coaches reported teenagers being physically abused at NBA-run training programs in China, the response was nothing alike. This paper asks the question of why a single institution like the NBA protects so differently. Specifically, I draw on the concepts of localization, which is how a business adapts its global products or services to fit local cultures, and international standards of human rights and due diligence. I argue that the NBA protects based on where it has direct control instead of protecting based on how severe the harm is. At the core of the league, protection is fast and the consequences are severe. Whereas at the edges, that protection falls apart because no one is clearly in charge of it. By comparing the NBA's core with its academies in China and a smaller case of India (Greater Noida) as well, I show that this gap is not a mistake the league can easily fix. It is a structural issue and built into the way that the NBA's markets went global. The protection of athletes at grassroots levels, I conclude, has been built without a sturdy perimeter. Until reform targets the structure of control rather than surface-level solutions like rewording the literature, the athletes at the edges will always remain out of the safeguards' reach.

Keywords glocalization, safeguarding, athlete protection, sport governance, human-rights due diligence, NBA, grassroots athletes

Introduction

In April of 2014, a recording surfaced in which Donald Sterling, the former owner of the Los Angeles Clippers, told an acquaintance not to bring Black people to his team's games. Four days later, NBA Commissioner Adam Silver, in New York, stated that he was "banning Mr. Sterling for life from any association with the Clippers organization or the NBA…" (NBA 2014). He imposed the largest fine the league's guidelines allowed and pushed Donald Sterling to sell the Clippers. All in four days. That speed showed that protection could be swift when the league needed it to be. Six years later, a very different set of complaints arrived from China. American coaches at NBA training academies had reported that the local coaches were physically abusing teenage players, and the schooling that the league promised to be "central" to the program was not actually taking place (Fainaru and Fainaru-Wada 2020). There was no four-day response in this matter. In fact, initially, there was no response at all. The complaints moved around inside the organizations and according to the NBA's deputy commissioner Mark Tatum, employees at NBA headquarters in New York weren't even briefed on the situation.

The distance between these two issues is the subject of my paper, and the question it raises is simple: what determines whether a sports institution protects an athlete from

discrimination and abuse? While it may seem easy to label this gap as a form of hypocrisy or even to say that the NBA just does not care about its grassroots athletes, neither of those answers is adequate in this intricate of a situation. Discrimination and abuse in sport are not only the acts of individuals; as Hylton (2010) argues, they are produced and allowed at the institutional level. For a long time now, sport has largely been a space where the interests of powerful organizations outweigh the rights of the young athletes (David 1999), and the modern sports community (global, commercial, and protective of its image) has only amplified this tension. The "whole-of-system" responsibility set by the International Olympic Committee (IOC) frames the protection of athletes from violence and abuse as something that should be run without interruption, from the individual athlete to the powerful institutions that employ them (Tuakli-Wosornu et al. 2024). The NBA is a prime example of this problem since it operates at both ends of this system simultaneously. It oversees an extremely profitable league at its core and runs youth academies at its edges, and the difference in protection between the two is strikingly immense.

To analyze this difference, I draw on the concept of glocalization. Giulianotti and Robertson (2004) use this term to describe how a global institution is not forced onto a local market but, in a way, is reshaped through it, so that the same global company gets different results in different geographic locations. Top sports organizations, in their account, work as "glocal transnational corporations…" (Giulianottie and Robertson 2004): global in reach, yet run locally at every point. Therefore, when the NBA expanded abroad, it did so in the same way: local partners. Partnerships with government-run facilities, national federations, and local staff that work directly with the athletes, even though the NBA itself does not employ them. Responding to the reports from China, NBA Deputy Commissioner Mark Tatum admitted that the league "didn't have the authority" over the local coaches and needed "more direct oversight" of academies that ran out of government facilities (Fainaru and Fainaru-Wada 2020). This admission names the problem directly: the same partnership that the NBA used to grow their markets had now allowed them to break the chain of responsibility, an example of that being that a complaint can travel partway up the organization and then disappear. And this pattern is not unique to basketball. It has also been documented in football academies across the global South, where young talent is developed for export with "weak duty of care", meaning that it looks like the organizations are developing athletes less for the players' benefit, and more for a larger organizations (Darby, Akindes, and Kirwin 2007). Protection, in summary, does not accidentally fail at the edges. It does because of how the system was built in the first place.

In this paper, I argue that the NBA hands out protection based on where it can keep direct control and not the severity of the harm. Where the league holds direct power (the core), and where discrimination can be attached instantly as a form of damage to the brand name of the NBA is where action is swift. At the edges, that responsibility is spread so extremely thin that no partner can alone be held accountable for it, a failure by the standards of human rights due diligence (Lin, Khan, and Song 2024; Byrne 2026). To further this case, I compare the NBA's core with its edges using three metrics: how fast and how severe the league's response is, who

actually controls the athlete (local partners or direct oversight), and where a complaint stops as it moves through the organization. This paper is split into three sections. First, I review what the current scholars have written about development, safeguarding at grassroots levels, and globalization in sport, and I show the gap that they have missed. Second, I build my framework using the aforementioned football academies as an earlier case. Third, I apply that framework to the Donald Sterling case at the core and the NBA academies in China and India at the edges, and I analyze the pattern these cases reveal. Overall, I conclude that the protection of star athletes and the disregard of the grassroots level athletes are not two different issues. They are the same one, just seen from two ends of the same spectrum.

Safeguarding, Development, and the Unasked Question

There are three bodies of scholarship that need to be understood to show how protection works in sport, and for the most part, they do not connect to each other. The first piece of scholarship explains what protection in sport should look like. The second shows where that protection fails, and the third argues that when it fails, the problem is built into the institution itself, not a matter of a few people who made bad decisions. Each of these bodies of literature explains one part of what my paper talks about, but none of them ask what decides how protection gets handed out in the first place. That gap is what I intend to answer.

The first body of literature talks about safeguarding. Safeguarding, as per the IOC consensus statement, states that athletes will receive protection from every form of interpersonal violence and abuse (Tuakli-Wosornu et al. 2024), and over the last twenty years , researchers and other bodies have mostly agreed that keeping athletes safe is a basic duty of the sport. The second part of the IOC consensus statement is the clearest example of this, calling protection a "whole-of-system" responsibility that should reach every level of a sport, from a youth athlete to the federation. This body of work has many strengths. It has a clear, shared, and extremely detailed standard for what protection should look like, and it also puts youth athletes at the forefront of that protection. More than twenty years ago, David (1999) pointed out that children provide a huge deal to competitive sport while getting the least protection from it, and newer data proves him right. Byrne (2026), writing on youth (specifically child) athletes at the Olympics, shows that a sports organization can say the right things about children's rights under international law but still fail to actually build those protections into their operations. However, a limitation of this scholarship is that it focuses solely on what protection *should* look like. It assumes that a single organization is the only one in control, and it treats gaps in protection as problems that can be fixed with better rules and vocational training. What it does not ask is what happens to that protection when a global institution, like the NBA, splits control across different countries (spreading its markets) and hands the everyday functions to foreign partners. This is exactly where the protection of athletes tends to fall apart.

The second body of literature fills in the other half of the picture, and it is far more precise about *where* that protection actually fails. Studies of sport-for-development find that programs that are supposed to aid youth athletes still fall short, and that sport tends to repeat the

existing inequalities rather than fix them (Hartmann and Kwauk 2011). Other reviews of this field describe the data as uneven and understudied, with a lot of activity but scarce evidence at the youth stage (Schulenkorf, Sherry, and Rowe 2016). Alongside this, a separate group of scholars who study race argue that discrimination in sport is not just individual bias: as Hylton (2010) puts it, racism works at the level of the institution, living in the systems and routines of sports organizations rather than simply in the actions of one person. The strengths of this work is that it refuses to accept the conventional norm. It also shows that protection often does not reach the grassroots level and that the reasons are more structural than previously thought. But it also has a blind spot. It tends to study the grassroots level alone and rarely connects that neglect to the heavy investment and fast protection at the top of the very same sport. This separation matters since treating these issues as two distinct ones does not allow the literature to ask whether these issues are connected or not. Hylton (2010) demonstrates racism as an institutional problem, but majorly within one country or one organization at a time. Neither asks how a single multinational institution can protect so strongly at its core and offer almost none at its edges.

That connection is the gap that my paper dives into. Putting the three bodies of scholarship together, we see a standard for what athletes are owed, strong evidence of where that standard fails, and the insight that the failures are institutional. However, what none of them offer is *why* the same organization reacts to harm at its center within days but not at all when the harm is peripheral. In other words, the literature has not paid enough attention to the possibility that protection does not follow the severity of the harm but rather where the institution holds direct control. This matters for three reasons. Firstly, if protection does follow control rather than harm, then it is no wonder that current reforms are constantly failing, because they are aiming at the wrong aspect. Secondly, leagues are expanding rapidly into markets where control has to be handed to local partners, which means the amount of youth athletes caught in this blind spot is increasing. Thirdly, when we view protection through the perspective of it being part of the system's design, we see the failure differently. It is not a small mistake that can be fixed at any given point in time. It is the natural result of how the system is built. Existing research still treats elite protection and grassroots disregard as two separate issues but I argue that they are actually two outcomes of the same structure.

Glocalization and the Extractive Academy

To formulate an argument, I need avenues in which I can adequately explain how one institution protects so selectively. This section builds that explanation out of three ideas: glocalization, which explains why a global institution produces uneven results in different geographic locations; the football academies in the global South as a place of extraction rather than development, which demonstrates what that unevenness caused by glocalization looks like; and human-rights due diligence, which shows that this kind of unevenness is a failure of accountability. To start, I build my framework on an older and better documented case than the NBA, so that the logic is sound before I apply it to the case of the NBA.

The first idea is glocalization. Giulianotti and Robertson (2004) use the term to describe how a global organization interacts with the local places it enters, so the same institution produces different results in different places. Essentially, top sports organizations work as glocal "transnational corporations" (Giulianotti and Robertson 2004, p. 545), meaning they are global in reach yet locally run. The consequence for protection is important: a single institution can have elite safeguards at its top and relatively none at its bottom without contradicting itself, since that unevenness can be attributed to the effect of globalization. This reframes the core-and-edges gap that my paper takes on. It is not proof that the institution is being inconsistent but rather proof of a larger theme: consistency was never part of the structure.

The second idea gives that abstraction a hold in reality. The clearest existing picture of what happens to grassroots-level athletes at the edges of a global sport comes from football's, or soccer's, academies in the global South. Darby, Akindes, and Kirwin (2007) demonstrate a system where young athletes are developed for export to richer markets under weak "duty of care"; an arrangement they call a form of extraction, where the value of the athlete flows globally but the responsibility for their protection and wellbeing remains regional. The academy is built less for the development and support for a young athlete's dreams and more for the pipeline for the institution. That is the template that I am going to carry over to the NBA: an academy structure that gathers up talent while spreading responsibility so thin that it disappears. And also to show that this pattern is not limited to one sport. At the protective end too, the same pattern shows up across football and basketball: when Paraguayan senator Celeste Amarilla directed racist remarks at Kylian Mbappe after France's victory over Paraguay in the 2026 FIFA World Cup, the response came quick and from the top: the Paraguayan Senate passed a motion to reject her "racist and discriminatory" remarks, the national government condemned said remarks publicly, and Mbappe's defense made its way across international media in mere days (CBC News 2026; Washington Post 2026). Meanwhile, the young athletes who face the same discrimination in the global South get no response. Therefore, the level of protection offered is centered more around multinational corporations than real damage done to people.

The third idea is turning that description into a metric that I can judge the NBA against. If protection at the edges fails because responsibility is delegated to local partners, then the league could argue that it cannot be blamed for any harm not caused directly by them. However, international human-rights governance rejects that notion. The framework of due-diligence associated with the UN Guiding Principles on Business and Human Rights maintains the fact that a global institution's responsibility runs through **all** of its business relationships (partners included) because in those kind of arrangements, direct legal authority is often missing (Lin, Khan, and Song 2024; Byrne 2026). Scholars in sports governance have labeled this difficulty as one of shared responsibility: the problem of pinning accountability down when harm occurs across a chain or partners, none of whom control the entire program (Rook, Prado, and Heerdt 2022). Therefore, by this standard, handing operations to a local partner does not hand off the responsibility of protection. Overall, this standard lets me move beyond describing the NBA's protection gap to analyzing it under three consistent subcriteria.

When synthesized, these three ideas provide a feasible framework and method. I treat the NBA as one institution seen at two different levels, core and edges, and I compare them across three measures: how fast and how severe the institution's response is, who holds the everyday control of the athletes, and where a complaint stops as it moves through the organization. My evidence is pulled from policy documents, other scholarship, and also investigative reporting since this field is one where information is majorly qualitative and unevenly documented (Schulenkorf, Sherry, and Rowe 2016). With this framework now developed and put into practice, I turn to the case of the NBA.

Protection at the Core, Abandonment at the Edges

Since protection follows control, then the clearest place to see the extent of the NBA's protective power is at its core, where it holds absolute control. The Donald Sterling case is the sharpest example of this. When a recording of the former Los Angeles Clippers owner's racist remarks became public in 2014, the league activated in just four days: Commissioner Adam Silver banned Sterling "for life", imposed the maximum fine that the NBA constitution permitted, and announced that he himself would push the other owners to force a sale of the team (NBA 2014). The severity was not improvised. Article 13 of the NBA Constitution allows the league the power to terminate an owner's franchise, and Article 24 lets the commissioner act against conduct deemed harmful to the league (NBA Constitution and By-Laws). When we run this against the three aforementioned measures that I developed earlier, we see that the league's response was swift and the consequences were severe. The control held by the NBA was direct since Donald Sterling was an owner directly bound to the constitution of the NBA, so the NBA had the authority to act on its own. And the complaint had nowhere to get lost since the harm was public, it landed at the center of the NBA organization, and the commissioner himself answered it. It is also worth noting that the league initiated its response so quickly due to the fact that there were concerns about Sterling before. What changed in 2014 was that the harm became public and actively attacked the brand name. Therefore, protection at the core switched on the instant the harm became the institution's issue.

This is glocalization in practice. As the NBA becomes more of a global entity, it does not simply extend its authority overseas unchanged; it re-forms itself around the local partners, and its safeguards get spread thinner at those points where direct control does. In 2020, an ESPN investigation had reported that American coaches at NBA training academies in China had told league officials that local coaches were physically abusing teenage players and that the schooling the league promised was not being delivered. The details were striking: although Adam Silver (NBA Commissioner) had said that education would be "central" to the program, coaches reported that two of the three academies provided no real schooling, and several athletes described their time there as harassment (Fainaru and Fainaru-Wada 2020). The response was neither fast nor severe: there was no ban, no fine, no four-day timeline. The league had simply said it passed concerns to local authorities and eventually shut down one of the academies. The control held was not direct. Since the academies were halfway across the world, they ran out of

local government-run facilities, staffed with local coaches that the NBA did not employ. And most importantly, the complaint stopped before it reached the top of the NBA organization. By the deputy commissioner's own words, staff at headquarters in New York were never informed of the reports at that time. This is the same mechanism as seen earlier in Section 2: the same partnership structure that let the league even operate in China also localized the complaint, so that it never traveled to the one place that had the power to act on it. When Mark Tatum admitted that the league lacked the authority and oversight to control those coaches, he was describing exactly how the structure of the system spreads responsibility so thin that it basically disappears.

In addition, India offers a smaller but equally revealing version of the exact same pattern, and it is very clearly seen from the public record. The NBA opened an elite academy in Greater Noida in 2017, part of an initiative that pushed the NBA into a market that they considered essential for global growth (NBA 2017). The academy ran on a familiar model: a local facility, local partners, and a promised "holistic" program of education and development. Yet in 2024 the league shut the Noida facility down (Deccan Herald 2024), and in 2025 it restructured its academy program worldwide, closing sites in Australia and Mexico and reallocating resources/funds towards markets that it deemed most important for expanding the league globally and generating long-term streams of revenue, a list that included countries like China and India (ESPN 2025). The league maintained this change as a strategic one and that is the whole point. At the core, the NBA's authority is fixed into a place that does not care about the business cycle whereas at the edges, the very infrastructure meant to develop and protect grassroots-level athletes is temporary. It is only built when a market looks promising and dismantled once the business strategy shifts. So protection at the edges is not only weaker than at the core, but also rests on the profitability of the market it sits in.

When we analyze across the three sites, the pattern formed is extremely consistent. Protection is swift, severe, and direct where the NBA holds control and where harm attacks the brand name. However, it is slow, mild, and very easily lost where control has been delegated to local partners and the harm stays far away from headquarters. The mechanism that connects them is the one that is predicted by the immense effect of glocalization. To recap, glocalization says that when a global institution expands into local settings, it does not stay intact. It gets reshaped by each place it enters, and its reach and control begin to come apart, the NBA and its academies being a prime example of this (Giulianotti and Robertson 2004). To grow, the NBA expands through local partners, and that expansion is now doing two things simultaneously: it moves the day-to-day control over grassroots-level athletes to staff that the league cannot control, and it shields the league's brand in the eventuality that a case like China and India arises. The same move that allows for the NBA to expand globally is also one that breaks the chain of responsibility. This is why the gap is structural. It is produced not in spite of expansion, but by its very design.

But a strong objection still lies in the case of the NBA. It is one that the NBA itself raised: the fact that it simply did not have the authority to discipline foreign coaches employed at those government facilities, so comparing China to Donal Sterling is unfair. To answer this

objection, there are two possible routes. Firstly, the standard of due diligence was written precisely for situations like this. The UN Guiding Principles hold an institution responsible across all its business relationships exactly because said institution often lacks legal authority over its partners. The absence of authority is not a defense but rather the reason the responsibility exists in the first place (Lin, Khan, and Song 2024; Byrne 2026). The second answer is much simpler and also harder to escape. The NBA chose this structure. Nothing forced the league to run its academies through partners it could not oversee, so the loss of accountability was not something that happened in this situation. It happened because of it. That makes the structure itself a failure, not an excuse for it. Measured against the "whole-of-system" standard that the league attempts to portray (Tuakli-Wosornu et al. 2024), we see that the NBA has built a system where protection only reaches the core.

Conclusion

This paper began with two reactions to harm inside one institution. It started with a lifetime ban handed down in a mere four days, and a set of abuse complaints that never reached the people with the power to act on them. I have argued that the distance between these two responses is not hypocrisy or indifference but something much larger: structure. The NBA protects fast and hard at its core, where it holds direct control and where harm can threaten the brand name. On the other hand, it protects weakly or not at all at its edges, where it has handed control to local partners and where harm could never gain enough traction to attack the brand name. Once again, the idea of localization explains why: a global institution grows by reshaping itself in local cultures, and in doing so, it spreads the responsibility so thin that it basically disappears. The China academies show that. The academy in Greater Noida shows how the infrastructure itself is viewed as temporary, opened and closed according to the commercial value of the market; all dependent on profitability.

Synthesizing the sections, the shape of the argument forms. The existing literature provided a standard for what athletes are owed, evidence of where that standard intersected with its limitations, and the extremely crucial insight that the failures are institutional. Reading them through the lens of glocalization and the extractive academy shows these two issues as two ends of a single structure. And the standard of due diligence shows that this is not a limit the league can hide behind. It shows that responsibility cannot disappear at the border of a partnership that the institution chose to create.

However, there are also several questions that my paper cannot fully answer. Does the same pattern appear in other sports and other regions across the world? What would a protection system that does not depend on profitability actually require? And would better data on grassroots athletes that the field still lacks (Schulenkorf, Sherry, and Rowe 2016), change the picture we see? These are all large questions, but they share a common theme in my paper. The whole-of-system standard is already existing. The only missing piece is the will to extend this standard to the young athletes furthest from the brand name, not just to the NBA superstars. What is missing is to build protection that holds at the edges just like it does at its core.

Works Cited

Byrne, Seamus. 2026. "The Rights of Olympic Child Athletes: The Need for Rights-Based Reform." The International Sports Law Journal. Published online April 7. https://doi.org/10.1007/s40318-026-00339-x.

CBC News. 2026. "Paraguayan Senator Defends Racist Rant Against Kylian Mbappé—and Insults Him Again." July 10. https://www.cbc.ca/news/world/mbappe-world-cup-racist-senator-9.7265511.

Darby, Paul, Gerard Akindes, and Matthew Kirwin. 2007. "Football Academies and the Migration of African Football Labour to Europe." Journal of Sport and Social Issues 31 (2): 143–161 https://www.academia.edu/107524938/Sport_Migration_and_Gender_in_the_Neoliberal_Age?source=swp_share

David, Paul. 1999. "Children's Rights and Sports: Young Athletes and Competitive Sports—Exploit and Exploitation." The International Journal of Children's Rights 7 (1): 53–81. https://www.researchgate.net/publication/249570493_Children%27s_rights_and_sports_Young_athletes_and_competitive_sports_exploit_and_exploitation

Deccan Herald. 2024. "NBA Academy's Noida Base to Cease Operations." March 24. https://www.deccanherald.com/sports/other-sports/nba-academy-s-noida-base-to-cease-operations-2950849

ESPN. 2025. "NBA Academy Venture to Undergo Significant Restructuring." ESPN. https://www.espn.com/nba/story/_/id/42270842/nba-academy-venture-undergo-significant-restructuring

Fainaru, Steve, and Mark Fainaru-Wada. 2020. "ESPN Investigation Finds Coaches at NBA China Academies Complained of Player Abuse, Lack of Schooling." ESPN, July 29. https://www.espn.com/nba/story/_/id/29553829/.

Giulianotti, Richard, and Roland Robertson. 2004. "The Globalization of Football: A Study in the Glocalization of the 'Serious Life.'" The British Journal of Sociology 55 (4): 545–568. https://doi.org/10.1111/j.1468-4446.2004.00037.x

Hartmann, Douglas, and Christina Kwauk. 2011. "Sport and Development: An Overview, Critique, and Reconstruction." Journal of Sport and Social Issues 35 (3): 284–305. https://www.researchgate.net/publication/258158933_Sport_and_Development_An_Overview_Critique_and_Reconstruction

Hylton, Kevin. 2010. "How a Turn to Critical Race Theory Can Contribute to Our Understanding of 'Race', Racism and Anti-Racism in Sport." International Review for the Sociology of Sport 45 (3): 335–354. https://www.researchgate.net/publication/249734000_How_a_turn_to_critical_race_theory_can_contribute_to_our_understanding_of_'race'_racism_and_anti-racism_in_sport

Lin, Shuang, Asif Khan, and Yu Song. 2024. "Upholding Human Rights in Mega Sports: A Study of Governance Practices within the IOC and FIFA through the Lens of the Ruggie Principle." Heliyon 10 (15): e35607. https://doi.org/10.1016/j.heliyon.2024.e35607.

National Basketball Association (NBA). 2014. "Adam Silver Bans Donald Sterling for Life." NBA press release, April 29. https://www.nba.com/cavaliers/releases/nba-bans-sterling-140429

National Basketball Association (NBA). 2017. "NBA Officially Opens First Academy of Its Kind in India." NBA press release, May 2017. https://www.nba.com/news/nba-academy-india-official-release

National Basketball Association (NBA). n.d. Constitution and By-Laws of the National Basketball Association. https://atlhawksfanatic.github.io/NBA-BY-LAWS/

Rook, William, Mónica Prado, and Daniela Heerdt. 2022. "Responsible Sport: No Going Back." The International Sports Law Journal 22: 85-89. https://doi.org/10.1007/s40318-022-00231-4

Schulenkorf, Nico, Emma Sherry, and Katie Rowe. 2016. "Sport-for-Development: An Integrated Literature Review." Journal of Sport Management 30 (1): 22–39. https://www.researchgate.net/publication/279283205_Sport-for-Development_An_Integrated_Literature_Review

Tuakli-Wosornu, Yetsa A., et al. 2024. "IOC Consensus Statement: Interpersonal Violence and Safeguarding in Sport." British Journal of Sports Medicine 58 (22): 1322–1344. https://pubmed.ncbi.nlm.nih.gov/39586634/

Washington Post. 2026. "Kylian Mbappé Condemns Paraguayan Senator over Racist Remarks After World Cup Match." July 8, 2026. https://www.washingtonpost.com/sports/soccer/2026/07/06/mbappe-paraguay-racism-world-cup/d23ea260-797f-11f1-b194-f872dd4ec5aa_story.html

Balancing Potency, Stability, and Solubility In MMAE-Based Antibody-Drug Conjugates: Dual-Payload and Advanced Linker Strategies By Yen Kai Chan

Abstract

Advanced linker strategies in monomethyl auristatin E (MMAE)-based antibody-drug conjugates (ADCs) have produced substantial gains in preclinical performance. Across the literature reviewed here, selected platforms report therapeutic indices as high as 60-fold, complete tumor regression in particular xenograft models, and the ability to tolerate higher drug loading than conventional vedotin-type conjugates. These outcomes are associated with approaches including tandem-cleavage linkers, tumor-responsive activation, hydrophilicity-enhancing modifications, stabilized conjugation chemistries, and, more recently, dual-payload architectures.

These gains are interdependent rather than additive. Increasing drug loading can increase cytotoxic capacity, but it also increases hydrophobic burden and may promote aggregation or accelerated clearance. Stabilizing a linker can reduce premature payload release, yet excessive dependence on intracellular processing may limit activity in settings with inefficient internalization. Hydrophilic modifications can improve pharmacokinetics while altering payload diffusion and bystander activity.

This review therefore evaluates MMAE-based ADC optimization as a constrained design problem governed by interactions among potency and drug loading, stability, and solubility. Dual-payload systems are considered as one emerging strategy within a broader set of advanced linker and conjugation approaches. The literature suggests that no single structure maximizes all three properties simultaneously; instead, effective design depends on target biology, internalization behavior, payload characteristics, and the intended therapeutic profile.

Methods

Literature Search Strategy

A structured qualitative literature search was conducted using Semantic Scholar and OpenAlex, supplemented by manual reference tracing from relevant review articles. Searches were performed between March and April 2026.

Primary search terms included combinations of:

- “MMAE antibody drug conjugate”
- “dual payload ADC”
- “linker stability ADC”
- “PEGylation ADC MMAE”
- “tandem cleavage linker ADC”

Searches were restricted to publications between 2003 and 2025 to capture both early development of MMAE-based ADCs and more recent advances in linker design. No language restrictions were applied, though all included studies were available in English.
The search produced roughly 100 candidate records. Because discovery-platform searches and manual reference tracing were conducted iteratively, an exact PRISMA-style identification count was not maintained. Titles and abstracts were screened to exclude studies that:

- did not involve MMAE-based ADCs
- lacked experimental validation
- focused exclusively on non-therapeutic applications

After title, abstract, and relevance screening, 25 publications were retained for full-text review and qualitative synthesis.

This review is qualitative in nature. Efforts were made to maintain consistency in data extraction by prioritizing studies reporting quantitative measures of efficacy, stability, and solubility. Because experimental systems differed, comparisons were interpreted within the context of each study rather than treated as directly equivalent.

Study Selection Criteria

Studies were included if they met all of the following criteria:
1. Investigated MMAE-based antibody drug conjugates
2. Evaluated linker systems (including dual-payload, cleavable, non-cleavable, or PEGylated designs)
3. Reported quantitative or semi-quantitative measures of performance (e.g., cytotoxicity, efficacy, therapeutic index)
4. Included data relevant to stability (e.g., plasma stability, premature release) or solubility (e.g., aggregation, hydrophobicity)

Studies were excluded if they:

- lacked sufficient methodological detail
- reported only theoretical or computational models without experimental validation
- focused on unrelated payload systems

Peer-reviewed primary studies were prioritized. Preprints and conference abstracts were considered only where they addressed a relevant platform for which peer-reviewed evidence was limited; evidence from these sources was treated as preliminary and was not given the same evidentiary weight as full peer-reviewed studies.

Data Extraction and Analysis

Data were extracted with attention to:

- linker architecture and conjugation strategy

- drug-to-antibody ratio (DAR)
- measures of efficacy (IC_{50}, EC_{50}, tumor regression)
- stability metrics (plasma half-life, premature release)
- solubility indicators (aggregation, hydrophilicity modifications)

Due to heterogeneity across experimental systems, direct numerical comparisons between studies were avoided unless conducted under comparable conditions. Instead, results were interpreted in relative terms within each study.
A qualitative synthesis approach was used to identify recurring design patterns and mechanistic relationships.

Synthesis Approach

A qualitative synthesis was performed to evaluate how advanced MMAE linker and conjugation strategies influence therapeutic performance relative to conventional designs. Dual-payload systems were treated as a distinct emerging subgroup rather than as a label for the wider linker literature. Particular attention was given to interactions among drug loading, stability, solubility, pharmacokinetics, and payload release.

Limitations

This review is limited by heterogeneity in experimental models, targets, dosing regimens, assay conditions, and reporting standards. Numerical results from different studies therefore should not be interpreted as direct head-to-head comparisons unless the original experiments used comparable conditions. The literature is also predominantly preclinical, and several advanced platforms have limited long-term toxicology or human validation. Finally, the search was structured but was not conducted as a formal systematic review, so the evidence base should be interpreted as a focused qualitative synthesis rather than an exhaustive meta-analysis.

Results

Studies employed diverse experimental systems, with HER2 being the most common target antigen tested across multiple studies (Wang et al., “From Prodrug”; Wang et al., “Ionized Cys-Linker”; McKertish and Kayser; Goldberg et al.). CD30-targeting ADCs were evaluated in several studies using both Hodgkin lymphoma and other hematologic malignancy models (Doronina et al.; Burke et al.; Ochtrop et al.; Alley et al.). Animal models predominantly used mice with tumor xenografts, with select studies incorporating rat models for pharmacokinetic and tolerability assessments (Chuprakov et al.; Burke et al.; Ochtrop et al.).

Therapeutic Efficiency Improvements

In Vitro Cytotoxicity and Potency

Several studies reported strong cytotoxic activity with advanced ADC architectures. In one dual-payload trastuzumab conjugate, the reported GR50 was 0.29 nM in HER2-high SK-BR-3 cells, compared with 1967 nM for trastuzumab alone (McKertish and Kayser). This comparison demonstrates the activity of the complete conjugate relative to the unconjugated antibody within that assay, but it should not be interpreted as a general fold-improvement attributable solely to dual-payload design because the conjugate also contained two cytotoxic agents and the result is specific to the tested cell model.

Ionized Cys-linker MMAE systems maintained high potency at approximately 10^{-11} M while showing reduced bystander toxicity around 10^{-9} M (Wang et al., "Ionized Cys-Linker"). Hypoxia-sensitive systems showed even more context-dependent behavior. Under hypoxic conditions, EC50 values ranged from 0.00012 nM to 0.46 nM, while remaining much less active under normoxic conditions (Wang et al., "From Prodrug"). This oxygen-dependent activation highlights an important point. Potency in these systems is not fixed, but depends on the surrounding environment.

In Vivo Efficacy and Tumor Regression

In animal models, several linker strategies achieved strong tumor control. Tandem-cleavage ADCs produced complete regressions in a reported treatment cohort, while PEGylated glucuronide-MMAE conjugates achieved 10/10 cures in a Raji lymphoma xenograft model at 3 mg/kg (Chuprakov et al.; Burke et al.). High-DAR dual-payload ADCs have also shown improved tumor suppression relative to comparator therapies in selected HER2-positive preclinical models (Watanabe et al.). These findings demonstrate the potential of the respective platforms, but differences in tumor model, antibody, payload composition, dosing, and endpoint prevent direct ranking across studies.

Therapeutic Index Expansion

Therapeutic-window improvements were reported across several platforms. Doronina et al. reported therapeutic indices as high as 60-fold for peptide-linked auristatin conjugates, establishing early evidence that linker design can substantially alter efficacy and tolerability. Other advanced systems have reported increased tolerated drug loading or high maximum tolerated doses in preclinical models. These values are platform- and species-specific and should not be treated as directly comparable measures of safety.

Mechanistic Advantages

Controlled release mechanisms contributed to improved therapeutic efficiency across multiple platforms. The tandem-cleavage strategy employed two sequential enzymatic cleavage events, with glucuronide moiety removal upon internalization, followed by dipeptide degradation to release the payload specifically inside target cells (Chuprakov et al.). This approach reduced premature payload release in circulation, minimizing off-target toxicities, particularly myelosuppression (Chuprakov et al.).

The hypoxia-activated azobenzene linker remained non-cleavable in healthy tissues with oxygen levels greater than 10%, while cleaving to release MMAE in hypoxic tumor microenvironments with oxygen levels less than 1% (Xiao et al.). This tumor-specific activation mechanism reduced off-target toxicity by preventing payload release in non-tumor cells (Xiao et al.).

Site-specific targeting advantages were demonstrated through various approaches. The ionized Cys-linker-MMAE utilized receptor-mediated endocytosis for targeted delivery while its ionized nature reduced membrane permeability and off-target toxicity (Wang et al., "Ionized Cys-Linker"). The system maintained high potency at the tubulin molecular level while reducing bystander effects (Wang et al., "Ionized Cys-Linker"). Non-internalizing antibody conjugates with cleavable linkers demonstrated that controlled release mechanisms were essential for therapeutic activity, with diffusion of released payloads within the tumor mass and eventual internalization by cancer cells (Dal Corso et al.).

Stability Challenges and Solutions

Plasma Stability and Premature Release

Premature drug release remains a major challenge in ADC design. In a peer-reviewed study of tandem-cleavage linkers, monocleavage constructs lost up to 20% of payload during seven days in rat serum, whereas the tested tandem-cleavage constructs showed no detectable payload loss over the same period (Chuprakov et al.).

The glucuronide-protected tandem-cleavage system addressed the sensitivity of peptide linkers to extracellular enzymes such as elastase through steric protection (Chuprakov et al.). The glucuronide moiety was removed upon ADC internalization and lysosomal degradation, allowing controlled dipeptide degradation and payload release (Chuprakov et al.).

Self-stabilizing maleimide chemistry addressed thiosuccinimide instability under physiological conditions. The mDPR-based ADCs demonstrated high stability over 7 days in rat plasma compared to conventional MC-based ADCs (Burke et al.). This approach prevented premature drug release caused by reversible reaction of the MC stretcher unit with cysteine thiols and capture by endogenous albumin (Burke et al.).

The stabilized glycopeptide linker with maleimide stabilization showed higher resistance to cleavage by serum and lysosomal enzymes compared to valine-citrulline conjugated ADCs (Goldberg et al.). The hypoxia-sensitive ADCs (mil40-5 and NAC-5) exhibited excellent stability in plasma with less than 2% MMAE release, and NAC-5 demonstrated approximately 77% of the conjugate remained unchanged in liver microsomes with no MMAE release (Wang et al., "From Prodrug").

Comparative Linker Stability

Multiple studies compared stability across different linker architectures. Valine-alanine and valine-citrulline linkers showed comparable stability in PBS with minimal loss over seven

days (97.92% and 97.05% AUC retention, respectively) (Wang et al., “Valine Alanine”). However, non-cleavable linkers, while stable in circulation, lacked therapeutic activity in non-internalizing antibody contexts due to the inability to release the drug through proteolytic degradation pathways (Dal Corso et al.).

The Val-Ala linker demonstrated superior performance compared to Val-Cit, Val-Lys, and Val-Arg variants in vivo stability assessments (Dal Corso et al.). Mass spectrometric analysis revealed differences not only in overall stability but also in specific cleavage sites among the four linker variants (Dal Corso et al.).

Solubility Issues and Formulation Strategies

Hydrophobicity and Aggregation

The high lipophilicity of MMAE posed fundamental challenges for ADC development. The hydrophobic nature of MMAE led to poor hydrophilicity and physicochemical properties with conventional linkers like vc-PAB, typically limiting such ADCs to low DAR values up to 4 with narrow therapeutic indices (Liu et al.). Hydrophobic payloads were consistently prone to aggregation and faster plasma clearance, which in turn reduced efficacy and increased toxicity (Long et al.).

Excessive protein aggregation resulted from exposure of hydrophobic regions and structural disruption when conjugating via cysteine residues (McKertish and Kayser). High-loaded ADCs showed inferior in vivo performance due to increased hydrophobicity, leading to accumulation and aggregation in the liver, triggering fast clearance and reduced efficacy (Ochtrop et al.).

PEGylation Strategies

Multiple studies employed PEGylation to address hydrophobicity challenges. The LD343 linker platform incorporated high hydrophilicity to fully offset MMAE's hydrophobicity, enabling homogeneous DAR8 ADCs with markedly improved physicochemical properties compared to vedotin-based DAR4 conjugates (Liu et al.). The mPEG24 side chain on the valine-lysine-PAB linker achieved maximum hydrophilicity and biophysical stability (Long et al.).

PEG chain length can materially affect ADC pharmacology. Burke et al. reported progressively slower clearance with increasing PEG length and identified PEG8 as a threshold beyond which additional chain length did not further reduce clearance in their tested system. A PEG12 glucuronide-MMAE linker incorporating a self-stabilizing maleimide emerged as the lead design for homogeneous DAR8 conjugates in that study. These values should be interpreted as properties of the specific platform rather than universal PEG-length rules.

The phosphonamidate-based conjugation handle combined a discrete hydrophilic PEG-substituent with an established linker-payload in a compact building block (Ochtrop et al.). This compact branched PEG architecture introduced hydrophilicity without increasing the

distance between the antibody and the payload, enabling the first homogeneous DAR8 ADC from VC PAB-MMAE without increased in vivo clearance rates (Ochtrop et al.).

Comparative Performance Across Linker Systems

Head-to-Head Efficacy Comparisons

Direct comparisons demonstrated the advantages of advanced linker systems. The PEG12-glucuronide-MMAE system showed comparable antitumor activity to conventional Val-Cit MMAE at one-third of the antibody dose (Burke et al.). The valine-alanine linker showed comparable performance to valine-citrulline in efficacy and selectivity, with increased selectivity between HER2-positive and HER2-negative cells by 2400- to 26,000-fold (Wang et al., "Valine Alanine").

The DAR4 cysteine rebridged conjugate demonstrated the highest in vitro potency compared to lower DAR conjugates and superior in vivo efficacy versus trastuzumab-DM1 in low HER2-expressing models (Bryant et al.). Trastuzumab-MMAU showed superior efficacy to trastuzumab-vc-MMAE in xenograft tumor growth inhibition (Goldberg et al.).

Pharmacokinetic Improvements

Advanced linker systems achieved favorable pharmacokinetic profiles. The LD343 platform conferred stable pharmacokinetics indistinguishable from unconjugated parent antibody and superior to vedotin conjugates (Liu et al.). The trastuzumab-MMAU conjugate exhibited increased terminal half-life and exposure compared to conventional systems (Goldberg et al.).

Longer PEG chains resulted in slower clearance with threshold effects at PEG8 length (Burke et al.). Conjugates bearing sufficient PEG length to minimize plasma clearance provided wider therapeutic windows relative to faster-clearing conjugates with shorter PEGs (Burke et al.). The phosphonamidate-based DAR8 ADC showed slow clearance rates similar to lower DAR ADCs despite higher drug loading (Ochtrop et al.).

Safety and Tolerability

Multiple systems demonstrated improved safety profiles. Tandem-cleavage linker ADCs showed significantly better safety with reduced myelosuppression and no evidence of hematologic toxicity compared to conventional systems (Chuprakov et al.). The ionized Cys-linker-MMAE showed no significant hematologic toxicity with MTD approaching that of naked antibody at 160 mg/kg (Wang et al., "Ionized Cys-Linker").

The hypoxia-sensitive linker achieved an MTD exceeding 160 mg/kg with reduced hepatotoxicity and no significant tissue lesions at 20 mg/kg dosing (Wang et al., "From Prodrug"). The valine-alanine system demonstrated negligible toxicity with minimal hematological effects and no significant differences in hematological parameters between treatment and vehicle groups (Wang et al., "Valine Alanine").

Cross-linking approaches using dibromomaleimide linkers demonstrated improved pharmacokinetics, superior efficacy, and reduced toxicity in vivo compared to conventional heterogeneous ADCs (Behrens et al.). The stabilized glycopeptide linker was tolerated at doses up to 12 mg/kg in cynomolgus monkeys, representing 2- to 4-fold higher dosing than vedotin ADCs (Goldberg et al.).

Critical Evaluation

Advanced MMAE linker and conjugation systems can improve drug loading, stability, pharmacokinetics, or therapeutic index, but the evidence does not support treating these approaches as interchangeable. Dual-payload ADCs are one emerging architecture within this wider design space. Increasing DAR can increase delivered cytotoxic payload, yet it can also increase hydrophobicity, aggregation, and clearance if the additional physicochemical burden is not controlled. Consequently, potency gains cannot be evaluated independently from formulation and pharmacokinetic behavior.

PEGylation strategies effectively address these limitations by increasing hydrophilicity and enabling the generation of homogeneous high-DAR conjugates with improved pharmacokinetics (Burke et al.; Ochtrop et al.). However, this modification is not without consequence. Increased hydrophilicity has been associated with reduced membrane permeability and diminished bystander killing effects, potentially limiting efficacy in heterogeneous tumor environments where payload diffusion is beneficial (Wang et al., “Ionized Cys-Linker”; Bryden et al.). This creates a trade-off between safety and tumor penetration, which becomes more relevant in heterogeneous tumor environments.

Similarly, advances in linker stability, particularly through tandem-cleavage and self-stabilizing chemistries, have significantly reduced premature payload release and improved tolerability (Chuprakov et al.; Burke et al.). While these systems enhance specificity and reduce systemic toxicity, they also introduce increased structural complexity and dependence on intracellular processing mechanisms. In cases where antibody internalization is inefficient, highly stable linkers may reduce therapeutic activity due to insufficient payload release (Dal Corso et al.).

Taken together, the studies support a context-dependent rather than universally optimal design strategy. Linker selection depends on the target antigen, internalization efficiency, payload properties, desired bystander effect, and acceptable systemic exposure. Dual-payload systems may broaden the available design space by combining distinct payloads or release behaviors, but their added manufacturing and characterization complexity, together with limited clinical validation, prevents concluding that they are categorically superior to optimized single-payload systems.

Synthesis

The findings in this review of the performance of MMAE-based ADCs can be conceptualized as a balance between three interdependent parameters: drug loading (drug-to-antibody ratio, DAR), molecular stability, and solubility.

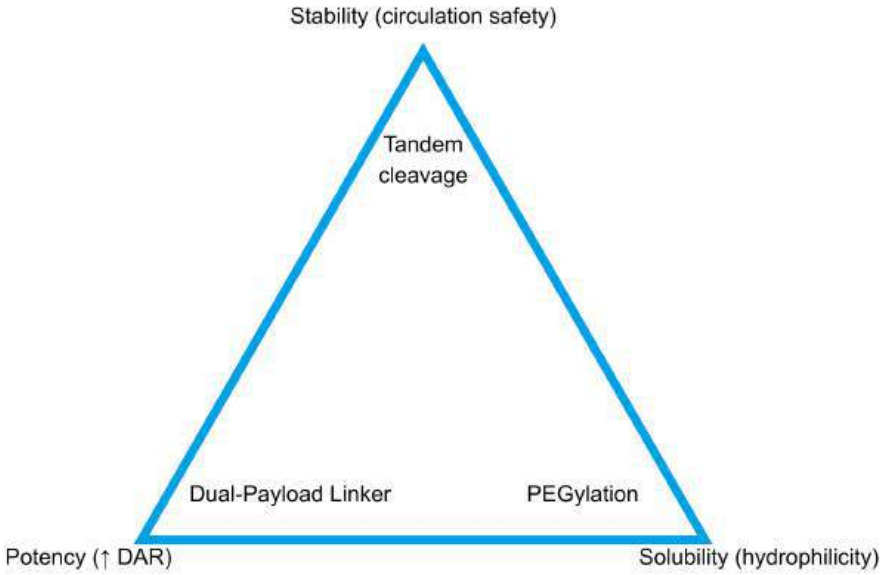

Figure 1.

Figure 1. Conceptual framework illustrating the trade-offs in MMAE-based ADC design. Increasing drug-to-antibody ratio (DAR) enhances potency but introduces hydrophobicity and instability. PEGylation improves solubility, while tandem-cleavage and stabilized linker chemistries enhance stability. Taken together, these results suggest that optimal ADC performance depends on balancing these competing parameters, rather than simply maximizing one (Chuprakov et al.; Liu et al.; Burke et al.; Ochtrop et al.).

To counteract these effects, hydrophilicity-enhancing strategies such as PEGylation are introduced, which improve solubility and pharmacokinetic profiles but may influence payload diffusion and bystander activity (Burke et al.; Bryden et al.). Concurrently, stability must be maintained through linker design, with tandem-cleavage systems and self-stabilizing chemistries reducing premature payload release while preserving intracellular activation (Chuprakov et al.; Burke et al.). Thus, optimal ADC design requires coordinated modulation of these three variables rather than maximization of any single parameter.

The literature supports three recurring design objectives for MMAE-based ADCs: managing hydrophobic burden through PEGylation or other hydrophilic modifications, controlling payload release through cleavable or tumor-responsive mechanisms, and improving circulating stability through stabilized conjugation chemistry or protective linker features. These strategies address different failure modes and can be combined, but their benefits are conditional rather than universally additive.

Context-Specific Optimization

The optimal linker configuration depends significantly on antibody internalization properties and target biology. For internalizing antibodies targeting HER2, CD30, and similar antigens, both cleavable peptide-based systems and non-cleavable systems with amino acid appendages demonstrated efficacy (Wang et al., "Ionized Cys-Linker"; Doronina et al., "Enhanced Activity"; Alley et al.). However, for non-internalizing antibodies such as the tenascin-C-targeting F16, cleavable linkers proved essential for therapeutic activity, with non-cleavable linkers showing no anticancer effect (Dal Corso et al.). This distinction reflects fundamental differences in payload release requirements based on antibody trafficking.

For high-DAR conjugates, sufficient hydrophilicity is important for maintaining favorable pharmacokinetics. In the PEGylated glucuronide-MMAE platform studied by Burke et al., PEG8 was the reported clearance threshold and PEG12 was selected for the lead linker. Other compact hydrophilic architectures have likewise enabled high drug loading without the clearance penalty expected from strongly hydrophobic conjugates. These findings support platform-specific optimization of hydrophilicity rather than a universal PEG8-12 design rule.

Stability-Efficacy Trade-offs

The relationship between linker stability and efficacy reveals important mechanistic insights. Highly stable non-cleavable linkers with ionized payloads achieved excellent plasma stability but required intracellular antibody degradation for drug release, resulting in reduced bystander effects (Wang et al., "Ionized Cys-Linker"). Conversely, cleavable linkers balanced controlled release with stability requirements, with tandem-cleavage systems offering superior protection against premature degradation while maintaining efficient intracellular payload liberation (Chuprakov et al.).

The hypoxia-activated systems represent a distinct approach where stability is context-dependent rather than absolute. These linkers remained stable under normoxic conditions typical of healthy tissues (>10% O_2) but activated specifically in hypoxic tumor microenvironments (<1% O_2) (Wang et al., "From Prodrug"; Xiao et al.). This tumor-selective activation achieved hypoxia cytotoxicity ratios of 10- to 27-fold, expanding therapeutic windows substantially beyond conventional systems (Wang et al., "From Prodrug").

Manufacturing and Formulation Considerations

The practical challenges of preparing high-DAR ADCs with hydrophobic payloads emerged as a recurring theme. Multiple studies noted aggregation issues and reduced recovery rates with conventional linkers at DAR values above 4 (Liu et al.; Ochtrop et al.; McKertish and Kayser). Solutions included PEGylation strategies that improved not only final ADC properties but also conjugation efficiency and labeling yields (Ochtrop et al.; Li et al.).

The valine-alanine linker offered advantages in the generation process with higher limpidity and increased recovery rates compared to valine-citrulline, suggesting practical

manufacturing benefits beyond therapeutic performance (Wang et al., “Valine Alanine”). Similarly, the compact phosphonamidate-based building blocks provided efficient one-pot reduction and alkylation protocols for generating homogeneous DAR8 conjugates from non-engineered antibodies (Ochtrop et al.).

Unresolved Challenges

Despite substantial progress, certain limitations persist across platforms. The optimal balance between bystander effect and off-target toxicity remains challenging, with hydrophilic modifications generally reducing bystander killing potential (Wang et al., “Ionized Cys-Linker”) while improving safety profiles (Liu et al.; Burke et al.). The dual-payload linker approach offers one solution by combining different mechanisms and release kinetics (McKertish and Kayser; Watanabe et al.), though the complexity of manufacturing and characterization increases substantially.

Long-term stability data and detailed toxicity profiles in non-human primates remain limited for many advanced systems. While several studies demonstrated improved short-term plasma stability and rat tolerability (Chuprakov et al.; Burke et al.; Goldberg et al.), comprehensive toxicology programs comparable to those supporting approved ADCs were not extensively reported. The clinical translatability of preclinical therapeutic index improvements requires validation in human studies for most platforms described.

Conclusion

Advanced linker and conjugation strategies, including dual-payload architectures, have expanded the design space of MMAE-based ADCs and improved performance in specific preclinical settings.

However, the evidence reviewed here does not identify a universally superior linker architecture. Potency and drug loading, stability, and solubility are interdependent, and improving one property can create constraints elsewhere in the system.

Future optimization should therefore be matched to biological context, including target expression, antibody internalization, tumor heterogeneity, payload permeability, and the desired balance between bystander activity and systemic safety.

Most of the advanced systems discussed remain supported primarily by preclinical evidence. Translation will require consistent comparative evaluation, longer-term toxicology, manufacturability assessment, and clinical validation before improvements observed in experimental models can be assumed to produce superior patient outcomes.

Works Cited

Alley, Stephen C., et al. "Intracellular Metabolism of Antibody-Drug Conjugates: Identification and Quantitation of Released Drugs." Cancer Research, vol. 66, no. 8 Supplement, 2006, abstract 469.

Behrens, Christopher R., et al. "Antibody-Drug Conjugates (ADCs) Derived from Interchain Cysteine Cross-Linking Demonstrate Improved Homogeneity and Other Pharmacological Properties over Conventional Heterogeneous ADCs." Molecular Pharmaceutics, vol. 12, no. 11, 2015, pp. 3986-3998. https://doi.org/10.1021/acs.molpharmaceut.5b00432

Bryant, Peter, et al. "In Vitro and In Vivo Evaluation of Cysteine Rebridged Trastuzumab-MMAE Antibody Drug Conjugates with Defined Drug-to-Antibody Ratios." Molecular Pharmaceutics, vol. 12, no. 6, 2015, pp. 1872-1879. https://doi.org/10.1021/acs.molpharmaceut.5b00116

Bryden, Francesca, et al. "Impact of Cathepsin B-Sensitive Triggers and Hydrophilic Linkers on In Vitro Efficacy of Novel Site-Specific Antibody-Drug Conjugates." Organic & Biomolecular Chemistry, vol. 16, no. 11, 2018, pp. 1882-1889. https://doi.org/10.1039/C7OB02780J

Burke, Patrick J., et al. "Optimization of a PEGylated Glucuronide-Monomethylauristatin E Linker for Antibody-Drug Conjugates." Molecular Cancer Therapeutics, vol. 16, no. 1, 2017, pp. 116-123. https://doi.org/10.1158/1535-7163.MCT-16-0343

Chuprakov, Stepan, et al. "Tandem-Cleavage Linkers Improve the In Vivo Stability and Tolerability of Antibody-Drug Conjugates." Bioconjugate Chemistry, vol. 32, no. 4, 2021, pp. 746-754. https://doi.org/10.1021/acs.bioconjchem.1c00029

Dal Corso, Alberto, et al. "Protease-Cleavable Linkers Modulate the Anticancer Activity of Noninternalizing Antibody-Drug Conjugates." Bioconjugate Chemistry, vol. 28, no. 7, 2017, pp. 1826-1833. https://doi.org/10.1021/acs.bioconjchem.7b00304

Doronina, Svetlana O., et al. "Enhanced Activity of Monomethylauristatin F through Monoclonal Antibody Delivery: Effects of Linker Technology on Efficacy and Toxicity." Bioconjugate Chemistry, vol. 17, no. 1, 2006, pp. 114-124. https://doi.org/10.1021/bc0502917

Doronina, Svetlana O., et al. "Development of Potent Monoclonal Antibody Auristatin Conjugates for Cancer Therapy." Nature Biotechnology, vol. 21, no. 7, 2003, pp. 778-784. https://doi.org/10.1038/nbt832

Goldberg, Shalom D., et al. "Trastuzumab-MMAU Antibody-Auristatin Conjugates: Valine-Glucoserine Linker with Stabilized Maleimide Conjugation Improves In Vivo Efficacy and Tolerability." Molecular Cancer Therapeutics, vol. 23, no. 11, 2024, pp. 1530-1543. https://doi.org/10.1158/1535-7163.MCT-23-0591

Gray, Meghan E., et al. "A Comparison of the Activity, Lysosomal Stability, and Efficacy of Legumain-Cleavable and Cathepsin-Cleavable ADC Linkers." Xenobiotica, vol. 54, no. 8, 2024, pp. 458-468. https://doi.org/10.1080/00498254.2024.2352051

Hu, Xinyue, et al. "Synthesis, Characterization, and Targeted Chemotherapy of SCT200-Linker-Monomethyl Auristatin E Conjugates." European Journal of Medicinal Chemistry, vol. 216, 2021, article 113297. https://doi.org/10.1016/j.ejmech.2021.113297

Kumar, Amit, et al. “Synthesis of a Heterotrifunctional Linker for the Site-Specific Preparation of Antibody-Drug Conjugates with Two Distinct Warheads.” Bioorganic & Medicinal Chemistry Letters, vol. 28, nos. 23-24, 2018, pp. 3617-3621. https://doi.org/10.1016/j.bmcl.2018.10.043

Li, Xiaoyi, et al. “Benzyl Ammonium Carbamates Undergo Two-Step Linker Cleavage and Improve the Properties of Antibody Conjugates.” Angewandte Chemie International Edition, vol. 64, no. 6, 2025, article e202417651. https://doi.org/10.1002/anie.202417651

Liu, Haidong, et al. “Expanding the Therapeutic Index of MMAE-Based Antibody Drug Conjugates (ADCs) with a Novel Linker System (LD343).” Journal for ImmunoTherapy of Cancer, 2023. https://doi.org/10.1136/jitc-2023-sitc2023.1407

Long, Jing, et al. “PEGylation of Dipeptide Linker Improves Therapeutic Index and Pharmacokinetics of Antibody-Drug Conjugates.” Bioconjugate Chemistry, vol. 36, no. 2, 2025, pp. 179-189. https://doi.org/10.1021/acs.bioconjchem.4c00392

McKertish, Candice Maria, and Veysel Kayser. “A Novel Dual-Payload ADC for the Treatment of HER2+ Breast and Colon Cancer.” Pharmaceutics, vol. 15, no. 8, 2023, article 2020. https://doi.org/10.3390/pharmaceutics15082020

Ochtrop, Philipp, et al. “Compact Hydrophilic Electrophiles Enable Highly Efficacious High DAR ADCs with Excellent In Vivo PK Profile.” Chemical Science, vol. 14, no. 9, 2023, pp. 2259-2266. https://doi.org/10.1039/D2SC05678J

Wang, Yanming, et al. “Antibody-Drug Conjugate Using Ionized Cys-Linker-MMAE as the Potent Payload Shows Optimal Therapeutic Safety.” Cancers, vol. 12, no. 3, 2020, article 744. https://doi.org/10.3390/cancers12030744

Wang, Yanming, et al. “Development and Properties of Valine-Alanine Based Antibody-Drug Conjugates with Monomethyl Auristatin E as the Potent Payload.” International Journal of Molecular Sciences, vol. 18, no. 9, 2017, article 1860. https://doi.org/10.3390/ijms18091860

Wang, Yanming, et al. “Development of Applicable Thiol-Linked Antibody-Drug Conjugates with Improved Stability and Therapeutic Index.” Drug Delivery, vol. 29, no. 1, 2022, pp. 754-766. https://doi.org/10.1080/10717544.2022.2039807

Wang, Yanming, et al. “From Prodrug to Pro-Prodrug: Hypoxia-Sensitive Antibody-Drug Conjugates.” Signal Transduction and Targeted Therapy, vol. 7, 2022, article 20. https://doi.org/10.1038/s41392-021-00833-8

Watanabe, Tomohiro, et al. “Homogeneous Dual-Payload Antibody-Drug Conjugates Produced by Combined Distinct Conjugation Strategies.” ACS Medicinal Chemistry Letters, vol. 16, no. 7, 2025, pp. 1334-1339. https://doi.org/10.1021/acsmedchemlett.5c00209

Xiao, Dian, et al. “Azobenzene-Based Linker Strategy for Selective Activation of Antibody-Drug Conjugates.” Angewandte Chemie International Edition, vol. 63, no. 16, 2024, article e202310318. https://doi.org/10.1002/anie.202310318

Zhang, Donglu, et al. “Exposure-Efficacy Analysis of Antibody-Drug Conjugates Delivering an Excessive Level of Payload to Tissues.” Drug Metabolism and Disposition, vol. 47, no. 10, 2019, pp. 1146-1155. https://doi.org/10.1124/dmd.119.087023

Mental Health and Remembering Dreams: A Review on Applying Dream Recall Frequency Analysis to the Transdiagnostic Method By Emily Elizabeth Yam

Keywords

Neuroscience; Clinical and Developmental Psychiatry; Mental Health; Dreaming

Abstract

The neural correlates and neuropsychological function of dreaming are yet to be fully understood. Following advancements in functional imaging methods for brain lesion studies, perspectives in neuroscience on the fundamental role of dreams in emotional processing have begun to shift from theoretical neuropsychological frameworks to objective studies of sleep architecture. Now, new research is being conducted on the relationship between dream recall frequency, psychosomatics, and mental health, shedding light on the manifestation of emotional dysregulation in dreaming. Continued research into dreaming may unravel new insights on the interplay between emotion, experience, and memory in the brain not only during sleep, but throughout differing states of consciousness. There is promise for a more informed understanding of dream phenomenology to be applied for the benefit of neuroscientific methods of mental health diagnosis in clinical practices, such as psychotherapy. This review explores the potential for dream recall frequency to be introduced to the transdiagnostic method in clinical psychiatry as a quantifiable indicator of mental health disorder. The findings suggest that notably low dream recall frequency, although at times attributable to differences in memory ability and cognitive functions, can be interpreted to indicate mental disorder. Moreover, a significant limitation to analysing dream recall frequency is the lack of reliability as dream recall is subjective and unable to be externally verified.

Introduction

A dream is a stimulus-independent recreation of reality created through the continuous combination of memory and imagination. During sleep, high-activity processes in the default mode network (DMN), a multi-regional system in the brain responsible for memory recall and unconstrained emotional thought, generate the experience of dreaming.(Sacco; Zellner) Dreams are observed and recalled most vividly when they occur during rapid eye movement (REM) sleep, when brain regions responsible for emotional perception are activated, and those for constructing memory, logical reasoning, directed thought, and self-awareness are deactivated.(Hobson) Dreams can also occur during NREM sleep which involves suspended sensory activity (reduced sensory awareness) and a lack of responsiveness toward outward stimuli, although notably less vivid and with a lower frequency of recall.(Wamsley)

Currently, the neural correlates of dreaming (specific, observable activity patterns in the DMN which are causally related to the occurrence of dreams) have not been identified; however, numerous studies attempt to relate regional activity patterns in the DMN, and more recently, dream phenomenology with underlying brain activity during sleep.(Nir and Tononi) Major DMN

regions are now understood to have their respective roles in dream synthesis: medial temporal lobe structures such as the hippocampus are responsible for memory recall, imagination, and simulating future events; the mPFC (medial prefrontal cortex) is associated with dream recall frequency (how frequently the presence and content of dreams are recalled after awakening) and visualisation; the PCC (posterior cingulate cortex) and precuneus create the perceptual and narrative imagery of dreams; and the inferior parietal lobule processes spatial and sensory awareness.(Wamsley)

On the other hand, various neurocognitive theories derive such findings to hypothesise the origin and function of dreaming. The frequently mentioned activation-synthesis model, proposed in 1977 by Hobson and McCarley, explains dreaming as a result of forebrain stimulation catalysed by spontaneous electrical impulses from the brainstem, serving to maintain sleep during high brain activity.(Hobson and McCarley; Ruby) Alternatively, the threat-simulation theory by Revonsuo in 2002 speculates that dreaming is an evolutionary defense mechanism which simulates physically and psychologically traumatic events in order to rehearse threat response and avoidance. Furthermore, the information processing theory adopted by modern neuroscience and emotion regulation theory developed by Cartwright and further discussed by Scarpelli et al. both suggest that dreams are part of a process of organising and processing information and emotion in the brain during sleep.(Wamsley and Stickgold; Scarpelli et al.)

A majority of theories acknowledge how dreams relate to experiences during the waking state (when an individual is not asleep) implying a shared understanding of its notable role in emotional processing.(Schwartz and Maquet) Subsequently, the potential of certain features of dreams to imply signs of maladaptive emotional regulation could be assumed. While the Freudian psychoanalytic theory suggests that literal images, considered to be manifestations of unconscious emotions, are more indicative of mental state, others suggest that objective sleep architecture, notably dream recall frequency, is more reliable in indicating abnormalities.(Schwartz and Maquet; Merced)

Within traditional clinical psychiatry, mental health diagnosis largely utilises symptomatic analysis in order to pinpoint a specific illness. However, dream analysis alone cannot identify a particular disorder. The transdiagnostic method provides an alternative approach, focusing on underlying general emotional dysregulation which overlap across several disorders.(Fusar‐Poli)

This review discusses the current neuropsychological evidence for the emotional significance of dreams, how mental disorders can be diagnosed using dream recall frequency analysis through the transdiagnostic method, the limitations to the approach, and directions for future studies.

The Neuropsychology of Dreaming and Mental Health

The prevalence of emotional content in dreams is proven evident by studies on how neurological functions in the DMN and perception-based experiences during dreaming reflect

emotional processing. Neuroimaging data on regional brain activity gathered during REM sleep are now analysed to imply damage in regions of the brain; D'Agostino et al. suggests that they can also be used to interpret the neural basis of dreams.(Schwartz and Maquet; D'Agostino et al.) An empirical study by Zellner and Margaret R. on brain lesions in DMN regions respectively associates the elimination of dreams with damage to the ventromedial prefrontal cortex (vmPFC), a region of the mPFC heavily linked to the hippocampus and amygdala responsible for self-related emotional processing such as personal narrative construction and self-reflection, suggesting that dream presence itself is heavily linked to emotional motivations.(Zellner; Wamsley) Reviews by Lara-Carrasco & Jessica and Ross Levin & Tore Nielsen on REM awakenings from self-reported dreams conclude on the non-negligable function of nightmares (referring to vividly recalled dreams involving or evoked by strong emotional reactions during wakefulness) in the emotional regulation system.(Levin and Nielsen; Lara-Carrasco) While direct content may be associated with psychiatric conditions such as post-traumatic stress disorder (PTSD) or schizophrenia, the contrast between the neutral to positive content of dreams mid-depression and negative content post-depression found in studies such as Barrett & Loeffler and Miller, shows that nightmares can often provide a means of consolation after bereavement or indicators of emotional dysregulation to be targeted with improved therapeutic coping strategies.(Lara-Carrasco; Moffitt et al.) Furthermore, Scarpelli et al., which analyses awakenings from REM sleep and electroencephalography (EEG; a non-invasive neuroimaging method which records spontaneous electrical activity in the brain) correlates to associate neural correlates of emotional processing with dream presence, finds a relation between repeated content and perception of experiences in waking life, even proposing that dreams serve to modulate traumatic memory as a coping mechanism toward negative or frightening events.(Scarpelli et al.)

In sum, emotional dysregulation often influences and manifests in dreams as nightmares. However, direct dream content interpretation is infeasible due to the significant lack of concrete research within neuroscience.(Ruby) This paper discusses dream recall frequency as a plausible alternative for its quantifiability as well as valid correlation with psychosomatic symptoms stemming from mental disorder.

Applying Dream Recall Frequency to the Transdiagnostic Method

Dreams nearly always occur during sleep; yet, it is typical for individuals to retain little memory of them after.(Sacco) Dream recall frequency is a measurement of the ability to remember dreams upon awakening, which can be influenced by various factors including mental state, cognitive function, and psychosomatic condition. While the majority of studies demonstrate high dream recall frequency in healthy individuals when woken up at the REM sleep stage, Sacco and Diego Dal. presents several hypotheses to address a potential correlation between lower dream recall frequency and patients with psychosomatic disorders.(Sacco) The review highlights this through a summary of research by Ilias et al. on patients of psychosomatic dermatitis found a drastically lower rate of dream recall frequency(17.5%) compared to the

average rate in healthy individuals (80%). Conversely, an empirical study by Cartwright et al. on the relationship between successful dream recall and lower drop-out rate in psychotherapy finds that higher DRF is related to more positive mental health.(Cartwright et al.)

However, Cory et al. argues self-reported dream recall in patients may be more greatly associated with differences in memory ability as opposed to mental conditions, highlighting a major limitation in current research: the lack of reliability in dream recall frequency measurement.(Cory et al.)

Limitations

The ability for dream analyses to obtain accurate data is largely restricted by the fact that dreams are largely subjective experiences which individuals do not always have consciousness during and coherent memory after.(Ruby) Although, recent developments in brain state measurements such as fMRI (functional magnetic resonance imaging) have opened the door to accessing subjective awareness (such as perception and intention) during unconscious states to achieve cerebral communication and interpretation of mental state in individuals without consciousness or overt external behavior and speech.(Evers and Sigman) There is an emerging possibility that further research to identify reliable neuropsychological correlates of dreaming could lead to achieving direct dream content interpretation during dreaming.(Evers and Sigman)

However, such intervention would pose ethical considerations regarding privacy. As dreams are a significant process in emotional processing can reveal deeply repressed emotions and experiences, often out of the dreaming individual's control. Arguably, regardless of consent, the lack of volition during dreaming constitutes that such consent alone cannot justify accidental invasions of privacy which may have been outside of their consideration. Presently, this concern is not immediate as the neural correlates of dreaming are not yet understood, meaning studies of dreams are still subjective to the dreaming individual, who has full ability to decide what they share of their experience.(Ruby) Nevertheless, in the coming decades it will be crucial to address such potential ethical implications alongside the innovation of new technologies in neuroscience.

Conclusion

To conclude, dreaming is a significant process in emotional regulation which neural basis and content are heavily linked to and influenced by emotion.(Schwartz and Maquet; Zellner; Wamsley) Dream recall frequency, influenced by psychosomatic disorder, can be recorded similarly to conscious behaviors in the waking state, where low frequencies are treated as an indicator of mental disorder.(Sacco) Importantly, dream recall frequency measurement is better implemented as a transdiagnostic method, as specific neural correlates for mental health conditions in dreaming have not yet been studied, let alone identified.(Ruby)

The neural basis of dreams was discussed, covering literature on DMN regions, the hippocampus and vmPFC, responsible for the synthesis of recalled memory and imagination.(Zellner; Wamsley) Moreover, a summary was provided on the neuropsychology of dreams and their role in emotional regulation.(Levin and Nielsen; Lara-Carrasco) Primarily, the

relationship between dream recall frequency and mental disorder was explored regarding how low dream recall frequency is correlated with psychosomatic disorder, physical illness linked to negative mental state, and high dream recall frequency is associated with positive mental health.(Sacco; Cartwright et al.)

Finally, the review outlined limitations to dream studies and ethical considerations regarding the privacy of dreams.(Evers and Sigman; Cory et al.) Overall, a key direction for future research is to identify reliable neural correlates for dreaming. Recent decades have shown a particular interest in studying dream consciousness for its notable similarity to waking consciousness (aside from a few significant variances in cognition), and therefore potential to yield insights into the neurology of consciousness itself.(Nir and Tononi) Not only would such advancement in understanding allow the dream presence and dream recall to be detected and externally validated, but also guide clinical practices to introduce and develop methods for identifying consciousness within patients who are unable to communicate due to a restricted or comatose state.

Acknowledgements

The author would like to express their gratitude to Dr. Kimberly Rose Clark and Monika Rybak for their mentorship during their research period. The author further extends their appreciation to Indigo Research for further support during their enrollment in their Indigo Research Intensive Seasonal (IRIS) Program.

Data Availability

No new data has been analysed or generated in support of this paper.

Works Cited

Cory, Thomas L., et al. "Predicting the Frequency of Dream Recall." Journal of Abnormal Psychology [US], vol. 84, no. 3, 1975, pp. 261–66. APA PsycNet, https://doi.org/10.1037/h0076653.

D, Cartwright, Rosalind, et al. "Focusing on Dreams: A Preparation Program for Psychotherapy." DeepDyve, 1 Mar. 1980, https://www.deepdyve.com/lp/american-medical-association/focusing-on-dreams-a-preparation-program-for-psychotherapy-0pJpfUn6If.

D'Agostino, Armando, et al. "Dreaming and the Neurobiology of Self: Recent Advances and Implications for Psychiatry." Frontiers in Psychology, vol. 4, Sept. 2013. Frontiers, https://doi.org/10.3389/fpsyg.2013.00680.

Evers, Kathinka, and Mariano Sigman. "Possibilities and Limits of Mind-Reading: A Neurophilosophical Perspective." Consciousness and Cognition, vol. 22, no. 3, Sept. 2013, pp. 887–97. ScienceDirect, https://doi.org/10.1016/j.concog.2013.05.011.

Hobson, J. Allan. Dreaming: An Introduction to the Science of Sleep. Oxford University Press, 2002.

Hobson, J. Allan, and Robert W. McCarley. "The Brain as a Dream State Generator: An Activation-Synthesis Hypothesis of the Dream Process." The American Journal of Psychiatry [US], vol. 134, no. 12, 1977, pp. 1335–48. APA PsycNet, https://doi.org/10.1176/ajp.134.12.1335.

Lara-Carrasco, Jessica. Nightmares, Dreaming, and Emotion Regulation: A Review. Jan. 2007. www.academia.edu, https://www.academia.edu/28137606/Nightmares_dreaming_and_emotion_regulation_A_review.

Levin, Ross, and Tore Nielsen. "Nightmares, Bad Dreams, and Emotion Dysregulation: A Review and New Neurocognitive Model of Dreaming." Current Directions in Psychological Science, vol. 18, no. 2, 2009, pp. 84–88.

Merced, Matthew. "Dreaming: Physiological Sources, Biological Functions, Psychological Implications." The Journal of Mind and Behavior, vol. 33, no. 3/4, 2012, pp. 173–93.

Moffitt, Alan, et al. The Functions of Dreaming. State University of New York Press, 1993.

Nir, Yuval, and Giulio Tononi. "Dreaming and the Brain: From Phenomenology to Neurophysiology." Trends in Cognitive Sciences, vol. 14, no. 2, Feb. 2010, p. 88. PubMed Central, https://doi.org/10.1016/j.tics.2009.12.001.

Ruby, Perrine Marie. "The Neural Correlates of Dreaming Have Not Been Identified Yet. Commentary on 'The Neural Correlates of Dreaming. Nat Neurosci. 2017.'" Frontiers in Neuroscience, vol. 14, Oct. 2020. Frontiers, https://doi.org/10.3389/fnins.2020.585470.

Sacco, Diego Dal. "Dream Recall Frequency and Psychosomatics." Acta Biomedica Atenei Parmensis, vol. 93, no. 2, May 2022, pp. e2022046–e2022046. www.mattioli1885journals.com, https://doi.org/10.23750/abm.v93i2.11218.

Scarpelli, Serena, et al. “The Functional Role of Dreaming in Emotional Processes.” Frontiers in Psychology, vol. 10, Mar. 2019. Frontiers, https://doi.org/10.3389/fpsyg.2019.00459.

Schwartz, Sophie, and Pierre Maquet. “Sleep Imaging and the Neuro-Psychological Assessment of Dreams.” Trends in Cognitive Sciences, vol. 6, no. 1, Jan. 2002, pp. 23–30. www.cell.com, https://doi.org/10.1016/S1364-6613(00)01818-0.

Transdiagnostic Psychiatry: A Systematic Review - Fusar‐Poli - 2019 - World Psychiatry - Wiley Online Library. https://onlinelibrary.wiley.com/doi/full/10.1002/wps.20631. Accessed 6 Sept. 2026.

Wamsley, Erin J. “How the Brain Constructs Dreams.” eLife, vol. 9, June 2020, p. e58874. eLife, https://doi.org/10.7554/eLife.58874.

Wamsley, Erin J., and Robert Stickgold. “Dreaming and Offline Memory Processing.” Current Biology : CB, vol. 20, no. 23, Dec. 2010, pp. R1010–13. PubMed Central, https://doi.org/10.1016/j.cub.2010.10.045.

Zellner, Margaret R. “Dreaming and the Default Mode Network: Some Psychoanalytic Notes.” Contemporary Psychoanalysis, vol. 49, no. 2, Apr. 2013, pp. 226–32. Taylor and Francis+NEJM, https://doi.org/10.1080/00107530.2013.10746548.

Differences between Early-Onset Alzheimer's Disease and Late-Onset Alzheimer's Disease in Attention Networks of the Brain By William Mao

Abstract

Alzheimer's Disease (AD), the most prevalent form of dementia, primarily targets elderly populations. There are two main types of AD: Early-Onset Alzheimer's Disease (EOAD), affecting people <65, and Late-Onset Alzheimer's Disease (LOAD), affecting those >=65. While all AD patients exhibit similar deficiencies in delayed memory, the areas of the brain affected differ. Additionally, EOAD patients have often performed worse in tests evaluating processing speed and executive function. Therefore, it's crucial to analyze how EOAD and LOAD differ in their impact on attention. Previous studies have examined the deterioration of brain regions and cognitive patterns in EOAD versus LOAD cohorts using MRI and cognitive tests. This review aims to link and organize these deterioration patterns in the context of the Posner attention model to better understand the neuropathological differences between the two subtypes. Current cognitive and structural evidence suggests that EOAD damages the executive control and orienting systems more severely. However, subcortical evidence surrounding alerting doesn't align with associated cognitive deficiencies due to a lack of evidence, so further conclusions about alerting are unclear. This signifies the possibility that distinct networks are affected when AD manifests at different ages, allowing the possibility to more accurately diagnose patients.

Keywords Late-onset, Early-onset, Alzheimer's Disease, Attention Systems

Introduction

Late-Onset Alzheimer's Disease (LOAD) is the predominant form of Alzheimer's Disease (AD), and its symptoms are characterized by the common loss of memory that is often associated with AD. However, less well-known is Early-Onset Alzheimer's Disease (EOAD) and how its pathology fits into the larger puzzle of AD pathology. Current research suggests that the areas affected by EOAD are primarily less concentrated in the hippocampal region than in LOAD, but rather the main regions of atrophy lie in the precuneus and posterior parietal cortex (Möller et al., 2013). The precuneus is a default mode network hub located in the parietal lobe that connects various functions such as self-awareness and consciousness, while the posterior parietal cortex, another parietal region, functions in spatial reasoning and attention. In addition, through analysis of structural MRI of EOAD patients, scientists have dubbed the areas of the temporal, parietal, posterior cingulate complexes, middle frontal gyrus, and superior frontal gyrus as EOAD "signature regions" (Touroutoglou et al., 2023). Touroutoglou also notes that these "signature regions" encompass five major networks of the brain, including the frontoparietal network, which is heavily involved in executive function, and the dorsal attention network, which operates under orienting. However, although these regions are strongly connected to attention, there is still a gap in the connection between these areas of the brain and how they fit into the framework of a model of attention. Posner's model of attention was selected

for this review because it breaks attention down into distinct systems, and each of the systems mostly operate different parts of the brain as well as having their own neuromodulation. Therefore, this review examines how the patterns of cortical and subcortical regions map onto Posner's model to explain differences in attention deficiencies between EOAD and LOAD patients.

Because EOAD and LOAD differ pathologically and exhibit different cognitive deficiencies, standard memory evaluations are not the most favorable for all EOAD and LOAD patients alike. This generalization leads to diagnostic delays (Mendez, M.F, 2019). Additionally, understanding the difference in their neuromodulation can transform how AD is treated because EOAD and LOAD may respond differently to cholinergic versus noradrenergic treatments (Falgás et al., 2024). This will be specifically discussed when examining the alerting network since imaging will be targeted toward subcortical structures involved in network-specific modulation. These gray areas primarily attribute to the missing link between how cortical and subcortical data connect together in the attention model. Tests evaluating attention among patients also do not distinguish between impacts on individual systems but rather generalize it globally.

Based on previous research, interpreting the patterns of the structural and cognitive evidence together on Posner's three networks of attention reveals the differences in the pathology of EOAD and LOAD and solidifies concepts. These deficits in attention may contribute to the observed effects of memory loss shared between EOAD and LOAD patients by affecting the brain's ability to encode memories. Specifically, EOAD exhibits a significantly larger impact on orienting and executive control due to heavy atrophy in parietal and frontal areas, while alerting is also affected by subcortical atrophy in the locus coeruleus, but its implications in cognitive behavior aren't conclusive. This will be organized by introducing Posner's model of attention in the first section and how its three networks contribute to a global understanding of attention, as well as the neuromodulation associated with each network. Cognitive differences exhibited by EOAD and LOAD patients will be discussed in the second section to understand general differences between the two. The third section will examine the cortical and subcortical evidence from MRI scans and Neuromelanin-sensitive TSE. This evidence will be reorganized and mapped onto Posner's three networks in the fourth section to analyze system-specific patterns and implications. The last section will address the unresolved gaps in the understanding of these networks, future directions, and the conclusion.

Posner's Model of Attention: A Three Network System

In 2012, Posner and Petersen described a modern framework for attention consisting of three different networks: the alerting, orienting, and executive control networks (Petersen, S.E., & M.I. P12). This model of attention is specifically noteworthy because it decomposes attention into three networks, each with its own anatomy and neuromodulation. Because of this convenience, each network can be tested through the Attention Network Test (ANT), a computer-based assessment evaluating each network through various visual cues.

The alerting network involves phasic alertness, the rapid increase in vigilance created by a warning signal, and tonic alertness, sustained alertness over longer periods of time. This is primarily carried out by the locus coeruleus (LC), a small nucleus located in the brainstem, the right frontoparietal cortex, and the thalamus. Additionally, the LC is strictly made up of noradrenergic neurons, which in turn means that the alerting network is primarily influenced by fluctuations in the neurotransmitter noradrenaline.

The orienting network dictates where attention is directed by prioritizing input. This is divided into the dorsal attention network (DAN) and the ventral attention network (VAN). The DAN is composed of the intraparietal sulcus (IPS), superior parietal lobe, and frontal eye fields (FEF), functioning in sensory-motor coordination, spatial orientation, and visual attention to generate a goal-directed response (top-down). In contrast, the VAN (bottom-up) comprises the temporoparietal junction and ventral frontal cortex and functions in orienting to a stimulus. Globally, the orienting network is primarily cholinergic, modulated by acetylcholine produced by the nucleus basalis of Meynert (Liu et al., 2015)

Operating through top-down control, the executive control network consists of two distinct networks: the cingulo-opercular network and the frontoparietal network. The cingulo-opercular network is composed of the dorsal anterior cingulate cortex (dAAC), medial frontal lobe, and anterior insula and maintains current goals while adjusting to feedback and unexpected outcomes. In contrast, the frontoparietal network, consisting of the lateral prefrontal cortex, involves fast-paced task initiation, switching, and adjustments. Additionally, both networks demonstrate an overlapping function in resolving conflict and are dopaminergically modulated by areas in the midbrain. However, since they are anatomically separate, they may be differentially affected in EOAD and LOAD.

Although the three networks are distinct systems, they are not completely separate. Structurally, the frontoparietal networks function both in orienting and executive control. Additionally, Mahoney et al. (2010) found a correlation between alerting and orienting ($r = 0.27$) utilizing the Attention Network Test (ANT) among elderly patients. Therefore, regional atrophy must be considered in the context of all possible systems of attention.

Cognitive Differences between EOAD and LOAD

In prior research, specific cognitive tests among EOAD and LOAD cohorts have revealed significant differences in baseline function of EOAD and LOAD patients. These tests aimed to evaluate specific domain functions using neuropsychological tests between different cohorts. Although many parameters are not ANT tests, the nature of the tests can be interpreted in the context of the networks. Due to this limitation, the cognitive evidence might be non-uniform and other factors are needed to explain the inconsistencies.

One example of examining specific cognitive parameters between EOAD and LOAD cohorts was conducted by Sá et al. (2012). Specifically, Sá et al. examined these relationships using Battery of the Lisbon for the Assessment of Dementia (BLAD) subtests in EOAD (n=109) and LOAD (n=171) cohorts. Sá et al. noted significant differences in two subtests — confirmed

with Bonferroni correction: Orientation ($p=0.001$) and Praxis ($p<0.001$). EOAD patients performed significantly worse in Praxis (11.47 ± 0.98) and better in Orientation (11.59 ± 3.34) as compared to the LOAD cohort (Orientation = 10.56 ± 2.66 and Praxis = 11.90 ± 0.46). Sá suggests that these results may signal that EOAD causes the left posterior hemisphere to be more prone to damage and atrophy. However, it is important to note that neither of those parameters are inherently related to attention networks. Therefore, this study needs to be evaluated in the context of other like-studies.

In a similar study with larger EOAD and LOAD cohorts by Hammers et al. (2025), Z-score analysis found significantly lower performances of EOAD patients in processing speed/attention ($p<0.001$), episodic delayed memory ($p = 0.006$), visuospatial skills ($p<0.001$), executive functioning ($p<0.001$), and episodic immediate memory ($p<0.001$). Specifically noteworthy is the substantially large effect size, η^2, of the visuospatial skills domain ($\eta^2 = 0.151$). Unfortunately, because this domain relies on MoCA Clocks and MoCA Cube measures, specific conclusions relating to specific attentional networks can't be drawn. However, brain lesion evidence has suggested that these tests may activate right parietal regions (Tranel et al., 2008). These regions are closely associated with DAN networks and may hint toward a connection with the orienting network. In a follow-up study by Hammers et al. (2026), the domains are broken down into more specific parameters. Data from Trail Making Test Part B, a test evaluating visual search speed and cognitive flexibility, confirm greater executive control deficiencies in EOAD patients. This suggests greater dysfunction in the frontal-parietal network (executive control) due to the nature of the alternating stimuli of the parameter.

In the context of the global attention network, Sá et al. and Hammers et al. fundamentally disagree on the cognitive differences of EOAD versus LOAD patients. Sá et al. find non-significant attention-associated results in the letter cancellation ($p = 0.788$) and digit span ($p = 0.776$) subtests, which evaluate processing speed and sustained attention. On the other hand, Hammers et al. found significant differences in attention parameters between EOAD (-2.94) and LOAD (-1.27) patients. A potential explanation for the discrepancies between the two studies is Mini-Mental State Examination (MMSE) evaluations of the EOAD versus LOAD cohorts. While the two cohorts of participants in the Sá et al. study had non-significant differences in cognitive impairment ($p = 0.358$), the two groups in the Hammer et al. study exhibited significant differences, with EOAD (21.36) having greater cognitive impairment than LOAD (25.11) ($p<0.001$). This makes it difficult to determine the effect of the stratification of cognitive impairment between the two cohorts and requires further clarification with cortical evidence by Katsumi et al. (2025). Another factor that may have contributed to this discrepancy is the difference in attention parameters and the non-uniform nature of the tests.

Although Sá et al. and Hammers et al. illustrate important patterns in cognitive tests that evaluate parameters overlapping or encompassing attention, orienting, and executive control, studies that compare cognitive deficiencies of the decomposed networks of Posner's model of attention directly are scarce. However, the Hammers et al. (2025) and Hammers et al. (2026) studies strongly suggest that orienting and executive control are impacted more severely by

EOAD than LOAD in more moderate cases of EOAD.

Cortical and Subcortical Evidence

Through extensive imaging analysis of brain atrophy in AD patients, connections to associated alerting, orienting, and executive control networks can be drawn. Several studies have analyzed regions specific to EOAD cohorts in relation to LOAD that are more atrophied.

Several studies have examined cortical evidence on regions associated with orienting and executive networks concentrated in parietal, DAN, and VAN regions. Using a W-scoring analysis, Katsumi et al. (2025) analyzed structural MRI data between EOAD and LOAD patients and discovered that the data suggested greater damage in the inferior parietal lobule, precuneus, and caudal middle frontal gyrus — regions associated with spatial and cognitive control — amongst EOAD patients while LOAD patients exhibited more significant damage to the entorhinal cortex and temporal pole, two regions that function in memory. Additionally, significant differences between the dorsal attention, ventral attention network, and frontoparietal networks were also found. Katsumi et al. concluded that these patterns were double dissociable, rendering the possibility that the differences in MMSE (greater cognitive impairment in the EOAD sample group) accounted for these differences to be unlikely. In general, Katsumi et al. suggest that EOAD and LOAD exhibit different specialized patterns of atrophy, further supporting possible explanations for the cognitive differences. Building on the Katsumi et al. study, atrophy in the precuneus in EOAD and medial temporal atrophy in LOAD was also found in a study conducted by Möller et al. (2013) through MRI imaging. Both of these studies examine and note regions associated with executive control and orienting but don't provide specific data on alerting regions and structure. Therefore, analyzing the alerting network will require deeper subcortical scanning.

Examining the locus coeruleus (LC), a major nucleus in the alerting network, Falgàs et al. (2024) used neuromelanin-sensitive MRI imaging. From the MRI statistics, it was determined that the distinction between EOAD and LOAD significantly predicted worse LC integrity and lower LC volume among EOAD patients. It was also depicted that the AD stage and impairment level didn't significantly impact LC volume ($p = 0.659$), which attributes atrophy mainly to EOAD compared to LOAD. Interestingly, when examining the effect of this atrophy on neuromodulation, Falgàs et al. found significant increases in noradrenaline in AD patients with lower LC integrity ($p = 0.040$), an unexpected result considering that the LC, the major hub of noradrenaline production, has increased atrophy in AD patients. Additionally, CSF noradrenaline levels were consistent in both EOAD and LOAD patients, further contributing to the odd pattern. Therefore, Falgàs et al. propose that synaptic structural integrity may be more useful for predicting the severity of symptoms in EOAD than neurotransmitter levels. Overall, Falgàs et al. determined that LC integrity significantly predicted greater Neuropsychiatric Inventory (NPI) total score (worse neuropsychiatric symptoms), suggesting that the significant difference ($p = 0.028$) between EOAD(13.3 ± 18.9) and LOAD(9.9 ± 13.7) scores could have a causal link with the type of AD. This study illustrates the effect of EOAD on the LC, a major structure in the

alerting network, and contributes to the idea of specialized EOAD atrophy patterns.

However, what's more comprehensive than piecing together these studies is found in a Longitudinal Early-Onset Alzheimer's Disease (LEADS) study conducted by Touroutoglou et al. (2023) using structural MRI to search for common patterns of atrophy among EOAD patients. Specifically, Touroutoglou et al. identified 9 "EOAD signature regions", categorizing atrophy through mean thickness difference. Parietal regions exhibited the greatest mean thickness difference (mm): inferior parietal (-0.3), precuneus (-0.3), and superior parietal (-0.25). This is followed by lateral temporal regions: Caudal lateral temporal (-0.29) and Mid lateral temporal (-0.24). The frontal regions exhibited the least atrophy, with <0.20 mm of mean difference. Touroutoglou et al.'s depiction of "EOAD signature regions" strengthens the pattern of atrophy in the parietal regions as well as introduces atrophy in major local default mode networks (DMN): the precuneus and the posterior cingulate.

In conclusion, several imaging scans have determined that EOAD exhibits specific areas of greater atrophy than LOAD, suggesting that the type of onset-AD itself influences atrophy despite age differences.

Mapping Imaging and Cognitive Evidence onto Posner's Model

Table 1: Mapping of evidence onto the three components of Posner's model of attention in EOAD and LOAD.

Posner Model Network	Associated Structures	Neuromodulator	Cognitive Evidence	Cortical and Subcortical Evidence
Alerting	Right frontoparietal networks, thalamus, LC	Noradrenaline	Very limited	Lower LC integrity ($\beta = 0.27$, $p = .016$) and volume ($\beta = 0.40$, $p < .001$) in EOAD (Falgàs et al., 2024)
Orienting	IPS, FEF, and superior parietal lobe	Acetylcholine	Worse visuospatial skills among EOAD patients ($\eta^2 = 0.151$; Hammers et al., 2025; see also Smits et al., 2012)	Increased EOAD atrophy in the parietal and precuneus regions (Katsumi et al., 2025; Möller et al., 2013). Increased DAN ($d = -0.31$, $p < .0001$) and VAN ($d = -0.17$, $p = .03$) atrophy (Katsumi et al., 2025)
Executive Control	dACC, medial frontal lobe, anterior insula, lateral prefrontal cortex	Dopamine	Worse performance by EOAD patients on Trail Making Test B ($\eta^2 = 0.067$; Hammers et al., 2025)	Increased EOAD atrophy in the frontoparietal network ($d = -0.57$, $p < .0001$)(Katsumi et al., 2025). Frontal regions included in the EOAD atrophy signature (Touroutoglou et al., 2023)

Note: Although worse visuospatial and Trail Making Test B performance was observed in the Hammers et al. studies, the EOAD cohort included a greater proportion of patients with dementia rather than mild cognitive impairment. Additionally, although frontal regions from part of the EOAD atrophy signature, they showed the weakest effects among the signature regions in Touroutoglou et al.

Compiling current evidence onto the three attention networks is crucial in understanding how each is affected and the true complexities of how the two different types of AD affect attention (Table 1). Therefore, each of the networks needs to be broken down and analyzed separately, considering both imaging evidence and cognitive evidence.

In the orienting networks, clear patterns surrounding parietal and precuneus as well as DAN and VAN atrophy suggest major implications for top-down and bottom-up orienting function. Evidence from Hammers et al. 2025) demonstrates significant results in visuospatial skill. Although the parameters that were tested were not directly orienting, they have been shown to activate right parietal regions (Tranel et al., 2008). This cortical theory about the right parietal regions is strongly confirmed by Katsumi et al. (2025) and Möller et al. (2013), suggesting that these cognitive impairment patterns are a sign of orienting network dysfunction (Table 1).

Similar to the orienting network, executive control is also strongly implicated (Table 1) by the Katsumi et al. (2025) study showing broad patterns of atrophy in EOAD patients, with some limitations. Cognitively, this difference between EOAD and LOAD is supported by the Trail Making Test B parameter in the Hammers et al. (2025) study, which involved alternating stimuli processed by the frontoparietal networks. However, because the frontoparietal networks and cingulo-opercular networks are anatomically separate, this evidence doesn't apply to the latter. Additionally, although frontal regions, which are associated with cingulo-opercular regions, are implicated by Touroutoglou et al. 2023, they are among the weakest signatures compared to others such as parietal regions. Despite this, the orienting network is well-represented on both sides and is more impacted by EOAD.

The alerting network contains the most gaps in knowledge since cognitive testing for alerting is very limited among EOAD versus LOAD cohorts. Subcortically, MRI imaging by Falgàs et al. does suggest both lower LC integrity and volume among EOAD patients. However, as a hub of noradrenaline release, Falgás et al. found constant levels in both EOAD and LOAD, which raises the significance of neuromodulation. Additionally, because prior research has found that the LC is the most reliable neuromodulatory hub in predicting attention in AD compared to others (Plini et al., 2021), the lack of cognitive evidence characterizing these LC deficiencies in EOAD patients compared to LOAD (Table 1) is notable and significant.

Overall, all three of the networks in Posner's model of attention have been implicated by cognitive evidence, structural evidence, or both to be more severe in EOAD, and the compilation of evidence serves to fill gaps in understanding and illustrate differences between the two types of AD. For example, while the cognitive studies from Hammers et al. and Sá et al. couldn't fully explain the impact of impairment severity stratification on attention processes, data from Katsumi et al. (2025) suggest, through imaging analysis, that DAN, VAN, and frontoparietal regions have more significant and greater atrophy at mild dementia (more cognitively impaired) than at the MCI stage for EOAD patients. In addition to EOAD specificity, regions that are more affected in LOAD have also been clearly illustrated (Katsumi et al., 2025). For example, the limbic network, anterior MTL epicenter, and hippocampus were found to be more atrophied in LOAD. These regions, functioning in memory, are also confirmed by the episodic delayed and

immediate memory evaluations ($p = < 0.006$; $\eta^2 = 0.013$ and $p < 0.001$; $\eta^2 = 0.021$, respectively) in the Hammers et al. 2025 study. Although these regions aren't associated with attention, it still suggests different pathological profiles between EOAD and LOAD. This creates a clear double dissociation between EOAD's stronger presence in attention networks and LOAD's impact on memory.

Discussion and Conclusion

Because of a lack of network specific analysis between EOAD and LOAD on attention, this review aimed to answer the question of how current research on EOAD and LOAD differs on their impact on individual attention networks. By connecting cognitive and imaging evidence to observed attention globally and in a decomposed format, many associations can be made. Posner and Petersen (2012) described a three-network model of attention consisting of separate networks that are inherently unique. Using the understanding of each network's function, cognitive evidence was compiled. While the Sá et al. (2012) study found no significant differences in attention-associated parameters, Hammers et al. (2025 & 2026) found significantly greater attention-related cognitive deficiencies in EOAD patients. This raised a question about how cognitive impairment severity affected attention differences. Cortically, Katsumi et al. (2025) and Möller et al. (2013) used statistical analysis and MRI imaging to determine significantly greater atrophy in EOAD patients compared to LOAD in parietal, VAN, and DAN regions. Additionally, data from the Katsumi et al. study suggest significant or greater atrophy in mild dementia cases in EOAD than in MCI cases compared to LOAD patients. This means that the degree of cognitive impairment does contribute to attention differences, which was hinted at in the cognitive analysis. Next, Touroutoglou et al. (2023) further identified specific EOAD signature regions in the parietal, temporal, and frontal areas of the brain. After mapping on Posner's model (Table 1), strong associations with orienting and executive control were made, but it was noted that there was a lack of alerting-specific cognitive evidence. Despite this overall strong global evidence, there are still gaps that future studies should address.

Although this review maps out rough patterns and finds associations between symptoms and brain structure, a comprehensive analysis directly on all 3 networks has yet to be done to analyze differences between EOAD and LOAD patients. Therefore, without further evidence, the degree of impact on each network can't be ranked in relation to each other. In addition, considering the importance of the LC in attentional processes, the gap in tests evaluating alerting also makes it difficult to understand cross-network implications comprehensively. For example, frontal parietal networks have been demonstrated to be atrophied in EOAD (Katsumi et al., 2025; Möller et al., 2013; Touroutoglou et al., 2023), but the lack of cognitive analysis of alerting doesn't differentiate the frontal parietal network's cross-implications in both alerting and executive control.

This review raises two further discussion and future direction points related to stage progression and neuromodulation. Firstly, a greater understanding of how AD stage progression affects overall symptoms between the two groups needs to be studied in order to fully understand

the inconsistencies between the Sá et al. and Hammers et al. studies. Secondly, Falgàs et al. found an increase in noradrenaline in AD patients compared to normal subjects and constant levels between EOAD and LOAD despite worse LC integrity in AD and EOAD compared to their respective groups. Falgàs suggests that this could be a result of compensatory action of remaining neurons, but a true action mechanism can't be determined until further analysis.

In conclusion, after mapping current cognitive and structural evidence in EOAD versus LOAD patients, implications of EOAD versus LOAD on the different attention networks are significant but not yet clearly understood. Researching and further analyzing the difference between EOAD and LOAD is important because reducing memory-related screening may reduce diagnostic delay (Hammers et al., 2026; Mendez, M.F, 2019). Treatment is also implicated because EOAD patients experience particularly disastrous effects and may not respond to antidepressant and anxiolytic medications, which are mainly used to treat LOAD patients (Touroutoglou et al., 2023; Falgàs et al., 2022). This review finds, in the context of Posner's attention model, a divergence between the 2 types of AD in attention networks.

Work Cited

Falgàs, N., Allen, I. E., Spina, S., Grant, H., Piña Escudero, S. D., Merrilees, J., Gearhart, R., Rosen, H. J., Kramer, J. H., Seeley, W. W., Neylan, T. C., Miller, B. L., Rabinovici, G. D., Grinberg, L. T., & Walsh, C. M. (2022). The severity of neuropsychiatric symptoms is higher in early-onset than late-onset Alzheimer's disease. European Journal of Neurology, 29(4), 957–967. https://doi.org/10.1111/ene.15203

Falgàs, N., Peña‐González, M., Val‐Guardiola, A., Pérez‐Millan, A., Guillén, N., Sarto, J., Esteller, D., Bosch, B., Fernández‐Villullas, G., Tort‐Merino, A., Mayà, G., Augé, J. M., Iranzo, A., Balasa, M., Lladó, A., Morales‐Ruiz, M., Bargalló, N., Muñoz‐Moreno, E., Grinberg, L. T., & Sánchez‐Valle, R. (2024). Locus coeruleus integrity and neuropsychiatric symptoms in a cohort of early‐ and late‐onset Alzheimer's disease. Alzheimer's & Dementia, 20(9), 6351–6364. https://doi.org/10.1002/alz.14131

Hammers, D. B., Eloyan, A., Thangarajah, M., Taurone, A., Beckett, L., Gao, S., Polsinelli, A. J., Kirby, K., Dage, J. L., Nudelman, K., Aisen, P., Reman, R., La Joie, R., Lagarde, J., Atri, A., Clark, D., Day, G. S., Duara, R., Graff-Radford, N. R., … Apostolova, L. G. (2025). Differences in baseline cognitive performance between participants with early-onset and late-onset Alzheimer's disease: Comparison of LEADS and ADNI. Alzheimer's & Dementia, 21(1), e14218. https://doi.org/10.1002/alz.14218

Hammers, D. B., Eloyan, A., Thangarajah, M., Taurone, A., Gao, S., Beckett, L., Kirby, K., Dage, J. L., Nudelman, K., Aisen, P., Reman, R., Vemuri, P., La Joie, R., Touroutoglou, A., Atri, A., Clark, D., Day, G. S., Duara, R., Graff-Radford, N. R., … Apostolova, L. G. (2026). Relationship between age and severity of cognitive impairment at diagnosis for early-onset and late-onset Alzheimer's disease: Comparison of LEADS and ADNI. Alzheimer's & Dementia, 22(2), e71160. https://doi.org/10.1002/alz.71160

Katsumi, Y., Touroutoglou, A., Brickhouse, M., Eckbo, R., La Joie, R., Eloyan, A., Nudelman, K. N., Foroud, T. M., Dage, J. L., Carrillo, M. C., Rabinovici, G. D., Apostolova, L. G., & Dickerson, B. C. (2024). Dissociable spatial topography of neurodegeneration in early-onset and late-onset Alzheimer's disease: A head-to-head comparison of MRI-derived atrophy measures between the LEADS and ADNI cohorts. Alzheimer's & Dementia, 20(S2), e091818. https://doi.org/10.1002/alz.091818

Katsumi, Y., Touroutoglou, A., Brickhouse, M., Eloyan, A., Eckbo, R., Zaitsev, A., La Joie, R., Lagarde, J., Schonhaut, D., Thangarajah, M., Taurone, A., Vemuri, P., Jack, C. R., Dage, J. L., Nudelman, K. N. H., Foroud, T., Hammers, D. B., Ghetti, B., Murray, M. E., … Dickerson, B. C. (2025). Dissociable spatial topography of cortical atrophy in early-onset and late-onset Alzheimer's disease: A head-to-head comparison of the LEADS and ADNI cohorts. Alzheimer's & Dementia, 21, e14489. https://doi.org/10.1002/alz.14489

Mahoney, J. R., Verghese, J., Goldin, Y., Lipton, R., & Holtzer, R. (2010). Alerting, orienting, and executive attention in older adults. Journal of the International Neuropsychological Society, 16(5), 877–889. https://doi.org/10.1017/S1355617710000767

Mendez, M. F. (2019). Early-onset Alzheimer disease and its variants. CONTINUUM: Lifelong Learning in Neurology, 25(1), 34–51. https://doi.org/10.1212/con.0000000000000687

Möller, C., Vrenken, H., Jiskoot, L., Versteeg, A., Barkhof, F., Scheltens, P., & van der Flier, W. M. (2013). Different patterns of gray matter atrophy in early- and late-onset Alzheimer's disease. Neurobiology of Aging, 34(8), 2014–2022. https://doi.org/10.1016/j.neurobiolaging.2013.02.013

Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 Years after. Annual Review of Neuroscience, 35(1), 73–89. https://doi.org/10.1146/annurev-neuro-062111-150525

Plini, E. R. G., O'Hanlon, E., Boyle, R., Sibilia, F., Rikhye, G., Kenney, J., Whelan, R., Melnychuk, M. C., Robertson, I. H., & Dockree, P. M. (2021). Examining the role of the noradrenergic locus coeruleus for predicting attention and brain maintenance in healthy old age and disease: An MRI structural study for the Alzheimer's disease neuroimaging initiative. Cells, 10(7), 1829. https://doi.org/10.3390/cells10071829

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. Annual Review of Neuroscience, 13, 25–42.

Sá, F., Pinto, P., Cunha, C., Lemos, R., Letra, L., Simões, M., & Santana, I. (2012). Differences between early and late-onset Alzheimer's disease in neuropsychological tests. Frontiers in Neurology, 3, 81. https://doi.org/10.3389/fneur.2012.00081

Smits, L. L., Pijnenburg, Y. A. L., Koedam, E. L. G. E., van der Vlies, A. E., Reuling, I. E. W., Koene, T., Teunissen, C. E., Scheltens, P., & van der Flier, W. M. (2012). Early onset Alzheimer's disease is associated with a distinct neuropsychological profile. Journal of Alzheimer's Disease, 30(1), 101–108. https://doi.org/10.3233/jad-2012-111934

Touroutoglou, A., Katsumi, Y., Brickhouse, M., Zaitsev, A., Eckbo, R., Aisen, P. S., Beckett, L. A., Dage, J. L., Eloyan, A., Foroud, T., Ghetti, B., Griffin, P., Hammers, D. B., Jack, C. R., Kramer, J. H., Iaccarino, L., La Joie, R., Mundada, N. S., Koeppe, R. A., … Dickerson, B. C. (2023). The sporadic early-onset Alzheimer's disease signature of atrophy: Preliminary findings from the Longitudinal Early-Onset Alzheimer's Disease Study (LEADS) cohort. Alzheimer's & Dementia, 19(Suppl. 9), S74–S88. https://doi.org/10.1002/alz.13466

Tranel, D., Rudrauf, D., Vianna, E. P. M., & Damasio, H. (2008). Does the clock drawing test have focal neuroanatomical correlates? Neuropsychology, 22(5), 553–562. https://doi.org/10.1037/0894-4105.22.5.553

Investigating the Performance of Computational Methods in Predicting Off-target Effects of CRISPR-Cas9: A Review By Ashwath Kannan

Abstract

CRISPR-Cas9 has been one of the most significant breakthroughs in gene editing, allowing scientists to alter the genome with unprecedented precision. However, off-target effects - unintended edits at genomic sites that resemble but do not exactly match the intended target - remain a limiting factor to realizing its full potential, carrying the risk of a wide range of undesirable mutations. A growing number of machine learning (ML) and deep learning (DL) models have been developed to predict these off-target effects computationally, evolving from models built on hand-engineered sequence and structural features toward models that learn their own representations directly from raw sequence data. In this paper, we present a comparative analysis of five models spanning this transition - CRISTA (2017), Elevation (2018), CRISPR-Net (2020), CRISPR-IP (2022), and Crispr-SGRU (2024) - selected to represent a chronological and architectural arc across both eras. We compare these models across six criteria: core algorithm and architecture, sequence encoding scheme, indel handling, class imbalance handling, interpretability, and reported performance on standard metrics. We highlight known challenges affecting the field as a whole, including the lack of a standardized, high-quality benchmark dataset, which limits how directly performance can be compared across studies. This paper aims to provide researchers and students entering this field with an accessible, rigorously sourced synthesis of how off-target prediction methodology has evolved, and where its open challenges remain.

Keywords CRISPR-Cas9, machine learning, deep learning, off-targets

Introduction

The Clustered Regularly Interspaced Short Palindromic Repeats (CRISPR) gene editing system, co-invented in 2012 by Emmanuelle Charpentier and Jennifer Doudna, has become a widely adopted tool for gene editing and the suppression or activation of gene expression (Jinek et al., 2012). Initially, it was discovered in bacteria as a defensive mechanism against phage infection (Barrangou et al., 2007). As DNA repairing mechanisms are evolutionally conserved in all organisms, researchers realised it could be introduced into the mammalian system to perform similar functions. Subsequently, CRISPR/Cas9 has proved itself useful in a variety of fields, including treating human disease, plant genetic engineering and animal disease treatment.

The key feature of CRISPR-Cas9 is that the edit to the gene (mutation) can be pre-designed or custom-made in a lab and targeted to a specific location anywhere in the genome. This means that the editing of the gene can be ‘programmed’. The CRISPR/Cas9 system works by delivering the Cas9 enzyme together with a guide RNA into a cell, which enables precise cutting at a specific location in the genome.[1] This targeted cut creates an

opportunity for genetic material to be edited *in vivo* - existing genes can be removed, and new sequences can be introduced. The guide RNA - or, in engineered type II systems, the single-guide RNA (sgRNA) - is what locates the correct target sequence within the genome and directs Cas9 to cut at that exact site.The Protospacer-Adjacent Motif (PAM), located at the end of the DNA target site, is a nucleobase sequence, typically three-to-five nucleotides long, serving as the binding signal of the Cas protein (Anders et al., 2014).

However, the sgRNA is not always accurate and is susceptible to error. For example, it may lead the Cas9 nuclease to an unintended location resulting in the cutting of a functional gene. Despite its high programmability, the sgRNA is prone to leading the Cas9 to genomic loci sharing a sequence similar (but not the same) as the sgRNA spacer. This results in the sgRNA being ‘off-target’. The consequences of this "off-target" binding are not always predictable: a functional gene may be knocked out unintentionally, mutations may build up elsewhere in the genome, and the resulting biological effects can be difficult to foresee (Huang et al., 2025). Therefore, effectively predicting and reducing off-target effects is key to improving the reliability and applications of this technology.

A number of computational models have been developed that use available datasets to better predict and prevent such off-target effects. These can be broadly categorised into Machine Learning (ML) and Deep Learning (DL) models. Examples of ML models include Random Forest (RF) and Support Vector Machine (SVM), while examples of DL models include Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN) and Transformer-based architectures. These models are constantly evolving as the underlying data improves, both in quality and quantity, and as computational methods advance. This paper discusses the current state of ML and DL models in predicting off-target effects, their effectiveness and the key areas being addressed to improve the performance of the models.

The CRISPR-Cas9 system

The CRISPR-Cas9 system has transformed biomedical research by allowing precise, low-cost, and efficient editing of genomic sequences. Researchers can use it to introduce corrections at specific DNA sites or to activate and silence target genes in cells and whole organisms. This versatility has made the technology central to several experimental workflows, including the rapid production of cell and animal models for disease research, large-scale functional genomic screening, and real-time visualization of genomic elements within living cells (Shalem et al., 2014; Chen et al., 2013). Beyond the laboratory, CRISPR-Cas9 has shown therapeutic promise, including studies in mouse models demonstrating correction of disease-causing mutations, and the technology has also entered broader ethical and scientific discussion around its potential application to human embryos(Yin et al, 2014. Emerging clinical applications under investigation include gene therapy, treatment of infectious diseases such as human immunodeficiency virus (HIV), and the engineering of a patient's own cells for use in cancer immunotherapy (Atkins et al., 2021; Stadtmauer et al., 2020). CRISPR/Cas9 is a gene-editing technology which involves two essential components: a guide RNA to match a

desired target gene, and Cas9 (CRISPR-associated protein 9) - an endonuclease which causes a double-stranded DNA break, allowing modifications to the genome.

Off-Target effects

CRISPR/Cas9 is not without shortcomings, and one of the most significant is the off-target effects. Understanding why it occurs is central to reducing how often it happens. The tolerance built into complementary base pairing between sgRNA and its target DNA is itself the root cause: because a perfect match is not strictly required, an sgRNA can end up binding a site that merely resembles its intended target closely, rather than matching it exactly. Given how many similar sequences exist across a large, complex genome, this creates real opportunities for editing to occur at the wrong location. Off-target activity has been directly measured using experimental tests capable of detecting off-target cleavage, confirming that it occurs across a range of experimental contexts, rather than being an isolated or one-off error.

Mismatches between sgRNA and target DNA generally fall into two categories. In the first, the sgRNA and the off-target site are the same length, differing only at a handful of individual positions. In the second, a difference in length between the two forces a small loop, or bulge, into the paired structure. Neither type is uniformly tolerated: off-target activity tends to decrease as the number of mismatches increases, and Cas9 shows less tolerance for mismatches positioned close to the PAM than for those further away (Li et al., 2019). Left unaddressed, either type of off-target binding can introduce experimental error or produce misleading results, limiting the technology's reliability.

In response, two broad strategies have been developed. One is biochemical: engineering higher-fidelity Cas9 variants that are inherently more selective, alongside more careful guide RNA design intended to reduce off-target activity at the source (Kleinstiver et al., 2016). The other is computational: using algorithms and bioinformatics tools to predict likely off-target sites before editing takes place, so that a target sequence can be chosen with a lower expected risk of unintended effects.

A summary of studies applying traditional machine learning methods for off-target prediction in CRISPR-Cas9

Study	Year	ML model used	Encoding	Data	Link for data/software	Prediction metric/results
Listgarten et al.	2018	Elevation (Boosted regression trees, L1 regression)	One-hot	GUIDE-Seq, Doench V2 ,CRISPOR	http://research.microsoft.com/en-us/projects/crispr	AUROC: 0.98
Lazzarotto et al.	2020	CHANGE-seq (GTB)	One-hot + sequence features	High-throughput sequencing data generated (CHANGE-seq),GUIDE-Seq,CIRCLE-seq	https://github.com/tsailabSJ/changeseq	AUROC: 0.995 AUPRC:0.881
Abadi et al.	2017	CRISTA	GC content + sgRNA secondary structure features	GUIDE-Seq, HTGTS, BLESS	http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1005807	Spearman: 0.81 AUROC: 0.96 AUPRC: 0.96 R2 = 0.8.
Peng et al.	2018	Ensemble SVM	One-hot + GC content +position-specific features	CRISPOR	https://github.com/penn-hui/OfftargetPredict	AUROC: 0.99 AUPRC: 0.45
Dhanjal et al.	2020	. L1 Logistic regression, L2 Logistic regression, RF, xgboost	One-hot + GC content +position-specific features	CIRCLE-seq, CRISPcut	http://web.iitd.ac.in/crispcut/off-targets	Accuracy: 91.49% AUROC: 0.97
Zhang et al.	2019	Ensemble scoring framework of scoring features using AdaBoost	One-hot + GC count + position-specific features	GUIDE-Seq, CRISPOR	https://github.com/Alexzsx/CRISPR	AUROC: 0.938 AUPRC: 0.299

A summary of studies applying deep learning models for off-target prediction in CRISPR-Cas9

Study	Year	ML model used	Encoding	Data	Link for data/software	Prediction metric/ results
Zhang et al.	2020	DL-CRISPR	One-hot	A series of in vitro and cell-based assays	https://github.com/yuuuuzhang/DL-CRISPR_of	Accuracy:98.57%

				were collected to generate data using a new data augmentation method	ftarget_prediction	Sensitivity:95.13 %
Liu et al.	2020	CnnCrispr	Embedding	DeepCRISPR	https://github.com/LQYoLH/CnnCrispr	AUROC: 0.957 AUPRC: 0.429
Lin et al.	2020	CRISPR-Net(LRCN)	one-hot	GUIDE-seq, Doench V2, CRISPOR, CIRCLE-seq, SITE-Seq	https://codeocean.com/capsule/9553651/tree/v1	AUROC: 0.995 AUPRC: 0.317
Zhang et al.	2022	CRISPR-IP(CNN, BLSTM)	one-hot	CIRCLE-Seq, SITE-Seq	https://github.com/BioinfoVirgo/CRISPR-IP	AUROC: 0.982 AUPRC: 0.751 Accuracy: 0.990
Fu et al.	2022	MOFF (2 CNN regression models)	One-hot	GUIDE-seq data],CHANGE-seq, TTISS	https://github.com/MDhewei/MOFF	Spearman: 0.5
Sun et al.	2024	CRISPR-M	one -hot	CIRCLE, GUIDE_I, SITE, PKD, GUIDE_II, GUIDE_III, HEK293T, K562	https://github.com/lyotvincent/CRISPR-M	AUROC : 0.9683 AUPRC : 0.51
Niu et al.	2021	R-CRISPR	One-hot	GUIDE-seq, Doench V2, CRISPOR , CIRCLE-seq, SITE-Seq	https://codeocean.com/capsule/9553651/tree/v1	AUROC: 0.991 AUPRC: 0.319
Zhang et al.	2024	CRISPR-SGRU	one-hot	CHANGE-seq, I1, K562, HEK293t, BE3, II5, I2, and II6	https://github.com/BrokenStringx/Crispr-SGRU	Recall: 0.693 Precision: 0.607 F1 score: 0.641 MCC: 0.644

Methodology

Table 1 and table 2 show that a wide range of ML and DL tools exist for CRISPR/Cas9 off-target prediction, spanning over a decade of methodological development. To provide an overview of how these tools have improved over the years, this paper deep-dives into 5 chosen models that represent a chronological and architectural arc: CRISTA (2017), Elevation (2018), CRISPR-Net (2020), CRISPR-IP (2022), and Crispr-SGRU (2024).

These five were selected for three reasons:

1. They cover a wide seven-year window of intense development and competition, allowing a direct before-and-after comparison.

2. The DL models (CRISPR-Net, CRISPR-IP, Crispr-SGRU) show a clear architectural progression: CRISPR-Net starts with a baseline CNN-plus-recurrent-network approach, CRISPR-IP adds an attention mechanism on top of that same backbone, and Crispr-SGRU extends the recurrent component (via stacking) while directly addressing class imbalance and interpretability.
3. These 5 models are widely used in many papers as benchmarks, meaning the field uses them commonly as references.

To effectively provide a comparison, we focus on six criteria that visibly show the differences in past and current practices:

1. core algorithm/architecture
2. sequence encoding scheme
3. indel handling
4. class imbalance handling
5. interpretability
6. reported performance on standard metrics — Precision, Recall, F1-score, Matthews Correlation Coefficient (MCC), Area Under the ROC Curve (AUROC) and Area Under the Precision-Recall Curve (PRAUC).

As off-target datasets are usually imbalanced, AUROC can be misleadingly high due to its focus on the large majority class (non-off-target). PRAUC, however, is determined through precision and recall for the rare positive class. Therefore, it can be seen as more reliable, alongside MCC that accounts for all four confusion-matrix outcomes. Hence, we take PRAUC and MCC as the more relevant criteria.

Machine Learning models

CRISTA (Abadi et al., 2017)

CRISTA(CRISPR Target Assessment), introduced by Abadi, Yan, Amar, and Mayrose in *PLOS Computational Biology* in 2017, sees off-target prediction as a regression problem, rather than a classification problem. It uses a Random Forest Regressor - that consists of many decision trees - to compute a value (the cleavage propensity score) on a continuous scale instead of making a binary decision.

The model's input representation starts off with a pairwise sequence alignment between the sgRNA and the candidate genomic DNA site. This is computed with a scoring scheme : +1 for a match, 0 for a mismatch, and -1.25 for a gap. However, researchers found that the alignment score only accounts for 34% of the variation in cleavage activity, meaning that other features are necessary to provide a more accurate prediction. The final model with these added features achieved substantially higher accuracy. The researchers' report identified spatial

structure/rigidity and PAM-region features as having the largest impact on predictive performance. Overall, these features-alignment score, positional mismatch information, structural/rigidity descriptors, and PAM-region features form the input to the Random Forest model. Handling of indels is already built into the model as a side effect of sequence alignment, unlike current DL models where indels are a problem that need to be addressed separately. Class imbalance was addressed by undersampling the majority class and oversampling the minority class. On reported performance, CRISTA achieves a Spearman correlation of 0.81, an AUROC of 0.96, an AUPRC (PRAUC) of 0.96, and an R^2 of 0.8 .

Elevation (Listgarten et al., 2018)

Elevation, introduced by Listgarten et al. in *Nature Biomedical Engineering* in 2018, is structured as a two-layer stacked regression model. The first layer (Elevation-score) uses gradient-boosted regression trees - each decision tree assigns a score and the next tree improves its output by correcting the mistakes the tree before it made, and so on - to predict the off-target activity of guide-target pairs that have a single mismatch. The second layer, called the "combiner", uses L1-regularized linear regression (non-essential features' weights are set to zero to prevent overcomplication) to combine multiple single-mismatch predictions into one score for pairs that have several mismatches. A separate model, Elevation-aggregate, also built using gradient-boosted regression trees, then takes all of a guide RNA's individual site scores and rolls them up into one overall guide-level score. This addresses a real practical need: a researcher choosing a guide RNA usually doesn't want hundreds of individual site scores, they want one number to rank guide RNAs by. This same aggregation idea shows up again later in the deep learning era, through CRISPR-Net's own companion aggregation tool.

The model's input features are all hand-engineered: the position of a mismatch, the one-hot-encoded identity of the mismatched nucleotides, and whether the mismatch was a transition or a transversion (a chemically similar substitution vs. a chemically different one). The researchers chose not to model indels as they believed that indels did not affect off-target activity as much as mismatches do. To prevent class imbalance, they used subsampling so that the number of active samples equaled the number of inactive samples.On interpretability, Elevation provides Gini-importance feature rankings that assign a score depending on the relevance of a particular feature in reducing prediction error. However, the researchers state that these rankings are not completely reliable as features affect each other and their value cannot be ascertained independently.

On reported performance, Elevation is evaluated using weighted Spearman correlation rather than AUROC, AUPRC, F1, or MCC - consistent with its framing as a ranking/regression tool rather than a classifier, similar to CRISTA. It's reported to consistently outperform competing methods (CFD, Hsu-Zhang, CCTop) on individual guide-target scoring, and achieved up to an eightfold improvement in Spearman correlation over the best competing method on the guide-level aggregation task .

ML era - Summary

While CRISTA and Elevation both differ in encoding and output type, they share the defining characteristic of the ML era: they rely on hand-engineered features that a human designed and added.The quality of these features depends on the humans's understanding of what influences Cas9 binding and cleavage. This allows the models to have interpretability, but also means that they have limits based on the completeness of these features. This was an obstacle that learned representations in DL models were designed to address.

Deep Learning models

CRISPR-Net (Lin et al., 2020)

CRISPR-Net, published by Lin, Zhang, Zhang, Chen, and Wong in *Advanced Science* in 2020, is a Long-term Recurrent Convolutional Network (LRCN) - a CNN followed by a recurrent layer - represents the first model in this comparison that moves from hand-engineered features to raw, learned sequence encoding.

Its encoding represents each sgRNA–DNA sequence pair as a fixed-length seven-channel binary matrix. A five-bit channel classifies the identity of the nucleotide: A, T, G, C or an indel. It uses a logical OR operation to record the nucleotide identity of both the sgRNA and DNA so that a conclusion (match or mismatch) can be reached. Two additional bits are used to identify the direction and type of any mismatch or indel. The overall matrix is of size (sequence length) x 7, which allows the model to represent both indels and mismatches - this was the first model to incorporate using learned, rather than alignment-based, encoding to handle indels. This encoded matrix is passed through an Inception-style convolutional layer (many parallel filter sizes scanning for sequence patterns at different scales) to catch sequence patterns at several different scales at once, rather than only one fixed pattern size. This output is concatenated to the original encoded input to allow the model to preserve accurate positional information. This combined representation is then passed through a bidirectional Long Short-Term Memory (LSTM) layer to add information on positions further apart, followed by two neural networks that combine all inputs so far. Finally, a single sigmoid output neuron compresses all the information to a score between 0 and 1, that represents the probability that a certain site is off-target. To represent many individual site scores as one overall guide score, a companion tool known as CRISPR-Net Aggregate was developed.

Like CRISTA and Elevation, class imbalance is a real concern for CRISPR-Net's training data, but unlike both classical ML models, no dedicated resampling or reweighting mechanism has been used. On interpretability, CRISPR-Net's learned matrix representation does not offer the built-in inspectability that CRISTA and Elevation's hand-engineered features provide - understanding why the model made a specific prediction would require additional external tooling, a trade-off that comes with moving away from engineered features.

On reported performance, CRISPR-Net trained on three combined datasets outperformed six existing off-target prediction models, reaching the highest AUROC of 0.995 and the highest AUPRC of 0.317 .

CRISPR-IP (Zhang & Jiang, 2022)

CRISPR-IP, published by Zhang and Jiang in *Computational and Structural Biotechnology Journal* in 2022, keeps the same CNN-plus-BiLSTM backbone as CRISPR-Net, along with a compressed version of its encoding scheme, but introduces one new addition: attention. This allows the model to learn which sequence positions should be weighted most heavily, and which can be seen as less important. For example, a mismatch at a sequence position near the PAM would be given more weight if it carried heavier biological consequences.This would take place without a human having to encode "PAM region = important", as attention allows the model to assign these weights based on the training data. Essentially, the BiLSTM's output is weighted by this mechanism, in contrast to the output being treated uniformly as in CRISPR-Net.

Like CRISPR-Net, class imbalance is not addressed as a dedicated design feature. Interpretability is only partially improved relative to CRISPR-Net: the attention weights offer some insight into which positions the model considered important for a given prediction, though this is a weaker and less direct form of interpretability than the engineered feature importances available in CRISTA and Elevation.

The paper provides a direct comparison between CRISPR-Net and CRISPR-IP, using the same datasets, evaluation metrics and training. This allows for a clear representation of the effectiveness of this added attention mechanism. On the CIRCLE-Seq dataset, CRISPR-IP improved over CRISPR-Net by 0.1% in Accuracy, 19.4% in PR-AUC, 3.7% in ROC-AUC, 3.4% in F1-score, 18.5% in Precision, and 2.5% in Recall. On the SITE-Seq dataset, the improvements over CRISPR-Net were 1% in Accuracy, 12.6% in PR-AUC, 2.9% in ROC-AUC, 14.2% in F1-score, 19.4% in Precision, and 2.6% in Recall. We see that attention has greatly improved the model's PRAUC and precision, indicative of the value of "attention".

Crispr-SGRU (Zhang et al., 2024)

Crispr-SGRU, published in the *International Journal of Molecular Sciences* in 2024, continues the same CNN-plus-recurrent-network lineage but modifies it in two significant ways. Firstly, it replaces the BiLSTM layer used in CRISPR-Net and CRISPR-IP with three stacked bidirectional GRU (BiGRU) layers with 30, 20, and 10 units respectively. This allows the model to progressively analyse patterns and representations at a deeper level. Secondly, the model is the first to directly address the issue of class imbalance. It achieves this through a dice loss function; it allows the model to focus on the rare positive class by computing the overlap between the predicted positives and true positives. This function was adapted from medical image segmentation.

Crispr-SGRU consists of two models:

1) The main model: this reuses CRISPR-Net's seven-channel encoding.
2) The student model: unlike the main model, this uses a 24x16 encoding - one channel per base-pair combination - which is trained to mimic the main model's predictions. Deep SHAP, which assigns importance scores to the model's inputs, is used to identify the significance of the contribution of a particular base pair in determining the model's prediction. The purpose of the student model is to address an issue that the other deep learning models were unable to: interpretability.

Crispr-SGRU reports mean Recall of 0.693, Precision of 0.607, F1-score of 0.641, and MCC of 0.644. On the most severely imbalanced datasets (II5, imbalance ratio 1,773.6; II6, imbalance ratio 6,846.6), it reports PRAUC of 0.521 and 0.442 respectively - improvements of 11% and 4.9% over the second-best compared method on each. On dataset I1 (mismatches and indels combined), it reports AUROC of 0.986 and PRAUC of 0.757. An ablation study found that removing the stacked BiGRU component reduced performance by an average of 19.7% on the K562 dataset and 4.8% on HEK293t, underscoring the specific contribution of the stacking design .

DL ERA - Summary

CRISPR-Net, CRISPR-IP, and Crispr-SGRU share a common foundation - a CNN-plus-recurrent-network backbone operating on a learned, rather than hand-engineered, sequence representation - but each addresses a different limitation left open by the one before it. This foundation is introduced by CRISPR-Net, the first in this comparison to handle indels by having a dedicated "gap" slot in its encoding rather than having an extra sequence alignment step (as in CRISTA). However, it does leave class imbalance and interpretability unaddressed.

CRISPR-IP builds directly on this backbone by adding an attention mechanism, and its same-dataset comparison against CRISPR-Net offers concrete evidence of the value this adds: the largest gains appear in PR-AUC and Precision rather than raw Accuracy, suggesting attention specifically improves the model's ability to distinguish rare true positives, rather than only marginally improving overall correctness. Both models so far have yet failed to address class imbalance - a problem that CRISTA and Elevation in the ML era had already tackled through different sampling techniques- as well as interpretability.

Crispr-SGRU, however, closed these gaps by using a dice loss function, which is considerably more sophisticated than subsampling, to address class imbalance and brought back interpretability through a combination of a second "student"model and DeepSHAP. These three models show a clear architectural progression - representational flexibility in the first, a weighting system in the second, and sophisticated handling of class imbalance and interpretability in the third- showing that the evolution of deep learning is not only about increasing accuracy, but regaining capabilities that hand-engineered features in the ML models had by default.

Discussion

The clearest gain of the DL era was representational flexibility: learning directly from a low-level sequence encoding allowed CRISPR-Net, CRISPR-IP, and Crispr-SGRU to incorporate indel handling without requiring a researcher to have anticipated and engineered an indel-specific feature in advance, as CRISTA's alignment-based approach required. However, DL models initially lost the interpretability that Random Forest and gradient-boosted tree models provided through feature importance. CRISPR-IP's attention mechanism offered a partial solution to this, but only through Crispr-SGRU's "student model" and Deep SHAP approach was interpretability fully regained. This shows that recovering this capability required deliberate additional engineering, rather than something that came automatically with better architecture.

One of the most striking findings of this comparison is that class imbalance was not a problem the deep learning era discovered - it was a problem the ML era had already solved twice, independently, before DL models arrived at all. Both CRISTA (via combined oversampling and undersampling) and Elevation (via simple undersampling) built dedicated mechanisms to address the fact that true off-target sites are vastly outnumbered by non-off-target candidates. Yet neither CRISPR-Net nor CRISPR-IP addressed this, indicating that the advancement in architectural sophistication did not translate into a similar advance on a problem the field had already recognized as important. Only with the introduction of Crispr-SGRU's dice loss function did a solution more sophisticated than either ML model emerge. This suggests that architectural innovation (CNNs, attention, stacked recurrence) and problem-specific engineering (imbalance handling, interpretability) progressed somewhat independently of one another.

One comparison in this set avoids this ambiguity entirely. Because CRISPR-IP's authors evaluated their model against CRISPR-Net on identical datasets and metrics, this pairing isolates the effect of a single architectural change (attention) far more cleanly than any other comparison available here. The finding that gains concentrated in PR-AUC and Precision rather than Accuracy is a useful illustration of the broader point above: overall accuracy can mask what is actually improving, or not, in an imbalanced-data setting, reinforcing why this paper treats PRAUC and MCC as more diagnostic than AUROC or Accuracy alone.

A related but distinct divide concerns how these models are evaluated in the first place. In the ML era, off-target prediction was framed as a ranking or regression problem and evaluated primarily with Pearson r^2 and Spearman correlation; CRISPR-Net, CRISPR-IP, and Crispr-SGRU, framing it instead as binary classification, are evaluated with AUROC, AUPRC, F1 and MCC. This reflects an unresolved disagreement across the field: should off-target prediction fundamentally be treated as a ranking task (how do these candidate sites compare to one another in risk) or a classification task (is this specific site an off-target, yes or no)? The distinction carries practical consequences for how a given tool would actually be used in guide RNA design.

Despite these evaluation difficulties, a striking point of agreement emerges across models built on entirely different methodological philosophies. CRISTA's feature-importance analysis, using human-designed features and a forward-selection procedure, found that PAM-region characteristics were among the strongest predictors of cleavage activity. CRISPR-IP's attention

mechanism, by contrast, learns entirely from raw sequence data with no hand-coded biological knowledge at all, yet tends to assign greater weight to positions near the PAM region when making predictions. That two methodologically unrelated approaches - one explicitly told where to look, the other discovering it independently through training - converge on the same region of biological importance is a meaningful form of validation. It suggests this signal reflects underlying biology rather than a side effect of either model's particular design choices, offering a rare point of agreement in a field where direct comparison between models is otherwise difficult.

One task, however, remains comparatively understudied despite its practical importance: aggregation. Both Elevation and CRISPR-Net include a dedicated companion tool - Elevation-aggregate and CRISPR-Net-Aggregate, respectively - built specifically to combine many individual site-level scores into a single guide-level summary score. This is not a minor implementation detail: for a researcher designing a guide RNA, the aggregate score is the more practically relevant output, since it is what determines whether a given guide RNA is judged safe to use overall. Yet this comparison, like the wider literature it draws on, has treated aggregation as a secondary feature of each model rather than a task warranting its own rigorous, cross-model benchmarking. Neither CRISPR-IP nor Crispr-SGRU's papers, as reviewed here, report an aggregation-level evaluation comparable to Elevation's or CRISPR-Net's. This represents a gap distinct from the dataset standardization problem discussed in the Limitations section: even with better datasets, the field would still lack a shared standard for evaluating how well models perform at this second, arguably more decision-relevant, task.

Limitations in the CRISPR research domain

The central limitation of this review - as well as the field it surveys - is the lack of a standardized, high-quality benchmark dataset for CRISPR-Cas9 off-target prediction. Each of the five models compared here was trained and evaluated on data its own authors assembled, using their own criteria for confirming a true off-target site, their own rules for selecting sites that were not off-targets, and used whatever scale of data was practically achievable given the cost of the wet-lab techniques used to detect real off-target activity in the first place, such as GUIDE-seq and CIRCLE-seq, each of which must be run separately for every guide RNA and cell type tested. With no shared, community-wide dataset, every reported performance figure in this paper is a measurement on that model's own data, not a directly comparable score on common ground - which is why the quantitative comparisons here cannot be read as a strict ranking.
A recent study introducing a further off-target tool, CCLMoff (Du et al., 2025), confirms how severe this is in practice: CRISPR-Net's own training dataset contained only 145 sgRNAs, and different groups apply inconsistent criteria when constructing negative samples - CRISPR-Net and CrisprDNT, for instance, both used Cas-OFFinder to generate negative sites, but with different mismatch and bulge parameters, so their "negative" classes are not defined identically.

This problem affects performance metrics across the five models examined here in several ways. First, dataset size varies enormously - from CRISPR-Net's 145 sgRNAs to Crispr-SGRU's evaluation across eight separate datasets of varying scale. Second, the balance

between positive and negative examples varies dramatically across datasets: some are close to evenly split between confirmed off-target sites and non-off-target sites (as in some of CRISTA's evaluation data), while others are extremely skewed, with true off-targets making up a small fraction of the total (Crispr-SGRU's II6 dataset, for example, has roughly 6,847 non-off-target sites for every 1 true off-target). This matters because metrics like PRAUC and AUROC are directly affected by how skewed a dataset is - a model tested on a more evenly balanced dataset will tend to post a higher PRAUC than the same model tested on a heavily skewed one, simply because the task is inherently easier when positive examples aren't so rare. CRISTA's high AUPRC of 0.96, for instance, is likely explained in large part by its own dataset being more balanced, rather than indicating it completely outperforms Crispr-SGRU, which was specifically tested on much harder, more skewed data. Third, there is no agreed-upon standard for how researchers decide which genomic sites to include as "negative" examples - meaning sites they treat as confirmed non-off-targets. Every researcher generates their own list of candidate sites to test (using tools that scan the genome for sequences similar enough to the guide RNA to be worth checking), and different research teams set different rules for how similar a site has to be before it even gets included in that list. As a result, one paper's negative examples might all be close, tricky look-alikes to the guide RNA, while another paper's negative examples might include many sites that are more obviously different - meaning two models' "negative" categories may not actually represent the same kind of genomic sites at all, even though both are labeled the same way. Fourth, some models (CRISTA, Elevation) evaluate using ranking-based metrics on continuous scores, while others (CRISPR-Net, CRISPR-IP, Crispr-SGRU) evaluate using classification metrics on binary labels, a difference in evaluation philosophy layered on top of the dataset differences already described.

Beyond the dataset problem, the field faces several other open challenges. One is generalization to conditions not well represented in existing training data. Crispr-SGRU, for instance, restricts its evaluation to sites with an NGG PAM sequence, explicitly leaving alternative PAM sequences (such as NAG or NGA) as future work. This means the model's reported performance is only established for NGG sites - how it would perform on other PAM variants remains untested and unknown.

A related concern is that standard cross-validation may overstate real-world reliability. In typical cross-validation, held-out test data often still contains guide RNAs that resemble ones the model encountered during training, giving it an easier task than it would face with a truly novel guide RNA in practice. A stricter alternative, leave-one-sgRNA-out testing, withholds every site associated with a given guide RNA from training entirely, so that guide RNA is completely new at test time. When Crispr-SGRU's own authors applied this stricter test, performance dropped measurably compared to standard cross-validation on the same data. This indicates a real gap between a model's reported benchmark performance and how reliably it would perform when applied to entirely new guide RNAs in practice.

Another challenge is the reliability of interpretability methods themselves. Even the models in this comparison that make an explicit effort toward interpretability are not without

flaws: Elevation's Gini-importance rankings cannot be cleanly separated when features interact, and Crispr-SGRU's SHAP-based attributions are computed on a smaller "student" model trained to approximate the main model, rather than on the main model directly. This suggests interpretability remains an unresolved problem in this field, despite the existence of tools designed to address it.

Conclusion

Across the seven years spanned by CRISTA, Elevation, CRISPR-Net, CRISPR-IP, and Crispr-SGRU, CRISPR-Cas9 off-target prediction has moved from models built on hand-engineered sequence and structural features toward models that learn their own representations directly from raw sequence data. This shift brought genuine gains but it was not a clean, linear march toward a strictly better tool at every step. Interpretability, present by default in CRISTA and Elevation's engineered-feature approach, had to be deliberately rebuilt by Crispr-SGRU through knowledge distillation and SHAP. Class imbalance handling, already solved independently by both classical ML models through different resampling strategies, was set aside entirely by the first two deep learning models before Crispr-SGRU reintroduced it in a more sophisticated form. This pattern points to something worth taking seriously going forward: a newer or more architecturally advanced model should not be assumed to have inherited every strength of the models before it. Verifying this directly - checking whether a new architecture actually preserves imbalance handling and interpretability, rather than assuming it does because it improves accuracy - would be a small but meaningful addition to how new models are evaluated and reported.

The most immediate priority this comparison points to is the adoption of a shared, standardized benchmark dataset, with transparent and consistent rules for what counts as a confirmed off-target and what counts as a valid negative comparison site. As long as every research group builds its own dataset using its own criteria, a reported improvement in accuracy or PRAUC could rather reflect a difference in how the data was assembled than a genuine improvement in the model itself. Establishing a common benchmark would let future comparisons attribute performance differences to the model itself.

Closely related is the question of how models are tested for real-world reliability. Much of the reported performance comes from standard cross-validation, which tends to understate how a model will perform on conditions genuinely absent from its training data - a new guide RNA it has never encountered in any form, or a PAM sequence outside the range it was trained on. Future work in this space would benefit from treating stricter tests, such as evaluating on entirely withheld guide RNAs or on underrepresented PAM variants, as a standard part of reporting a new model's performance, rather than as an optional, secondary check. This distinction matters practically: a tool intended to guide real experimental decisions is only as useful as its performance on cases it has not already seen.

There are many directions of future research that could meaningfully advance the field. One is architectural: transformer-based models, which extend the attention mechanism

CRISPR-IP first introduced into this comparison, represent a natural next step, and would be most valuable if evaluated not only on raw accuracy but on whether they retain the imbalance-handling and interpretability advances later models in this comparison had to work to establish. A direction not yet explored by any of the models compared here is joint prediction of on-target efficiency and off-target risk within a single model. Guide RNA design in practice requires balancing both concerns at once - a guide that minimizes off-target activity is not useful if it also cuts poorly at its intended site - yet every model in this comparison addresses off-target prediction in isolation. A model trained to predict both properties simultaneously, rather than relying on two independently developed tools attached together, could better reflect how these predictions are actually used, and might also reveal whether the sequence features that drive on-target efficiency and off-target risk interact in ways that neither prediction task captures on its own.

Acknowledgments

I would like to thank Dr. Eric Sakk for his mentorship and guidance throughout the development of this paper, and for his valuable feedback on the research and analysis presented here.

Works Cited

Abadi, Shiran, et al. “A Machine Learning Approach for Predicting CRISPR-Cas9 Cleavage Efficiencies and Patterns Underlying Its Mechanism of Action.” PLoS Computational Biology, vol. 13, no. 10, Oct. 2017, p. e1005807

Anders, C., et al. "Structural Basis of PAM-Dependent Target DNA Recognition by the Cas9 Endonuclease." Nature, vol. 513, no. 7519, 2014, pp. 569-73.

Atkins, A., et al. "Off-Target Analysis in Gene Editing and Applications for Clinical Translation of CRISPR/Cas9 in HIV-1 Therapy." Frontiers in Genome Editing, vol. 3, 2021, article 673022.

Barrangou, R., et al. "CRISPR Provides Acquired Resistance against Viruses in Prokaryotes." Science, vol. 315, no. 5819, 2007, pp. 1709-12.

Chen, B., et al. "Dynamic Imaging of Genomic Loci in Living Human Cells by an Optimized CRISPR/Cas System." Cell, vol. 155, no. 7, 2013, pp. 1479-91.

Dhanjal, J.K., et al. "Evaluation of Off-Targets Predicted by sgRNA Design Tools." Genomics, vol. 112, no. 5, 2020, pp. 3609-18.

Du, W., et al. "A Versatile CRISPR/Cas9 System Off-Target Prediction Tool Using Language Model." Communications Biology, vol. 8, 2025, article 882.

Fu, R., et al. "Systematic Decomposition of Sequence Determinants Governing CRISPR/Cas9 Specificity." Nature Communications, vol. 13, 2022, article 474.

Huang, W.Z. "Computational Prediction of Off-Target Effects in CRISPR Systems." Computational Molecular Biology, vol. 15, no. 2, 2025, pp. 102-11.

Jinek, M., et al. "A Programmable Dual-RNA-Guided DNA Endonuclease in Adaptive Bacterial Immunity." Science, vol. 337, no. 6096, 2012, pp. 816-21.

Kleinstiver, B.P., et al. "High-Fidelity CRISPR-Cas9 Nucleases with No Detectable Genome-Wide Off-Target Effects." Nature, vol. 529, no. 7587, 2016, pp. 490-95.

Lazzarotto, C.R., et al. "CHANGE-Seq Reveals Genetic and Epigenetic Effects on CRISPR–Cas9 Genome-Wide Activity." Nature Biotechnology, vol. 38, no. 11, 2020, pp. 1317-27.

Li, D., et al. "Battling CRISPR-Cas9 Off-Target Genome Editing." Cell Biology and Toxicology, vol. 35, no. 5, 2019, pp. 403-06.

Lin, J., et al. "CRISPR-Net: A Recurrent Convolutional Network Quantifies CRISPR Off-Target Activities with Mismatches and Indels." Advanced Science, vol. 7, no. 13, 2020, article 1903562.

Listgarten, J., et al. "Prediction of Off-Target Activities for the End-to-End Design of CRISPR Guide RNAs." Nature Biomedical Engineering, vol. 2, 2018, pp. 38-47.

Liu, Q., et al. "Deep Learning Improves the Ability of sgRNA Off-Target Propensity Prediction." BMC Bioinformatics, vol. 21, no. 1, 2020, article 51.

Niu, R., et al. "R-CRISPR: A Deep Learning Network to Predict Off-Target Activities with Mismatch, Insertion and Deletion in CRISPR-Cas9 System." Genes, vol. 12, no. 12, 2021, article 1878.

Peng, H., et al. "Recognition of CRISPR/Cas9 Off-Target Sites through Ensemble Learning of Uneven Mismatch Distributions." Bioinformatics, vol. 34, no. 17, 2018, pp. i757-i765.
Shalem, O., et al. "Genome-Scale CRISPR-Cas9 Knockout Screening in Human Cells." Science, vol. 343, no. 6166, 2014, pp. 84-87.
Stadtmauer, E.A., et al. "CRISPR-Engineered T Cells in Patients with Refractory Cancer." Science, vol. 367, no. 6481, 2020, eaba7365.
Sun, J., et al. "CRISPR-M: Predicting sgRNA Off-Target Effect Using a Multi-View Deep Learning Network." PLOS Computational Biology, vol. 20, no. 3, 2024, e1011972.
Yin, H., et al. "Genome Editing with Cas9 in Adult Mice Corrects a Disease Mutation and Phenotype." Nature Biotechnology, vol. 32, no. 6, 2014, pp. 551-53.
Zhang, G., et al. "Crispr-SGRU: Prediction of CRISPR/Cas9 Off-Target Activities with Mismatches and Indels Using Stacked BiGRU." International Journal of Molecular Sciences, vol. 25, no. 20, 2024, article 10945.
Zhang, S., et al. "Synergizing CRISPR/Cas9 Off-Target Predictions for Ensemble Insights and Practical Applications." Bioinformatics, vol. 35, no. 7, 2019, pp. 1108-15.
Zhang, Y., et al. "DL-CRISPR: A Deep Learning Method for Off-Target Activity Prediction in CRISPR/Cas9 with Data Augmentation." IEEE Access, vol. 8, 2020, pp. 76610-17.
Zhang, Z.-R., and Z.-R. Jiang. "Effective Use of Sequence Information to Predict CRISPR-Cas9 Off-Target." Computational and Structural Biotechnology Journal, vol. 20, 2022, pp. 650-61.

SiO2 Cloud-phase Ambiguity Through Nucleation and Annealing Kinetics on HD 189733b By Avni Jain

Abstract

The dayside atmosphere of exoplanet HD 189733b hosts high-altitude silicate clouds. Given HD 189733b's atmospheric temperature of ~1200 K, nucleation and annealing kinetics may give insight into the phase of these amorphous or crystalline SiO_2 clouds. This review evaluates whether existing nucleation and annealing constraints support crystalline SiO_2 clouds on HD 189733b. JWST/MIRI observations revealed a feature interpreted as evidence for small SiO_2 grains. However, two analysis methods within the same study reached different conclusions about their physical phase. A statistical retrieval demonstrated a slight preference for amorphous, while a separate grid model assumed a crystalline structure as an input. This ambiguity requires a separate method to determine the phase of the silicate clouds. Existing circulation modelling analyses suggest that vertical settling timescales exceed 1000 days for small (<0.1 μm) particles in the upper atmosphere. Furthermore, laboratory annealing experiments on amorphous magnesium silicate smoke show that above 1067 K, the stall in spectral evolution is no longer present. Grains formed above ~1100 K reach a fully crystalline state within ~10 minutes. HD 189733b's dayside equilibrium temperature, ~1191 K, exceeds both thresholds. These laboratory rates suggest that annealing may be rapid. Hence, this paper argues that time is unlikely to be the limiting factor, making a purely amorphous result difficult to explain. Resolving this ambiguity is important for accurate spectral analyses of the atmosphere of HD 189733b and other silicate-cloud exoplanets, since the phase of the silicate clouds uniquely determines how they interact with light.

1. Introduction

Hot Jupiters challenge standard models of planet formation. These gas giants orbit closer to their host stars than Mercury orbits the Sun. Some of these even completely orbit their host star in under four days. One example is KOI 1843.03, which has an orbital period of just 4.2 hours (Rappaport et al.). In the protoplanetary disk from which planets accrete, only rocky material can condense close to a star, since it is too hot for volatile ices to survive. Gas giants, on the other hand, are expected to form beyond the frost line (Kennedy & Kenyon), which is the point in the disk that is cold enough for ices like water, methane, and ammonia to freeze out of the gas and stick to dust grains. In other words, ices condense alongside rock, allowing planetesimals to accrete enough matter to pull in the surrounding hydrogen and helium and grow into a full gas giant (Brouwers and Ormel). Under this model, gas giants should not be found close to their host star, since there simply isn't enough solid material in this region to build the massive core a gas giant needs before it can start accreting in gas.

Hot Jupiters were, in fact, the first exoplanets discovered around main-sequence stars, and their existence immediately challenged this textbook, solar-system-derived picture of planet

formation (Dawson & Johnson). More than two decades later, there is still no consensus on their origin. Dawson and Johnson identify three proposed formation channels: in situ formation, disk migration, and high-eccentricity tidal migration. In situ formation suggests that hot Jupiters form close to their host star to begin with, rather than moving there afterwards. In contrast, disk migration posits that the planet forms further away from the host, near or beyond the frost line, and then gradually spirals inward as it interacts gravitationally with the leftover gas in the protoplanetary disk. Slowly, it starts losing orbital energy over time until it settles into a tight orbit. In high-eccentricity tidal migration, a gravitational interaction with another planet or star in the system flings the hot Jupiter onto a highly stretched, eccentric orbit. Tidal forces from the host star then gradually circularise the orbit and pull the planet into the tight orbit we observe today. Dawson and Johnson argue that no single channel can explain everything, but pairing two of the three major channels together may account for hot Jupiters. Their short orbital periods and large radii also make hot Jupiters easy to detect using radial velocity and transit methods, compared to smaller and more distant exoplanets. For this reason, they have played a central role in early exoplanet atmospheric characterisation.

Atmospheric characterisation of hot Jupiters has depended heavily on observations from Hubble and JWST. Hubble observations, largely in the optical and near-infrared, provided the first measurements of an exoplanetary atmosphere (Charbonneau et al.). JWST extends into the mid-infrared, where molecular and mineral absorption features become visible (Inglis et al., 2024). This includes a silicate feature near 8.7 μm that motivates this paper.

Clouds can only form in the presence of an atmosphere. They form when a vapour cools past its condensation point and nucleates into particles (Powell et al.). It is the same basic process that governs cloud formation on Earth. In some cases, instead of water vapour condensing into water droplets, it might be silicate vapour condensing into mineral grains (Powell et al.). On a hot Jupiter, however, this raises two difficult questions: what are these clouds made of, and what physical state are they in? Both are difficult to answer from Earth-based or even space-based observations alone, as it is not possible to test samples from these planets directly. Hence, they motivate this paper. We investigate whether nucleation and annealing kinetics favour amorphous or crystalline silicate grain formation at ~1200 K in the atmosphere of the planet HD 189733b.

2. Observational Evidence for Silicate Clouds

HD 189733b is one of the most heavily studied hot Jupiters (Pont, Sing, et al.). It is bright, close, and easy to observe compared to most exoplanets. JWST's 2024 observations provided strong evidence that the clouds in its atmosphere are dominated by SiO_2. Furthermore, the cloud optical depth has been constrained below 0.05 across the 1-8 μm range (Inglis et al.), which means that the clouds barely affect the planet's radiative balance even though they are detectable spectroscopically.

Even before JWST, indirect evidence already pointed toward sub-micron scattering particles sitting somewhere near HD 189733b's day-night terminator. The planet's transmission

spectrum shows a steep optical scattering slope paired with a weaker-than-expected near-infrared water absorption feature (Etangs et al.; Madhusudhan et al.; Pont, Knutson, et al.; Sing et al.; Sing et al.). Both signs point toward small scattering particles in the atmosphere rather than a clear sky. Wakeford and Sing further proposed that the scattering could come specifically from "small, sub-micron grains of enstatite, or $MgSiO_3$". Separately, Evans et al. used HST/STIS to measure the planet's geometric albedo. They observed that there was a sharp increase in wavelengths for reflected blue light, which they point out as likely being due to the presence of light-scattering sub-micron aerosols. However, this line of evidence now appears less robust than previously thought. Inglis et al. point out that more recent CHEOPS photometry from Krenn et al. measured a smaller geometric albedo than Evans et al. found. This lower albedo is more consistent with a dayside model that has no significant cloud scattering. Therefore, before JWST, empirical cloud models proposed an ambiguous picture, and the cloud composition remained uncertain.

To confirm the composition, a direct spectroscopic detection was required. JWST's MIRI LRS instrument produced the first direct composition-sensitive evidence for silicate clouds on this planet. This was presented in the form of an excess absorption feature at 8.7 μm (Inglis et al.). The study proposed that the absorption feature might be due to small SiO_2 grains of ~0.01 μm sitting in the uppermost layers of the planet's dayside atmosphere. This was the first direct detection of silicate clouds on HD 189733b, with the SiO_2 cloud model being preferred over a clear-atmosphere model at 6.0σ (Inglis et al.). The study also offers an explanation for why the clouds are located so high up in the atmosphere. They tie the formation of these clouds to the hottest regions of the

dayside, close to the substellar point. This detail gives insights into where nucleation is occurring. But the question of the physical phase, crystalline or amorphous, that SiO_2 clouds exist in remains.

To answer this question, Inglis et al. performed a retrieval analysis by running scattering coefficients for multiple candidate cloud species: MnS, ZnS, KCl, Fe, $MgSiO_3$, Mg_2SiO_4, and SiO_2. For the species where data were available, they tested both crystalline and amorphous forms directly against each other as free parameters. The results showed that there was little difference in phase, and amorphous SiO_2 clouds were slightly favoured over crystalline at a level of ≤1σ. This small preference is insufficient to confidently assert the physical state of silicate clouds. Hence, the researchers needed to use another approach, forward-modelling, to determine if the silicate clouds are crystalline or amorphous.

In their forward-modelling analysis, the researchers fit the spectrum for five free parameters: the base pressure level of the cloud deck, the sedimentation efficiency (f_{sed}), the cloud particle density, the mean particle radius, and the width of the particle size distribution. The cloud phase was not included as a free parameter. Instead, Inglis et al. assumed Clouds composed of solid SiO_2 (quartz) on the planet at a 1191 K equilibrium temperature, utilising α-crystal optical constants measured at 928 K by Zeidler et al. The crystalline assumption is

mainly due to this previous work, which only measured temperature-dependent optical constants for corundum, spinel, and alpha-quartz, which is the crystalline form of SiO2 (Zeidler et al.). Therefore, the forward model's crystalline assumption may not reflect the true phase of silicate clouds, since the only high-temperature optical constraints available when the fit was constructed were derived from crystalline SiO_2.

JWST observations have confirmed the presence of silicate clouds on HD 189733b at 6σ. But the apparent disagreement between the retrieval and forward-modelling results cannot, on their own, confirm the phase of the SiO_2. One method resulted in a weak, ≤1σ lean toward amorphous. The other assumed crystalline before fitting, because optical high-temperature constraints have only been reported for crystalline SiO_2. Therefore, the phase cannot be settled by either method alone. Instead, nucleation and annealing kinetics need to be considered as an independent approach (Helling et al.; Lee et al.; Powell et al.) as a free parameter within a forward-modelling framework rather than fixed beforehand.

3. Atmospheric Conditions and Dynamics on HD 189733b

HD 189733b is tidally locked, meaning one side always faces its host star throughout the exoplanet's orbit. Because of this, the planet does not have a consistent mean atmospheric temperature but rather two different mean temperatures for the day side and the night side. Understanding how large this day-night temperature difference actually is can help us determine whether the silicate grains detected earlier have the right conditions and enough time to crystallise.

The dayside and nightside of HD 189733b differ substantially in temperature. Knutson et al., analysing the half-orbit phase curve using data from the Infrared Array Camera (IRAC) on the Spitzer Space Telescope, measured a brightness temperature at 8 μm that reaches about 1212 ± 11 K at its hottest and drops to about 973 ± 33 K at its coolest, which is a gap of roughly 240 K across the planet. That same study also timed the peak in the planet's flux and found that it occurs 2.3 ± 0.8 hours before the centre of the secondary eclipse, or roughly 16 ± 6 degrees before opposition. However, this measurement does not directly locate the hotspot: the region of maximum temperature on the planet. Rather, it corresponds to the time of peak disk-integrated brightness. A more detailed fit of that same phase curve divided the dayside into twelve longitudinal strips and placed the actual hot spot roughly 30 degrees east of the substellar point (Knutson et al.). Furthermore, using a full-orbit, multi-wavelength phase curve, Knutson et al. found hot-spot offsets of 35 ± 4 degrees and 20 ± 6 degrees in two separate wavelength channels. Compared against the 30-degree figure from Knutson et. al., this result suggests a more complicated heat transport mechanism across the exoplanet's atmosphere than simple eastward wind-driven directions.

We must now determine the size of these particles in order to better understand the atmospheric circulation. Lee et al. modelled cloud particles on HD 189733b at pressures less than the 1 bar level and found that most of them fall somewhere between 0.01 and 0.5 μm. This

matches the roughly 0.01 μm grains in the atmosphere of HD 189733b detected by Inglis et al. (2024). Parmentier et al., modelling the exoplanet HD 209458b, show that strong vertical mixing can keep material airborne for particles of up to a few μm across. Smaller particles tend to remain well mixed with the gas, while larger particles settle more efficiently. Lee et al. (2016) further find that small particles can stay aloft for long stretches from atmospheric motion alone, for example, a 1,000-day settling time scale for particles under 0.1 μm. However, we note that this is a generalisation of a nightside-specific result, and should not be taken as a direct dayside measurement. Lee et al. also modelled HD 189733b and found that SiO_2 particles are the dominant grain type at low dayside latitudes and low pressures, and that particle growth is suppressed in the hottest parts of the dayside. That analysis helps explain whether these grains stay concentrated on the dayside of the planet rather than drifting to the cooler nightside during their ≥1000-day lifetime. Mullens and Lewis estimate horizontal winds of roughly 1 to 5 km/s for this class of planet, further demonstrating that grains moving between hemispheres is plausible.

Overall, HD 189733b's dayside has a mean atmospheric temperature of around 1191 K, and grains appear to persist in the upper atmosphere for over 1000 days, according to circulation modelling. This means that residence time on its own is unlikely to hinder the process of crystallisation. The Lee et al. result gives reasonable grounds to treat this long residence time as relevant to the hot dayside specifically. However, this conclusion only applies to the dayside. The average nightside brightness temperature is roughly 240 K lower than the dayside temperature (Knutson et al.), and the fraction of the grain residence lifetime spent on the cooler nightside is still an open question. It is therefore unclear how a higher-resolution circulation model would change where SiO_2 nucleation is predicted to happen.

4. Nucleation and Annealing Kinetics of Silicate Grains

When a vapour condenses into a solid faster than its atoms can settle into an ordered lattice, the result is an amorphous, glass-like structure rather than a crystal. Annealing is the process by which the atoms of an amorphous solid slowly rearrange themselves into a crystalline structure, given that the solid is held at a sufficiently high temperature for a long period of time. But before annealing can act, the grain must first be nucleated. This question is governed by the nucleation theory. This section discusses both the annealing and nucleation processes, first by looking at how fast existing amorphous SiO_2 grains crystallise once they are at HD 189733b's dayside temperature and, second, by understanding how and where these grains form.

Hallenbeck et al. performed laboratory annealing experiments on magnesium silicate smoke, which was condensed from vapour to mimic the grain formation process, and tracked how the infrared spectrum changed with time and temperature. Below approximately 1067 K, this evolution proceeds in two distinct phases: the spectrum changes, stalls partway through crystallisation, and then advances toward a fully crystalline structure. Above this thermal threshold, the intermediate stall is no longer present, allowing the grains to anneal continuously without interruption. At temperatures above 1100 K, the full process completes on timescales of

10 minutes or less. However, below 1000 K, grains rarely anneal enough in a typical astrophysical timeframe to reach even the stalled state. HD 189733b's day side sits at roughly 1191 K, which lies above the 1067 K threshold where the stall is no longer observed. The underlying crystallisation rate as a function of time and temperature is given by Harker & Desch.

$$\mathrm{SEI} \approx 1 + 5.88\times10^{69} \cdot t \cdot \exp(-177{,}430/T) \quad (1)$$

Here, t is time in seconds, T is temperature in Kelvin, and SEI is the spectral evolution index, with full crystallisation corresponding to SEI ≥ 100. At 1191 K, this expression implies that full crystallisation can occur on very short time scales (minutes or less), although the exact value depends on the adopted parameters. This trend aligns with the findings of Hill et al., suggesting that thermal annealing is not the bottleneck for grain crystallinity in hot Jupiter environments.

While annealing describes the growth rate of existing crystalline grains, nucleation theory explains their initial formation and accounts for variations in crystallisation kinetics. The findings of Inglis et al. align with this established framework. Their retrieved cloud parameters indicate that either homogeneous or heterogeneous nucleation occurs in the upper atmosphere, supporting the mechanisms proposed by Lee et al. and Powell et al. Notably, the atmospheric model of HD 189733b developed by Lee et al. directly incorporates this specific framework. In their model, TiO_2 clusters form first through homogeneous nucleation, and these clusters then act as seed particles that other species (including SiO_2) accrete onto through heterogeneous surface reactions. An astrophysically grounded nucleation pathway already exists specifically for this exoplanet, which is consistent with broader principles from terrestrial materials science. Wang et al., building on Uhlmann's application of the classical nucleation theory, calculated that pure quartz glass held at 1573 K would take about 5 x 10^4 seconds, roughly equal to 14 hours, to nucleate on its own. This time period is slow enough that homogeneous nucleation seems unlikely in practice. But actual quartz glass crystallises much faster than that number implies. Observations show that the crystallisation starts at the surface rather than spreading evenly through the volume, meaning real crystallisation is heterogeneous, seeded by existing surfaces and impurities. Jean et al. studied borosilicate glass-ceramic, where cristobalite, one of the crystalline forms of SiO_2, was one of two phases that formed. The researchers determined that a viscous flow process controls the initial incubation delay before crystallisation begins. During this preliminary stage, the nucleation activation energy ranges from 270 to 290 kJ/mol. Following this initial period, the subsequent crystallisation process exhibits growth behaviour that aligns with diffusion control mechanisms. However, the two studies above analyse melt-cooled industrial glasses. Hence, the formation pathway on HD 189733b is fundamentally different because the exoplanet's crystallisation environment consists of vapour-phase condensation within a low-pressure, low-gravity atmosphere dominated by hydrogen. Nevertheless, the general principle that unseeded nucleation is slow while seeded nucleation is

fast is still applicable in this context. The specific gestation period and the activation energy, however, should not be used as quantitative predictions for this exoplanet.

Applying existing laboratory data to the atmosphere of HD 189733b involves several important limitations. A primary challenge is the scarcity of experimental measurements of silicate properties at temperatures representative of exoplanetary atmospheres. Potapov and Bouwman note that studies investigating the optical properties of refractory materials under exoplanet-relevant high-temperature conditions remain limited. This is significant because temperature directly influences a material's optical properties. As temperature increases, structural and chemical changes can shift and broaden spectral features. Consequently, the use of room-temperature data in high-temperature atmospheric models can introduce substantial inaccuracies. The availability of temperature-specific laboratory data is therefore essential for accurate modelling. Current laboratory measurements of crystalline silicates, derived through the high-temperature annealing of amorphous materials, extend only to approximately 1000 K, which is insufficient for characterising the 1191 K dayside temperature of HD 189733b. Potapov and Bouwman further argue that experimental investigations spanning a temperature range of 300–2000 K are necessary. An additional limitation arises from post-experimental sample handling. Exposure of annealed samples to air may induce chemical reactions that were not present under the original high-temperature conditions. This could potentially introduce biases into laboratory measurements relative to the processes occurring in an exoplanet atmosphere.

Applying Hallenbeck's thresholds to HD 189733b's approximately 1191 K dayside temperature suggests crystallisation timescales on the order of minutes, potentially even shorter. This timescale is well within the cloud particle residence time of more than 1000 days established earlier. Potapov and Bouwman provide an important assessment of the uncertainties related to this interpretation. Given the limited availability of high-temperature silicate data, their review provides a useful benchmark for evaluating current constraints and the limits of existing observations and experiments. Other studies have similarly reported amorphous-to-crystalline annealing on relatively short timescales, supporting the general conclusion that crystallisation can occur efficiently under sufficiently hot conditions (Gail and Sedlmayr; Fabian et al.; Mullens and Lewis). However, laboratory measurements of silicate annealing at temperatures above approximately 1000 K remain scarce (Potapov and Bouwman). As a result, even physically motivated extrapolations may fail to capture important chemical and structural processes that occur at higher temperatures. Addressing this limitation will require additional in-situ experimental measurements under conditions similar to those of hot-Jupiter atmospheres.

5. Spectroscopic Signatures of SiO2 Phase on Hot-Jupiters

Although JWST detected SiO_2 in the atmosphere of HD 189733b, distinguishing between amorphous and crystalline grains in transmission spectra remains inherently difficult. This represents a fundamental limitation of the measurement itself rather than a data-quality issue that can be resolved through longer exposure times. At the spectral resolution provided by MIRI, the

differences between the two phases can be subtle and difficult to isolate. Consequently, the retrieval's weak preference for amorphous SiO_2 may reflect this observational limitation rather than a genuine physical signal. Before considering observations of other hot Jupiters, it is important to ask whether the two phases can even reliably be distinguished in principle using currently available instruments.

Amorphous and crystalline silicate grains are spectroscopically distinguishable in the mid-infrared. This distinction arises from differences in their internal atomic structure. In amorphous aerosols, atoms are arranged randomly, causing light absorption to be distributed across broad, smooth spectral features without any fine structure. In contrast, crystalline aerosols possess a regular, repeating lattice structure that produces sharp spectral features with pronounced fine structure and a much stronger temperature dependence (Mullens & Lewis). In principle, spectra with sufficient resolution and signal-to-noise ratio should be capable of distinguishing between the two phases based solely on the shape of their absorption features. Therefore, the phase ambiguity associated with the SiO_2 detection in HD 189733b likely does not reflect a fundamental physical limitation. Rather, it is an observational challenge that could plausibly be resolved with higher-quality data. If amorphous and crystalline silicates produced indistinguishable spectra, comparisons across different planets would provide no meaningful information about which phase is actually present.

Crystalline SiO_2 has already been observationally preferred in another hot Jupiter, and the temperature pattern across planets is consistent with HD 189733b also leaning toward a crystalline interpretation. This comparison is made directly by Inglis et al. Using equilibrium temperatures, HD 189733b is placed at approximately 1191 K, between WASP-17b at approximately 1760 K and WASP-107b at approximately 750 K (Southworth; Mullens & Lewis; Alderson et al.; Dyrek et al.). A slight preference for crystalline SiO_2 over the amorphous form was reported for WASP-17b (Grant et al.). The SiO_2 cloud model showed a 3.5-4.2σ improvement over a cloud-free model and a 2.6σ improvement over a generic aerosol model. These results are substantially less definitive than the 6.0σ SiO_2 detection reported for HD 189733b (Inglis et al.). In contrast, WASP-107b's silicate signature is not dominated by crystalline material. Instead, the spectrum was reproduced using a mixture of amorphous SiO, SiO_2, $MgSiO_3$, and carbon (Dyrek et al.). That ordering alone does not establish a particular phase state. WASP-17b's atmosphere also has a super-solar metallicity of up to 100 times solar, which shifts the SiO_2 condensation curve independently of temperature (Inglis et al.). Even though temperature is not the only variable, the available evidence remains consistent with the central argument of this paper. Data suggest that the cooler planet may contain an amorphous mixture, while the hotter planet might show a preference for crystalline structures. HD 189733b lies between these two cases and is closer to the crystalline end. Amorphous aerosols can also transition to crystalline through high-temperature annealing on timescales of hours (Mullens & Lewis), the same hours-scale estimate discussed earlier in relation to Hallenbeck's shorter crystallisation timescales. This precedent strengthens the case for HD 189733b being crystalline,

while still carrying the limitations previously discussed.

Grant et al. acknowledge a limitation that reinforces a point raised earlier in this paper. Crystalline SiO_2 provided a slightly better match to the 8.6 μm feature than amorphous SiO_2. However, neither material's optical constants have been measured in laboratory conditions representative of WASP-17b's atmosphere, indicating that additional experimental work is needed (Grant et al.). This serves as direct confirmation of the optical-constants gap highlighted by Potapov and Bouwman. Its presence in a second, independently analysed planet demonstrates that the issue is not unique to HD 189733b. It further shows that the shortage of laboratory data is not a consequence of how a single planet is analysed, but rather a genuine gap in the field's experimental record. The same limitation emerges whenever researchers attempt to constrain the silicate phase with high spectroscopic precision.

While mid-infrared spectroscopy can, in principle, distinguish between amorphous and crystalline silicate phases, observations across multiple planets illustrate how this may vary with temperature. WASP-17b, which is hotter than HD 189733b, shows a tentative preference for crystalline silicates, whereas the cooler WASP-107b favours an amorphous structure. This cross-planet temperature trend is consistent with the central thesis of this paper. However, the evidence from WASP-17b is statistically weaker than HD 189733b's own SiO_2 detection and is subject to the same high-temperature optical-constants gap discussed earlier. The case for HD 189733b's phase therefore continues to rest primarily on the kinetics argument derived from laboratory annealing data. Whether higher-resolution or broader-wavelength JWST observations can independently resolve this sub-sigma phase preference remains an open question. Testing this will determine whether observational constraints alone can break these physical degeneracies without relying on cloud kinetics or cross-planet comparisons.

6. Conclusion

In summary, the evidence points more toward HD 189733b's silicate clouds being crystalline than amorphous. Because a definitive measurement is not yet available, this finding cannot rest on a single direct observation. Instead, it comes from several separate lines of indirect evidence: observational, dynamical, kinetic, and spectroscopic. All of these suggest that the nucleation and annealing kinetics of these grains clear both the time and temperature thresholds needed to crystallise. We outline how these findings connect, while also addressing the areas where the evidence remains less definitive.

In section 2, we discussed the retrieval and forward model from Inglis et al. The retrieval tested the phase directly and only found a weak, $<1\sigma$, lean toward amorphous. The forward model assumed crystalline mostly because optical constants at high temperatures were only available for crystalline SiO_2, not for amorphous. JWST data alone cannot be used to determine the phase. Annealing and nucleation kinetics need to be considered, as discussed in Section 4. Section 3 provides another key detail: the dayside temperature reaches roughly 1191 K (Southworth), and grains remain suspended there for long periods, exceeding 1000 days

according to Parmentier et al. Hence, residence time is not the bottleneck for crystallisation. Furthermore, since the HD 189733b model presented by Lee et al. shows SiO_2 grains piling up at low dayside latitudes specifically, it is fair to treat that long residence time as something that happens on the dayside. Since residence time is not the limiting factor, the only variable that is left to examine is temperature. Hallenbeck et al. show that the crystallisation stall gets cleared within minutes above roughly 1100 K. HD 189733b's dayside equilibrium temperature is well above this value. Nucleation theory explains why this happens in practice: Lee et al.'s TiO_2-seeded heterogeneous nucleation model, built specifically for this planet, shows seeded nucleation moves fast. Lab work on terrestrial glass ceramics (Wang et al.; Jean et al.) supports this general principle. Homogeneous nucleation proceeds slowly while heterogeneous nucleation occurs quickly, though the precise rates measured in lab settings do not translate directly to a vapour condensing in a planetary atmosphere.

This kinetic prediction is consistent with other observations as well. Mid-infrared spectroscopy can distinguish between amorphous and crystalline phases (Mullens and Lewis). The cooler WASP-107b (736 to 770 K) shows an amorphous state, whereas the hotter WASP-17b (1750 to 1771 K) reveals a tentative crystalline preference at 3.5-4.2σ. HD 189733b's atmospheric temperature at 1191 K sits between those of WASP-107b and WASP-17b, but favours crystallisation. To be clear, this is supporting precedent, not solid proof. The real case for crystalline still rests on the kinetics in Section 4, and the evidence from WASP-17b is statistically weaker than HD 189733b's 6σ composition detection.

There are three key problems that need to be addressed. First, there is a lack of high-temperature laboratory data. Zeidler et al. give the only reported optical constants for crystalline SiO_2 at high temperatures. Potapov and Bouwman confirm this is a field-wide problem, with laboratory measurements only reaching about 1000 K when the relevant range is 300 to 2000 K. Grant et al. ran into the same hurdle studying WASP-17b. Second, the whole kinetics argument is dayside-specific, and the nightside temperature is about 240 K cooler (Knutson et al.). While Lee et al. found that grains concentrate at low dayside latitudes, suggesting they remain warm, this result does not reveal what fraction of a grain's 1000+ day lifetime is spent on the hot side versus the cold side. Third, there is a mismatch in annealing timescales. Hallenbeck et al. (2000) show crystallisation happening in minutes above roughly 1100 K. The measurements from Fabian et al. span from more than 50 hours at 1000 K to 3–5 minutes at 1121 K, indicating that Hallenbeck's minute-scale timeframe applies specifically to the hotter regime corresponding to HD 189733b's dayside. While this alignment has not been systematically cross-checked against every operational condition described in both studies, it offers the most physically consistent explanation available. While these three caveats do not

fundamentally invalidate the central conclusions of this work, they constitute standard methodological limitations that need to be considered together.

There are several remaining open questions. Higher-resolution atmospheric circulation models could identify where nucleation is occurring rather than assuming the uniform dayside

conditions adopted in this paper. In situ laboratory measurements of silicates above 1000 K remain unavailable (Potapov and Bouwman, 2022), requiring temperature-dependent kinetics to be estimated through extrapolation. Future JWST observations with higher resolution or broader wavelength coverage may also help distinguish amorphous and crystalline silicates more confidently.

The DRIFT model, a kinetic cloud formation model developed for HD 189733b (Lee et al., 2016), is not entirely silent on crystalline phases. The authors argue that long atmospheric residence times provide sufficient time for cloud particles to rearrange into crystalline forms. This raises an important question about the kinetic framework developed in Section 3. It examines whether the framework provides independent evidence for crystallisation, or primarily quantifies a mechanism already assumed by DRIFT.

Current evidence is consistent with crystallisation occurring in HD 189733b's atmosphere, but significant uncertainties remain. Resolving them will require improved laboratory measurements, atmospheric models, and observational constraints.

Works Cited

Alderson, L., et al. “A Comprehensive Analysis of WASP-17b’s Transmission Spectrum from Space-Based Observations.” Monthly Notices of the Royal Astronomical Society, vol. 512, no. 3, Apr. 2022, pp. 4185–209. arXiv.org, https://doi.org/10.1093/mnras/stac661.

Brouwers, M. G., and C. W. Ormel. “How Planets Grow by Pebble Accretion - II. Analytical Calculations on the Evolution of Polluted Envelopes.” Astronomy & Astrophysics, vol. 634, Feb. 2020, p. A15. www.aanda.org, https://doi.org/10.1051/0004-6361/201936480.

Charbonneau, David, et al. “Detection of an Extrasolar Planet Atmosphere.” The Astrophysical Journal, vol. 568, no. 1, Mar. 2002, pp. 377–84. DOI.org (Crossref), https://doi.org/10.1086/338770.

Dawson, Rebekah I., and John Asher Johnson. “Origins of Hot Jupiters.” Annual Review of Astronomy and Astrophysics, vol. 56, no. 1, Sept. 2018, pp. 175–221. arXiv.org, https://doi.org/10.1146/annurev-astro-081817-051853.

Dyrek, Achrene, et al. “SO2, Silicate Clouds, but No CH4 Detected in a Warm Neptune.” Nature, vol. 625, no. 7993, Jan. 2024, pp. 51–54. arXiv.org, https://doi.org/10.1038/s41586-023-06849-0.

Etangs, A. Lecavelier des, et al. “Rayleigh Scattering in the Transit Spectrum of HD 189733b.” Astronomy & Astrophysics, vol. 481, no. 2, Apr. 2008, pp. L83–86. arXiv.org, https://doi.org/10.1051/0004-6361:200809388.

Evans, Thomas M., et al. “THE DEEP BLUE COLOR OF HD 189733b: ALBEDO MEASUREMENTS WITH HUBBLE SPACE TELESCOPE/SPACE TELESCOPE IMAGING SPECTROGRAPH AT VISIBLE WAVELENGTHS.” The Astrophysical Journal Letters, vol. 772, no. 2, July 2013, p. L16. Institute of Physics, https://doi.org/10.1088/2041-8205/772/2/L16.

Fabian, D., et al. “Amorphous to Crystalline Transition of Magnesium Silicate and Silica Nanoparticles.” vol. 456, 2000, p. 347. NASA ADS, https://ui.adsabs.harvard.edu/abs/2000ESASP.456..347F.

Gail, H. P., and E. Sedlmayr. “Mineral Formation in Stellar Winds. I. Condensation Sequence of Silicate and Iron Grains in Stationary Oxygen Rich Outflows.” Astronomy and Astrophysics, vol. 347, July 1999, pp. 594–616.

Grant, David, et al. “JWST-TST DREAMS: Quartz Clouds in the Atmosphere of WASP-17b.” The Astrophysical Journal Letters, vol. 956, no. 2, Oct. 2023, p. L32. Institute of Physics, https://doi.org/10.3847/2041-8213/acfc3b.

Hallenbeck, Susan L., et al. “Evolving Optical Properties of Annealing Silicate Grains: From Amorphous Condensate to Crystalline Mineral.” The Astrophysical Journal, vol. 535, no. 1, May 2000, p. 247. Institute of Physics, https://doi.org/10.1086/308810.

Harker, David E., and Steven J. Desch. “Annealing of Silicate Dust by Nebular Shocks at 10 AU.” The Astrophysical Journal, vol. 565, no. 2, Jan. 2002, p. L109. Institute of Physics, https://doi.org/10.1086/339363.

Helling, Ch., et al. “The Mineral Clouds on HD 209458b and HD 189733b.” Monthly Notices of the Royal Astronomical Society, vol. 460, no. 1, July 2016, pp. 855–83. Silverchair, https://doi.org/10.1093/mnras/stw662.

Helling, Ch., and F. J. M. Rietmeijer. “Glittery Clouds in Exoplanetary Atmospheres?” International Journal of Astrobiology, vol. 8, Jan. 2009, pp. 3–8. NASA ADS, https://doi.org/10.1017/S1473550408004382.

Hill, Hugh G. M., et al. “Constraints on Nebular Dynamics and Chemistry Based on Observations of Annealed Magnesium Silicate Grains in Comets and in Disks Surrounding Herbig Ae/Be Stars.” Proceedings of the National Academy of Sciences, vol. 98, no. 5, Feb. 2001, pp. 2182–87. DOI.org (Crossref), https://doi.org/10.1073/pnas.051530998.

Inglis, Julie, et al. “Quartz Clouds in the Dayside Atmosphere of the Quintessential Hot Jupiter HD 189733 b.” The Astrophysical Journal Letters, vol. 973, no. 2, Sept. 2024, p. L41. Institute of Physics, https://doi.org/10.3847/2041-8213/ad725e.

Jean, Jau‐Ho, et al. “Devitrification Kinetics and Mechanism of K_2O–CaO–SrO–BaO–B_2O_3 –SiO_2 Glass‐Ceramic.” Journal of the American Ceramic Society, vol. 84, no. 6, June 2001, pp. 1354–60. DOI.org (Crossref), https://doi.org/10.1111/j.1151-2916.2001.tb00841.x.

Kennedy, Grant M., and Scott J. Kenyon. “Planet Formation around Stars of Various Masses: The Snow Line and the Frequency of Giant Planets.” The Astrophysical Journal, vol. 673, no. 1, Jan. 2008, p. 502. iopscience.iop.org, https://doi.org/10.1086/524130.

Knutson, Heather A., Nikole Lewis, et al. “3.6 AND 4.5 Mm PHASE CURVES AND EVIDENCE FOR NON-EQUILIBRIUM CHEMISTRY IN THE ATMOSPHERE OF EXTRASOLAR PLANET HD 189733b.” The Astrophysical Journal, vol. 754, no. 1, June 2012, p. 22. Institute of Physics, https://doi.org/10.1088/0004-637X/754/1/22.

Knutson, Heather A., David Charbonneau, et al. “A Map of the Day-Night Contrast of the Extrasolar Planet HD 189733b.” Nature, vol. 447, no. 7141, May 2007, pp. 183–86. arXiv.org, https://doi.org/10.1038/nature05782.

Krenn, A. F., et al. “The Geometric Albedo of the Hot Jupiter HD 189733b Measured with CHEOPS.” Astronomy & Astrophysics, vol. 672, Apr. 2023, p. A24. DOI.org (Crossref), https://doi.org/10.1051/0004-6361/202245016.

Lee, E., et al. “Dynamic Mineral Clouds on HD 189733b - I. 3D RHD with Kinetic, Non-Equilibrium Cloud Formation.” Astronomy & Astrophysics, vol. 594, Oct. 2016, p. A48. www.aanda.org, https://doi.org/10.1051/0004-6361/201628606.

Madhusudhan, Nikku, et al. Exoplanetary Atmospheres. 2014. arXiv.org, https://doi.org/10.2458/azu_uapress_9780816531240-ch032.

Mullens, Elijah, and Nikole K. Lewis. “Silicate Sundogs: Probing the Effects of Grain Directionality in Exoplanet Observations.” The Astrophysical Journal Letters, vol. 988, no. 2, Aug. 2025, p. L43. arXiv.org, https://doi.org/10.3847/2041-8213/ade885.

Parmentier, Vivien, et al. “3D Mixing in Hot Jupiters Atmospheres - I. Application to the Day/Night Cold Trap in HD 209458b.” Astronomy & Astrophysics, vol. 558, Oct. 2013,

p. A91. www.aanda.org, https://doi.org/10.1051/0004-6361/201321132.

Pont, F., H. Knutson, et al. “Detection of Atmospheric Haze on an Extrasolar Planet: The 0.55-1.05 Mm Transmission Spectrum of HD 189733b with the Hubble Space Telescope: Transmission Spectrum of HD 189733.” Monthly Notices of the Royal Astronomical Society, vol. 385, no. 1, Feb. 2008, pp. 109–18. DOI.org (Crossref), https://doi.org/10.1111/j.1365-2966.2008.12852.x.

Pont, F., D. K. Sing, et al. “The Prevalence of Dust on the Exoplanet HD 189733b from Hubble and Spitzer Observations.” Monthly Notices of the Royal Astronomical Society, vol. 432, no. 4, July 2013, pp. 2917–44. Silverchair, https://doi.org/10.1093/mnras/stt651.

Potapov, Alexey, and Jeroen Bouwman. “Importance of Laboratory Experimental Studies of Silicate Grains for Exoplanet Atmosphere Characterization.” Frontiers in Astronomy and Space Sciences, vol. 9, Oct. 2022. Frontiers, https://doi.org/10.3389/fspas.2022.912302.

Powell, Diana, et al. “Formation of Silicate and Titanium Clouds on Hot Jupiters.” The Astrophysical Journal, vol. 860, no. 1, June 2018, p. 18. Institute of Physics, https://doi.org/10.3847/1538-4357/aac215.

Rappaport, Saul, et al. “THE ROCHE LIMIT FOR CLOSE-ORBITING PLANETS: MINIMUM DENSITY, COMPOSITION CONSTRAINTS, AND APPLICATION TO THE 4.2 Hr PLANET KOI 1843.03.” The Astrophysical Journal Letters, vol. 773, no. 1, July 2013, p. L15. Institute of Physics, https://doi.org/10.1088/2041-8205/773/1/L15.

Sing, D. K., et al. “Hubble Space Telescope Transmission Spectroscopy of the Exoplanet HD 189733b: High-Altitude Atmospheric Haze in the Optical and near-UV with STIS.” Monthly Notices of the Royal Astronomical Society, vol. 416, no. 2, Sept. 2011, pp. 1443–55. arXiv.org, https://doi.org/10.1111/j.1365-2966.2011.19142.x.

Southworth, John. “Homogeneous Studies of Transiting Extrasolar Planets – IV. Thirty Systems with Space-Based Light Curves.” Monthly Notices of the Royal Astronomical Society, vol. 417, no. 3, Nov. 2011, pp. 2166–96. Silverchair, https://doi.org/10.1111/j.1365-2966.2011.19399.x.

Uhlmann, Donald R. “Nucleation and Crystallization in Glass-Forming Systems.” Glass … Current Issues, edited by A. F. Wright and J. Dupuy, Springer Netherlands, 1985, pp. 1–20. Springer Link, https://doi.org/10.1007/978-94-009-5107-5_1.

Wakeford, H. R., and D. K. Sing. “Transmission Spectral Properties of Clouds for Hot Jupiter Exoplanets.” Astronomy & Astrophysics, vol. 573, Jan. 2015, p. A122. www.aanda.org, https://doi.org/10.1051/0004-6361/201424207.

Wang, Wenbin, et al. “Crystallization Behavior and Properties of Quartz Glass-Ceramics Synthesized from Desert Sand.” Ceramics International, vol. 50, no. 1, Jan. 2024, pp. 370–83. DOI.org (Crossref), https://doi.org/10.1016/j.ceramint.2023.10.111.

Zeidler, S., et al. “Optical Constants of Refractory Oxides at High Temperatures - Mid-Infrared Properties of Corundum, Spinel, and α-Quartz, Potential Carriers of the 13 Mm Feature.” Astronomy & Astrophysics, vol. 553, May 2013, p. A81. www.aanda.org, https://doi.org/10.1051/0004-6361/201220459.

Semiconductors and Strategic Alignment: How U.S.-China Semiconductor Competition Has Influenced Taiwan's Foreign Policy Since 2018 By Nidhi Yellamraju

Abstract

The world's most advanced technologies increasingly depend on an island smaller than many U.S. states, placing Taiwan at the center of U.S.-China competition. Taiwan produces many of the world's most advanced semiconductors, so its chip industry has become tied to a much larger political question: what happens when the United States and China compete over technology that both depend on? This paper asks how that competition has affected Taiwan's foreign policy and its strategic alignment with the United States since 2018. The argument is that the chip conflict did not create Taiwan's relationship with Washington. That relationship was already built on security ties and decades of cooperation in the technology sector. What changed after 2018 was the scope of the relationship. Export controls, TSMC's investment in Arizona, protection of sensitive technology, and supply-chain policy became part of U.S.-Taiwan cooperation. Existing scholarship explains why Taiwan matters to the United States and China, especially through debates over the "silicon shield" and semiconductor dependence (Chou 2023; Diamond, Ellis, and Schell 2023). Less attention is given to how Taiwan itself responds to being placed at the center of this competition. Using a qualitative case study of policy documents, trade data, TSMC decisions, and existing scholarship, the paper identifies a strategy of selective alignment. Taiwan has moved closer to the United States in areas where advanced technology is treated as a security issue, while still keeping major economic ties with China and trying to preserve the domestic chip industry that gives Taiwan much of its strategic value.

Key Words U.S.-China semiconductor competition, Taiwan, foreign policy, strategic alignment, selective alignment, geoeconomics, economic security, supply-chain resilience, technological interdependence.

1. Introduction

Although the semiconductor is small in size, the political debate surrounding it is not. Chips with advanced technology are required for mobile phones, automobiles, artificial intelligence, and modern weapons. The industry responsible for producing them is located in a number of countries and no one country carries out all the stages. American companies are particularly strong in the areas of chip design, software, and in the production of key manufacturing equipment. Firms in Japan and the Netherlands provide other tools which are hard to replace. South Korea leads in the production of memory chips. Through TSMC and its broader semiconductor industry, Taiwan has become especially significant in the manufacture of advanced logic chips. Bown and Wang (2024) claim that this specialization reduced costs and increased efficiency, but it also led to dependence. Thadani and Allen (2023) make the same point regarding the Indo-Pacific region, where a large amount of the world's semiconductor production is concentrated. For a long time, this dependence was mainly seen as an economic

advantage. As relations between the U.S. and China deteriorated, it began to appear as a security issue as well.

It became obvious in 2018 when the U.S.-China trade war started with the imposition of tariffs and disputes about trade practices, but technology soon became one of the central issues. Huawei was an early case in point. The United States imposed tighter restrictions on the Chinese company and, in May 2020, amended the rule regarding foreign-made direct products so that some chips manufactured outside the United States could still be banned if American technology had been used in their production (U.S. Department of Commerce 2020; Bown 2020). This change was important since a number of advanced semiconductor companies outside the United States still depended on American equipment or software. Farrell and Newman (2019) refer to this kind of influence as "weaponized interdependence". A production network can appear to be global and decentralized, yet still rely on a small number of points that a single government can control. In the case of semiconductors, the United States had sufficient control over some of those key points to be able to affect companies located far outside its own borders. Taiwan found itself at the heart of these changes. Although TSMC had become the world's leading manufacturer of advanced logic chips, it had not achieved this status independently. According to Miller (2022), the company's rise can be attributed to Taiwanese industrial policy, a high level of technical expertise, and its close associations with American technology firms. Keller and Pauly (2017) also demonstrate how public institutions and private companies contributed to the development of Taiwan's semiconductor industry. This background has placed Taiwan in an unusual situation: it is economically significant since so many companies depend on its factories, yet it is also politically vulnerable because China claims the island and has not ruled out the use of force. At the same time, the United States maintains close but unofficial relations with Taipei. After chips were incorporated into national-security policy, the decisions that TSMC made regarding its customers, the sites of its factories, and the technology it used could no longer be clearly separated from Taiwan's foreign relations.

A great deal of scholarly work on Taiwan comes from outside the region. For instance, the 'silicon shield' debate considers whether the world's reliance on chips made in Taiwan reduces the likelihood of conflict since a war would be economically disastrous. Chou (2023) examines that argument but also highlights its shortcomings, while Diamond, Ellis, and Schell (2023) link the importance of Taiwan's semiconductors to broader U.S.-China security considerations. Other researchers look into why U.S. export controls are effective or why governments are attempting to build supply chains with trusted partners (Bown 2020; Benson, Quitzon, and Reinsch 2023; Yoon 2023). These are valuable questions. Yet, they tend to start with either Washington's or Beijing's perspective—asking how much Taiwan matters to them, what they might lose, or how they might respond. The present focus, by contrast, is different. It centers on what Taiwan does when its economic significance becomes an element of a geopolitical rivalry.

It doesn't follow that previous research overlooked Taiwan's decisions. Studies on the cross-Strait relationship have demonstrated that Taiwanese institutions are able to reduce

Chinese economic influence when officials perceive political risk. For example, Lee and Knoerich (2023) look at the extent to which mainland Chinese investment can be used as a tool of economic statecraft. Before the semiconductor conflict intensified, Taiwan was also working on diversifying its economic relationships. The New Southbound Policy, introduced in 2016, was intended for strengthening economic links with countries in South and Southeast Asia and for reducing its overreliance on China (Kennedy et al. 2018; Tso and Jung 2018). The gap in the analysis is what occurred when the earlier concern regarding dependence came up against the more recent struggle over advanced technology. After 2018, Taiwan was no longer just concerned about how much economic exposure to China was safe; it also had to consider which parts of its technology sector had become so politically sensitive as to no longer be treated as ordinary trade.

The paper's research question is thus: How has the competition between the United States and China over semiconductors influenced Taiwan's foreign policy and its strategic alignment with the United States since 2018? The term 'strategic alignment' has a fairly particular meaning here. It does not indicate that Taiwan and the United States agree on all matters, nor does it mean that they have a formal alliance; it refers to ongoing coordination in areas which both sides have come to regard as strategically important. In this instance, such areas include semiconductor investment, export controls, the protection of critical technology, supply-chain planning, and regular economic cooperation. The claim is that the struggle over chips has caused these issues to become more deeply embedded in the relationship between the United States and Taiwan and has given them a great deal more political significance than they had prior to 2018. At the same time, Taiwan has not merely joined an American economic group. China remains Taiwan's major trading partner since mainland China and Hong Kong together accounted for 31.7 percent of Taiwan's exports in 2024 (Taiwan Ministry of Finance 2025). There are good reasons for Taiwan to keep advanced research and manufacturing based on the island since the value of its semiconductor position depends on a broader network of skilled workers, suppliers, infrastructure, and accumulated manufacturing knowledge (Reinsch and Whitney 2025). This situation is referred to here as selective alignment. While Taiwan has moved nearer to the United States in areas where technology is closely linked to national security, it has not regarded all its economic links with China as posing a security threat. This difference helps to explain how Taiwan can maintain strict technology controls and accept U.S. investment while at the same time having reasons to keep the core of advanced production based on the island. The paper begins by reviewing the main literature and presents this framework. It then looks at the relationship before 2018, the restrictions imposed on Huawei, TSMC's expansion into Arizona, Taiwan's technology-protection policies, diversification, and the growth of bilateral economic institutions. The last section puts these cases together in order to evaluate how much room Taiwan has to influence the relationship.

2. Literature Review

Geoeconomics offers a useful starting point when it comes to the semiconductor conflict. Blackwill and Harris (2016) claim that economic measures such as trade, investment, and industrial policy can be employed in order to achieve strategic objectives. Farrell and Newman (2019) contribute a further dimension by concentrating on networks; they maintain that some states acquire power since they control parts of a global system which other states find difficult to avoid. Bown's (2020) examination of Huawei illustrates how this works in the case of semiconductors. A chip may be produced in Taiwan, but the manufacturing process could still depend on American software or machinery. This dependence allowed the United States to exert influence over production that occurred in a different location. The semiconductor conflict is therefore not merely a tale of governments getting more involved in business. It is also about the way the structure of the industry confers unusual leverage on certain governments. Taiwan doesn't quite match the model since it controls a different kind of chokepoint. Although it doesn't have the same legal reach as the United States, other countries still place great reliance on the manufacturing facilities in Taiwan. Reinsch and Whitney (2025) point out that Taiwan's advantage lies not just in TSMC but also in its experienced engineers, its specialized suppliers, the infrastructure it has, and the knowledge that has accumulated over the years. This results in an unequal relationship without making it one-sided. U.S. regulations can restrict the products that a Taiwanese company is allowed to sell, as the Huawei case demonstrates, but American technology companies themselves also place heavy reliance on manufacturing in Taiwan (Bown 2020; Reinsch and Whitney 2025). Taiwan is thus subject to the very network that makes it valuable. This kind of tension is one reason why Taipei has both incentives to work closely with Washington and to safeguard the industrial base which gives Taiwan its bargaining power. The study of supply-chain resilience brings with it another issue. The efficiency of semiconductor production has been partly achieved because different countries have specialised in the areas where they are most capable, but this has also made it hard to quickly reproduce the system (Miller 2022; Bown and Wang 2024). Benson, Quitzon and Reinsch (2023) claim that it is not realistic to say that every country should produce everything itself. Cooperation is essential. Taiwan gains when foreign investment brings its companies nearer to their key partners, but the interests of Taiwan and the United States are not the same. U.S. policy aims at reducing the risk arising from its heavy reliance on production in Taiwan, while Taiwan's strategic value is in part due to the fact that its advanced manufacturing cluster is hard to replace (Benson, Quitzon and Reinsch 2023; Reinsch and Whitney 2025). The policy question is therefore not merely whether overseas fabs are good or bad. It is how diversification can enhance supply security without undermining the concentration of skills and production which gives Taiwan influence.

The "silicon shield" literature also raises this question. Chou (2023) looks at whether reliance on chips from Taiwan might serve to prevent conflict, while Diamond, Ellis, and Schell (2023) investigate the role of semiconductors in the wider security relationship between the United States, Taiwan, and China. The main point is that a major disruption in production in

Taiwan would result in costs that would extend well beyond the island, thereby giving other countries a strong incentive to maintain stability. Nevertheless, this line of reasoning reflects more on the assessments of other governments than it does on Taiwan's own policy. TSMC's expansion overseas, Taiwan's laws which protect sensitive technology, and its economic discussions with Washington all involve decisions regarding how the importance of semiconductors should be handled. The later examples in this paper demonstrate that concentrated chip production does not automatically translate into political leverage; its value depends on the amount of capacity Taiwan shares abroad, the knowledge it protects, and whether the domestic ecosystem remains hard to replace.

Research across the Straits helps to show that Taiwan had already been thinking in this way before 2018. The New Southbound Policy was not a result of U.S. semiconductor restrictions; in fact, it was introduced earlier, in 2016, as a response to concerns over having too much of its economic activity linked to China (Kennedy et al. 2018; Tso and Jung 2018). Huang (2018) points out that in practice diversification proved to be more difficult than it appeared in policy terms, mainly since the economic attraction of China had remained strong. Lee and Knoerich (2023) also demonstrate that Taiwanese institutions will intervene when Chinese investment is regarded as politically risky. These earlier policies are important since they prevent the argument from becoming too simplistic. The semiconductor conflict did not suddenly cause Taiwan to become wary of its dependence on China; rather, it gave older concerns a new security dimension and tied them more closely to U.S. policy.

Put together, these bodies of literature allow for a Taiwan-focused argument. It has already been shown that the United States has chokepoint power and that it is difficult to replace Taiwan's chip factories (Farrell and Newman 2019; Reinsch and Whitney 2025). Studies comparing both sides also indicate that Taiwanese leaders have for a long time been concerned about their economic dependence on China (Kennedy et al. 2018; Lee and Knoerich 2023). However, the way in which these pressures have interacted in Taiwan's decisions after 2018 is not well developed. Here, strategic alignment is viewed as a pattern that can be seen in a number of different decisions rather than as something that is proven by a single major event. Selective alignment accounts for why this pattern has its limits. Although Taiwan has redirected some of its most sensitive technology relationships towards the United States and other reliable partners, it has not attempted to end trade with China or to shift the centre of its semiconductor industry overseas. These limitations are part of the pattern that is being explained, not proof that the alignment is unimportant.

3. From Commercial Interdependence to Strategic Alignment

Before 2018 Taiwan had already become deeply involved with the American technology industry, a connection that had developed over many decades. The Taiwanese government made substantial investments in higher-value industries, and organizations like the Industrial Technology Research Institute played a role in building up the technical knowledge required for semiconductor manufacturing. TSMC was established in 1987 and introduced the foundry

model, which enabled companies to design chips without having to own the factories that produced them (Miller 2022; Keller and Pauly 2017). This model closely linked Taiwanese manufacturing with foreign chip designers, such as major American companies (Miller 2022). Therefore, when the competition between U.S. and Chinese technology intensified, Taiwan was not having to decide between two entirely separate systems; it was already part of a technology network with strong ties to the United States. The change that occurred after 2018 was that the commercial relationships began to have much more obvious political and security implications.

Huawei is a clear example of this. Although the company had previously been one of TSMC's important customers, U.S. restrictions made it much more difficult to maintain that relationship. In 2020, the United States broadened its regulations so that foreign-made chips could still be restricted if American technology was used in the manufacturing process (U.S. Department of Commerce 2020). TSMC ceased taking new orders from Huawei when the rules came into effect (Bown 2020). At first sight, this appears to be a business decision prompted by American law. However, it also shows just how little room a company such as TSMC had to remain neutral once the semiconductor industry became part of the competition between the United States and China. Although manufacturing took place in Taiwan, the rules concerning access to key technology applied across borders.

The Huawei situation should not be seen as simply a case in which the United States had power and Taiwan had none. Although TSMC relied on American tools and customers, American technology companies in turn placed great importance on TSMC's advanced manufacturing (Bown 2020; Reinsch and Whitney 2025). The mutual dependence did not result in an equal relationship since Washington still possessed legal tools that Taipei did not, a point clearly shown by the 2020 rule change (U.S. Department of Commerce 2020). On the other hand, it did mean that the United States also had a strong interest in ensuring that TSMC remained healthy and linked to the American technology sector. Taiwan found itself in an uncomfortable but valuable position: it could be pressured because of its dependence on the network, while the network in turn depended on Taiwan. This combination is what helps to explain why the competition in the semiconductor sector brought Taiwan closer to U.S. technology policy without merely treating Taiwan as a passive target of American rules.

Yet the fact that TSMC has complied does not by itself prove that Taiwan's foreign policy has changed. After all, TSMC is a private enterprise and had no choice but to respond to the rules in question. The bigger picture gains greater credibility when the Huawei decision is compared with what came afterwards. In 2022 Taiwan revised its National Security Act in order to guard its designated core technologies against economic espionage and unauthorized overseas use (Taiwan Ministry of Justice 2022). Also, U.S.-Taiwan talks started to regularly cover supply-chain resilience, economic coercion, and technology cooperation (U.S. Department of State 2022; U.S. Department of State 2023). It is important that these decisions at the level of the state indicate that the change was not merely the result of one company's decision to comply. The competition in the semiconductor sector has been altering the policy environment surrounding

TSMC, and this environment is now increasingly linking economic decisions to Taiwan's relationship with the United States.

4. TSMC's U.S. Expansion & the Politics of Resilience

The Arizona investment caused the relationship to go beyond simply reacting to U.S. restrictions. The project had an initial amount of about $12 billion, after which it was significantly expanded. In November 2024, the U.S. Department of Commerce approved up to $6.6 billion in direct CHIPS Act funding and up to $5 billion in loans for more than $65 billion worth of planned investment by TSMC in three fabs in Arizona (U.S. Department of Commerce 2024). The United States referred to the project in terms of economic and national security. This wording is important; TSMC was not being seen as just another foreign company setting up a factory. Rather, its investment had become part of the American government's strategy to decrease its dependence on advanced chip production based in East Asia.

For Taiwan, the benefits from the investment in Arizona are real. The project has brought TSMC into closer contact with its U.S. customers, suppliers, and the industrial policy of the United States, and the U.S. government has pledged billions of dollars in direct financial support and loans for the expansion (U.S. Department of Commerce 2024). As a result, Washington now has a direct economic interest in the company's long-term success. This is in contrast to the situation with Huawei, where TSMC had mainly had to adjust to a rule set by Washington. In the case of Arizona, it is a long-term investment shared by both the company and the U.S. government. Since Taiwan has no formal diplomatic relations with the United States, such a large project establishes another lasting link in what is already an important relationship. The political value arises from the material relationship itself: factories, subsidies, suppliers, and customers all provide reasons for continued interaction that do not rely on formal diplomatic recognition.

There is also a drawback. The United States is seeking to have more advanced chip production carried out within the country since a heavy reliance on Taiwan poses a supply risk, a concern that is central to U.S. semiconductor-resilience policy (Benson, Quitzon, and Reinsch 2023; U.S. Department of Commerce 2024). Taiwan, on the other hand, runs the risk of becoming easier to replace. Reinsch and Whitney (2025) maintain that Taiwan's strength lies not just in its individual factories but in a dense network of skilled workers, suppliers, infrastructure, and experience. The 'silicon shield' literature is also based on the idea that Taiwan's production is unusually difficult to substitute (Chou 2023; Diamond, Ellis, and Schell 2023). Should too much of the leading-edge ecosystem be moved overseas, Taiwan could suffer a loss of economic significance which causes other countries to continue investing in its stability. Arizona thus generates both cooperation and tension.

Arizona thus shows the principle of selective alignment more clearly than a simple narrative of Taiwan becoming closer to the United States. The project serves U.S. supply-chain objectives and enhances TSMC's position within the American economy, but Reinsch and Whitney (2025) maintain that U.S. resilience should be achieved without undermining Taiwan's

semiconductor ecosystem. This distinction is important for Taiwan; although some overseas production can help to strengthen relations with customers and partners, the fact that advanced research, suppliers, and manufacturing are concentrated at home continues to be a major factor in Taiwan's economic significance. It is selective alignment that more accurately captures this balance, and while cooperation with Washington can be increased, Taiwan still has good reasons to safeguard the part of the semiconductor system which makes the island difficult to replace.

5. Selective Alignment: Technology Protection and Supply-Chain Diversification

The laws enacted by Taiwan themselves show that the government has begun to regard semiconductor knowledge as a matter of security. In 2022, Taiwan revised its National Security Act to make it a crime to engage in economic espionage and to use trade secrets relating to designated national core key technologies without authorization abroad. The Ministry of Justice described the change as a measure designed to safeguard Taiwan's high-technology competitiveness and to ensure that sensitive know-how does not get into the hands of foreign hostile forces, including those linked to mainland China (Taiwan Ministry of Justice 2022). This was not merely an effort to protect companies against ordinary theft; the law had included certain kinds of industrial knowledge within the realm of national security.

There are valid reasons for Taiwan to guard its sensitive technology even when this aligns with American concerns. Although advanced semiconductor manufacturing relies on equipment, it also depends on engineers, specialized suppliers, and production knowledge that has been built up over the years (Miller 2022; Reinsch and Whitney 2025). Should key knowledge pass to a competitor, Taiwan would suffer not just the loss of a private business's competitive edge but also a weakening of the industrial ecosystem which is what makes the island strategically significant. Yoon (2023) refers to "techno-geopolitics" when describing how advanced industries are becoming seen as security assets, and Taiwan is an example of this. It is not necessary for Taiwan to copy all of the U.S. export rules in order for alignment to occur. What is important is the visible overlap—the fact that Washington limits access to sensitive semiconductor capabilities while Taipei has established its own legal restrictions concerning technologies it considers nationally important (U.S. Department of Commerce 2020; Taiwan Ministry of Justice 2022).

Diversification follows the same course, even though the policy began earlier. The New Southbound Policy in Taiwan had already been aiming at reducing its concentrated exposure to China before the present chip rivalry (Kennedy et al. 2018; Tso and Jung 2018). After 2018, though, the goal of diversification fitted better with the U.S. effort to build resilient and trusted supply chains. Yoon (2023) cites the proposed "Fab 4" cooperation between the United States, Taiwan, Japan, and South Korea as an example of political trust becoming incorporated into semiconductor planning. This coincidence made Taiwan's earlier policy more valuable. While expanding economic relations beyond China could reduce dependence, it might also enhance Taiwan's position among the countries that were becoming more concerned about technology security.

The figures also show that the term decoupling goes too far. The Taiwan Ministry of Finance states that mainland China and Hong Kong received 43.9 percent of Taiwan's exports in 2020; by 2024 this proportion had dropped to 31.7 per cent (Taiwan Ministry of Finance 2025). A fall of over twelve percentage points is significant. Yet 31.7 per cent is still a very large part of total exports. Taiwan is clearly less dependent on the Chinese market than it was at the peak in 2020, but it has by no means cut itself off. The trend is also uneven, with the most stringent restrictions being found in areas relating to advanced technology and other sectors connected with national security, not in all types of trade. This is the reason why selective alignment is more in line with the evidence than decoupling.

Taken together, the cases are quite consistent. Although TSMC ceased taking new orders from Huawei following the U.S. restrictions, Taiwan has not withdrawn from the Chinese market (Bown 2020; Taiwan Ministry of Finance 2025). While TSMC is developing advanced manufacturing capacity in Arizona, the broader semiconductor ecosystem in Taiwan is still based on the island (U.S. Department of Commerce 2024; Reinsch and Whitney 2025). The government has also imposed tighter controls over sensitive technologies without making normal commercial exchanges impossible (Taiwan Ministry of Justice 2022). The pattern is not an attempt to end interdependence; rather, it is an effort to separate the aspects of interdependence that have become security-sensitive from those that are mainly commercial. That is the practical meaning of selective alignment as set out in this paper.

6. Economic Diplomacy and Institutionalized Alignment

The pattern is less easy to regard as merely the result of individual corporate decisions because of changes in U.S.-Taiwan economic diplomacy. The Economic Prosperity Partnership Dialogue was started in 2020 under the AIT and TECRO. At the third dialogue in 2022 the agenda covered supply-chain resilience and economic coercion (U.S. Department of State 2022). The two sides had also reached a science and technology agreement in 2020, and the first formal Science and Technology Cooperation Dialogue in 2023 involved semiconductor research and applications (U.S. Department of State 2023). Although these meetings are not major events, it is in part for this reason that they are important; they have made technology security a matter that officials discuss regularly rather than only when a crisis compels them to react.

It is because regular institutions alter the meaning of alignment since they demonstrate that cooperation is not a one-off occurrence but one that is repeated. The 2022 Economic Prosperity Partnership Dialogue covered supply-chain resilience and economic coercion, and the 2023 Science and Technology Cooperation Dialogue included semiconductor research and applications (U.S. Department of State 2022; U.S. Department of State 2023). By doing so, technology security is placed within a continuing government-to-government relationship. Although the United States and Taiwan still have different interests, particularly when it comes to the long-term balance between overseas production and maintaining advanced capacity on the island, there are now established regular channels for discussing semiconductor and supply-chain

issues. This provides stronger evidence of strategic alignment than a single investment announcement or one export-control decision.

The U.S.-Taiwan Initiative on 21st-Century Trade extends this network of institutional arrangements. The first agreement came into force in December 2024 and included matters like customs procedures and regulatory practices (Office of the United States Trade Representative 2024). These rules do not concern semiconductor controls, so it would be incorrect to say that the chip dispute caused them. The point is that as technology security has become more important, U.S.-Taiwan economic cooperation has also expanded. The relationship is not any longer based merely on TSMC or on defense issues; there are now a number of ordinary economic channels through which the two sides can collaborate.

Cooperation of that kind is particularly advantageous for Taiwan since its diplomatic status is unusual. Although technological cooperation will not lead to formal recognition, it does provide practical grounds for governments to continue dealing with Taipei. The U.S. policy documents increasingly refer to Taiwan's semiconductor ecosystem as being important for American economic and supply-chain security, and technology and resilience have now become regular topics in U.S.-Taiwan discussions (Reinsch and Whitney 2025; U.S. Department of State 2022; U.S. Department of State 2023). These interests allow working relationships to develop even in the absence of formal diplomatic relations. Taiwanese officials also take part in setting up the forums and determining the agenda. The central role of semiconductors does not eliminate Taiwan's diplomatic constraints, but it does give Taipei more specific avenues by which to maintain a broader relationship with Washington.

7. Discussion: Taiwan's Agency in Managed Interdependence

The evidence does not indicate a complete departure from the past; rather, it shows an extension of a relationship that already existed. Huawei demonstrated how U.S. control of upstream technology could affect TSMC's access to a Chinese customer (U.S. Department of Commerce 2020; Bown 2020). Although Arizona involved a higher level of cooperation, it also brought up the question of how much advanced production should be moved off the island (U.S. Department of Commerce 2024; Reinsch and Whitney 2025). The amendment to Taiwan's National Security Act in 2022 established a narrower boundary around critical technology (Taiwan Ministry of Justice 2022). Trade figures indicate that ties with China are still substantial even though their proportion has decreased, with mainland China and Hong Kong together accounting for 31.7 percent of exports in 2024 (Taiwan Ministry of Finance 2025). The new economic and scientific dialogues have given these developments a more lasting character (U.S. Department of State 2022; U.S. Department of State 2023). In all the cases, the same kind of change is evident: technology policy is increasingly functioning as foreign policy since decisions regarding factories, engineering knowledge, and supply chains now influence how Taiwan handles its relations with both the United States and China.

Taiwan does have a degree of autonomy in this situation, but this is not without limits. The authorities in Taipei cannot just cut themselves off from U.S.-connected semiconductor

technology, as the Huawei case has shown (Bown 2020; U.S. Department of Commerce 2020). Moreover, they cannot disregard the Chinese market, since it still made up a major portion of exports in 2024, nor can they ignore the broader security pressures associated with relations across the Straits (Taiwan Ministry of Finance 2025; Diamond, Ellis, and Schell 2023). Taipei operates within these constraints. The areas in which it does have some flexibility are evident in its policy actions: the government has enhanced legal safeguards for sensitive technology, kept on pursuing diversification, and established regular economic and scientific forums with the United States (Taiwan Ministry of Justice 2022; U.S. Department of State 2022; U.S. Department of State 2023). Lee and Knoerich (2023) point out a similar trend regarding Chinese investment, in which Taiwanese institutions are able to influence the impact of economic pressure coming from a more extensive state. The semiconductor cases also indicate that Taiwan can make use of its technological strengths to strengthen its relationships with partners it prefers.

Selective alignment is useful since the other main options do not match the evidence very well. Achieving full decoupling would mean separating from China to a greater extent than Taiwan has in fact done; in 2024 mainland China and Hong Kong still accounted for 31.7 per cent of Taiwan's exports (Taiwan Ministry of Finance 2025). A policy of merely following the United States would also fail to take into account the tensions arising from shifting advanced manufacturing overseas, since Taiwan's domestic cluster is still a major contributor to its strategic value (Reinsch and Whitney 2025). Maintaining an equal distance from both powers is difficult to reconcile with Taiwan's existing security arrangements with Washington and the fact that the major semiconductor companies depend on technology linked to the United States (Bown 2020; Diamond, Ellis, and Schell 2023). The evidence instead indicates that Taiwan is drawing boundaries. Coordination is most evident in the areas where technology and national security intersect, while greater room is left for ordinary economic exchanges. This is the reason why selective alignment captures both Taiwan's move towards Washington and the limits that Taiwan imposes on that movement.

If we consider the situation from Taiwan's point of view, this alters the standard narrative about semiconductors. In the context of the silicon-shield debate, Taiwan is usually seen as the location where valuable factories by chance happen to be. When it comes to supply chain talks, Taiwan can be viewed as representing a concentration risk which other countries have to alleviate (Chou 2023; Benson, Quitzon, and Reinsch 2023). While both of these descriptions reflect some truth, neither fully accounts for Taiwan's own decisions. Since 2018, those decisions have been to back major manufacturing investment in Arizona, to enhance legal protection for key technologies, and to reduce part of Taiwan's export dependence on China (U.S. Department of Commerce 2024; Taiwan Ministry of Justice 2022; Taiwan Ministry of Finance 2025). At the same time, the broader advanced-manufacturing ecosystem is still concentrated in Taiwan, a point which Reinsch and Whitney (2025) describe as central to the island's economic and strategic significance. Taiwan's chip industry is therefore not an automatic shield; its political value lies in how the government and companies handle the dependence associated with it.

8. Conclusion

Since 2018 the competition between the United States and China has caused a shift in the political significance of Taiwan's chip industry. All the fundamental elements of this situation had already been in place as TSMC had, for a long time, maintained close links with American technology firms, Taiwan had already got security associations with the United States, and the government had started making efforts to cut down its excessive reliance on China before the trade war (Miller 2022; Kennedy et al. 2018; Diamond, Ellis, and Schell 2023). The semiconductor conflict then brought these factors together more closely. When Washington began to regard advanced chips as a matter of national security, the choices made regarding customers, factory sites, trade secrets, and supply chains also came to have foreign-policy implications (Bown 2020; U.S. Department of Commerce 2020). Taiwan did not suddenly decide to side with the United States in 2018; rather, the existing relationship had become wider and more closely connected with economic security.

This pattern is evident in a number of different cases. The United States' restrictions on Huawei had an effect on one of TSMC's major customer relationships in China (Bown 2020). The Arizona project then linked TSMC more closely to U.S. industrial policy and at the same time raised doubts about how much advanced production should remain on the island (U.S. Department of Commerce 2024; Reinsch and Whitney 2025). Taiwan's legal changes in 2022 provided stronger national-security protection for critical technology (Taiwan Ministry of Justice 2022). Diversification away from China went on, but the 2024 trade figures indicate that mainland China and Hong Kong still made up 31.7 per cent of Taiwan's exports (Taiwan Ministry of Finance 2025). The regular U.S.-Taiwan economic and scientific dialogues helped to establish these developments in a more permanent way (U.S. Department of State 2022; U.S. Department of State 2023). No individual case demonstrates a new strategic pattern; the pattern is due to repetition, since after 2018 there was a shift towards a stronger security link in the areas of market access, investment, technology protection, trade diversification, and economic diplomacy.

The competition between the United States and China therefore strengthened Taiwan's strategic alignment with the United States and altered the areas in which this alignment takes place. The best way to describe the outcome is selective alignment: Taiwan has gotten nearer to Washington on the most sensitive aspects of the technology sector, but it has not completely severed ties with China and still maintains the domestic production capacity which is strategically important (Taiwan Ministry of Finance 2025; Reinsch and Whitney 2025). This Taiwan-focused approach contributes to the existing literature since the issue is viewed from the perspective of Taipei. Research into weaponized interdependence accounts for why the United States can make use of semiconductor chokepoints, and the silicon-shield debate explains why other states are concerned about Taiwan's factories (Farrell and Newman 2019; Chou 2023). Taiwan's challenge is how to stay closely connected to the network without losing control of the capabilities that provide it with value. The policies discussed here represent an attempt to strike a balance in managing that situation rather than trying to escape interdependence entirely.

Works Cited

Benson, Emily, Japhet Quitzon, and William Alan Reinsch. 2023. "Securing Semiconductor Supply Chains in the Indo-Pacific Economic Framework for Prosperity." Center for Strategic and International Studies.

Blackwill, Robert D., and Jennifer M. Harris. 2016. War by Other Means: Geoeconomics and Statecraft. Cambridge, MA: Belknap Press of Harvard University Press.

Bown, Chad P. 2020. "How the United States Marched the Semiconductor Industry into Its Trade War with China." East Asian Economic Review 24(4): 349–388.

Bown, Chad P., and Dan Wang. 2024. "Semiconductors and Modern Industrial Policy." Journal of Economic Perspectives 38(4): 81–110.

Chou, Kuanchu. 2023. "A Discussion on Geopolitical Risk of Taiwan Semiconductor Industry." APSA Preprints, August 18.

Diamond, Larry, James O. Ellis Jr., and Orville Schell, eds. 2023. Silicon Triangle: The United States, Taiwan, China, and Global Semiconductor Security. Hoover Institution and Asia Society Center on U.S.-China Relations.

Farrell, Henry, and Abraham L. Newman. 2019. "Weaponized Interdependence: How Global Economic Networks Shape State Coercion." International Security 44(1): 42–79.

Huang, Kwei-Bo. 2018. "Taiwan's New Southbound Policy: Background, Objectives, Framework and Limits." UNISCI Journal 46: 47–68.

Keller, William W., and Louis W. Pauly. 2017. "Crisis and Adaptation in East Asian Innovation Systems: The Case of the Semiconductor Industry in Taiwan and South Korea." Business and Politics.

Kennedy, Scott, Matthew P. Funaiole, Derek Mitchell, and Bonnie S. Glaser. 2018. The New Southbound Policy: Deepening Taiwan's Regional Integration. Washington, DC: Center for Strategic and International Studies.

Lee, Chun-Yi, and Jan Knoerich. 2023. "Buying Taiwan? The Limitations of Mainland Chinese Cross-Strait Direct Investments as a Tool of Economic Statecraft." The China Quarterly.

Miller, Chris. 2022. Chip War: The Fight for the World's Most Critical Technology. New York: Scribner.

Office of the United States Trade Representative. 2023. "United States and Taiwan Hold Second Negotiating Round for the U.S.-Taiwan Initiative on 21st Century Trade." August 18.

Office of the United States Trade Representative. 2024. "USTR announces Entry into Force of First Agreement under U.S.-Taiwan Initiative on 21st Century Trade." December 9.

Reinsch, William Alan, and Jack Whitney. 2025. "Silicon Island: Assessing Taiwan's Importance to U.S. Economic Growth and Security." Center for Strategic and International Studies.

Roussilhe, Gauthier, Thibault Pirson, Mathieu Chonneux, and David Bol. 2022. "From Silicon Shield to Carbon Lock-in? The Environmental Footprint of Electronic Components Manufacturing in Taiwan (2015–2020)." arXiv:2209.12523.

Sutter, Karen M. 2025. "Taiwan's Role in Global Semiconductor Supply Chains." Congressional Research Service, IG10073.

Sutter, Karen M. 2025. “U.S.-Taiwan Trade and Economic Relations.” Congressional Research Service, IF10256.
Taiwan Ministry of Finance. 2025. “In 2024, Taiwan’s Exports to Mainland China and Hong Kong Accounted for 31.7%, Down 12.2 Percentage Points from the Historical High in 2020.” Statistical Bulletin No. 1, January 16.
Taiwan Ministry of Justice. 2022. “National Security Act Amendment Passed Third Reading: Upgrading Trade Secret Protection and Protecting National Core Key Technologies.” May 20.
Thadani, Akhil, and Gregory C. Allen. 2023. “Mapping the Semiconductor Supply Chain: The Critical Role of the Indo-Pacific Region.” Center for Strategic and International Studies.
Tso, Chen-Dong, and Gratiana Jung. 2018. “Taiwan’s New Southbound Policy: One-Year Economic Performance Review.” National Bureau of Asian Research.
U.S. Department of Commerce. 2020. “Commerce Addresses Huawei’s Efforts to Undermine Entity List, Restricts Products Designed and Produced with U.S. Technologies.” May 15.
U.S. Department of Commerce. 2024. “Biden-Harris Administration Announces CHIPS Incentives Award with TSMC Arizona to Secure U.S. Leadership in Advanced Semiconductor Technology.” November 15.
U.S. Department of State. 2022. “The 2022 Third U.S.-Taiwan Economic Prosperity Partnership Dialogue.” December 15.
U.S. Department of State. 2023. “Inaugural U.S.-Taiwan Science and Technology Cooperation Dialogue Strengthens Collaboration.” May 22.
Yoon, Junghyun. 2023. “Supply Chain Security in the Age of Techno-Geopolitics: ‘Fab 4’ Case in the Semiconductor Industry.” Korean Journal of International Studies 21(1): 27–60.

The Deviation of the Triplets: a Comparative Study on Estonia, Kazakhstan, Uzbekistan's STEM Education By Wo Xuan Li

Abstract

The Programme for International Student Assessment (PISA) is an OECD test that measures how well fifteen-year-olds apply mathematics, reading, and science to unfamiliar problems. It has reflected that there are major differences in STEM outcomes between countries that gained independence from the Soviet Union. This study aims to find the factors that influence this divergence. It compares the STEM curriculum and examination demand of Estonia, Kazakhstan, and Uzbekistan - countries with Soviet backgrounds but different STEM education outcomes. By analyzing and comparing science curriculum objectives as well as analyzing each country's STEM related exam contents, this study found that STEM curriculum objectives and examination demand have direct influence over the three country's STEM education outcome. It found that the countries with curriculum and exam structure that are deemed as more challenging have higher scores on the PISA examination.

Introduction

Former Soviet states that became independent from the Soviet Union in 1991 inherited a similar education system. Their education system was centralized, heavy on content, and focused on memorization. Emphasis was placed on mathematics and science subjects. Thirty years later, however, the education systems in these countries no longer look alike. In particular, three post-Soviet countries took a distinctively different path in development of their education systems. Estonia devolved a large amount of autonomy to schools and teachers and invested in digital learning through programs like Tiger Leap. Meanwhile, Kazakhstan took on a more centralized approach, running visible reform projects such as the Nazarbayev Intellectual Schools and the trilingual policy while keeping a prescriptive national curriculum. On the other hand, the Uzbekistan government began serious education reform only in 2017.

The divergent outcomes of these educational systems manifest starkly in the performance of math and science subjects. In 2022, fifteen year olds Estonian students scored 510 in the mathematics in Programme for International Student Assessment (PISA) against an OECD average of 472. Additionally, 13% of students in Estonia were top performers in mathematics, meaning that they attained Level 5 or 6 in the PISA mathematics test. This was the highest mathematics score among all European countries. Kazakhstan students scored 425, with only 2% classified at Level 5 or 6 in the PISA mathematics test. Uzbekistan students took the test for the first time in 2022 and scored 364, virtually no one reaching the top level of 5 or 6 in PISA mathematics. The gap of 146 points between Estonia and Uzbekistan is larger than the gap between most of the countries that participated in the PISA examination, which is surprising, considering the fact that they all emerged from the same starting system.

Existing research does not explain this divergence well. Most of them examine one country at a time, which provides useful detail but does not allow horizontal comparison.

Another part of it treats the whole post-Soviet region as a single bloc going through transition, which neglects precisely the variation that needs explaining. A third group proposes general causes of educational inequality, such as government spending, national wealth, marketization, and the move toward meta subject skills, but these explanations are rarely tested against each other in countries that share common culture and history. Research on STEM education outcomes specifically, rather than education quality in general, is still underdeveloped.

This paper compares STEM education in Estonia, Kazakhstan, and Uzbekistan from independence to the present, analyzing the history of each country's education reformation. PISA mathematics score is the main benchmark. The comparison of science curriculum objectives and exam difficulty will be the evidence to evaluate the importance of these factors for the development of STEM education. The answer will be useful to any post-Soviet country trying to rebuild technical education rather than simply increase government spending.

Literature Review

Research on post-Soviet STEM education shares a common premise. Studies on Estonia, Kazakhstan, and Uzbekistan commonly argue that independence from the Soviet Union in 1991 left each country with the same institutional heritage. These countries first started with a centralized and content heavy education system organized around memorization and high stakes examination. In STEM, this allows students to build a solid foundation, readily retrieving information needed to pursue more advanced STEM education. However, it often neglects student's scientific skill development by skipping the process of understanding the principle behind the concept, acting as a double edged sword for countries that adopt this system. Scholars studying post-Soviet Union countries' STEM education also agree that the biggest limitation is difference in education quality rather than access to them. The literature converges on where the failure lies, pointing consistently to government spending, marketization, inequality in assessment structure. Existing studies note the same symptoms but identify different causes.
The first strand of literature establishes the correlation between government spending and STEM education success. Kangalakova and Rakhmetova focus on different factors affecting intellectual potential of students in Kazakhstan, and find that government investment and city development directly impact a region's STEM education development (Kangalakova, 2021). By comparing indicators across two time points ranking all 17 regions of Kazakhstan on each metric, they identified strong academic regions including Almaty City, Astana, and Karaganda, weak academic regions such as the Turkestan regions, and top research universities locations were one of the main drivers for the ranking. The paper found that cities with more investment on education and science facilities had higher development for the student's intellectual potential, with innovation and culture as lesser important sub-factor. In the meantime, Tompieva et al. (Tompieva, 2019) contend that Kazakhstan is currently catching up to world's leading education systems, using several examples of government investment projects on STEM education (Tompieva, 2019). They conclude that Kazakhstan's education development depends heavily on how much the government invests in education. Meanwhile, Shaturaev analyzes education in

Uzbekistan and frames Uzbekistan's low education development as a financing problem, attributing low PISA results to inadequate expenditure and teacher salaries, recommending further government investment. The paper suggests that spending 23.4 percent of the national budget on primary education is not sufficient. (Shaturaev, 2022) Although Shaturaev acknowledges the importance of preschool attendance, home literacy environment, and family support, he still emphasizes more funding. The author is contradicting himself, because if the factors listed above are not improved, just increasing the government spending will not improve the education in Uzbekistan.

The second strand examines a single national system in isolation. Bokayev (2016) surveys the quality of primary and secondary education in Kazakhstan and identifies weaknesses in curriculum delivery, teacher preparation, and administrative capacity. Toybazarova and Nazarova (2018) trace the trajectory of Kazakhstani modernization, cataloguing the recent trends alongside the problems that have accompanied them, including the establishment of NIS, a government initiative targeted to raise the new generation of Kazakhstani scientists (Farxodjonova, 2020). perform a similar exercise for Uzbekistan, describing the innovative processes introduced into the educational system and lauding its educational achievements (Farxodjonova, 2020). Lapping evaluates the democracy system of Estonia and discusses how its growth resulted in the development of its education system. (Lapping, 2004) The account of these papers is detailed and credible, and together they give the reader an accurate picture of each country's STEM education growth. A limitation however, is that readers retain focus on one country, without comparing one system against another, and consequently cannot distinguish whether problematic aspects are specifically pertaining to Kazakhstani or Uzbekistan or they are shared by every system that inherited the Soviet model.

The third strand of literature analyzes education reforms that occurred within the three countries and states the factors that influenced their success or failure, but without standardized metrics their conclusions about the current state of STEM education in Kazakhstan, Uzbekistan, and Estonia remain unreliable. Kozhabekov and colleagues state that as liberalization has taken hold in Kazakhstan's education sector, rapid private growth has challenged quality assurance and equitable access. (Kozhabekov, 2026) Though backed by a large amount of qualitative evidence, the problem is that the conclusion migrates from discourse to educational quality without any instrument capable of assessing quality. Yesnazar and Zhorabekova, meanwhile, discuss a concept called meta-subject skill, which is a student's cognitive capacity to acquire knowledge independently through a set of learning abilities that maintains their discipline while increasing their efficiency. Their research was conducted at the primary school level, examining what impact STEM-based instruction has on primary school students and their meta-subject skills. They concluded that STEM education in Kazakhstan needs greater focus on soft skills (Yesnazar, 2025). That conclusion is weak, however, because they used the Draw-A-Scientist Test to evaluate students' scientific ability rather than an established assessment such as PISA, which measures how students apply existing knowledge to problems they have not encountered before.

As analyzed above, the weaknesses in existing research on STEM education in Kazakhstan, Uzbekistan, and Estonia fall into three categories: overestimating government spending as an explanation, the analysis of a single country without comparing it to other systems. Moreover, they also analyze the three countries' education reform, and estimate their attribution without using universally verified metrics. There lies a common-overarching problem, which is the failure to consider the full range of factors that shape a country's education system. Rather than standing in isolation, all three systems are interconnected, sharing cultural background, the movement toward liberalization in education, and the emphasis placed on STEM. To address the weaknesses of the existing literature, this paper undertakes a comparative study of all three countries, examining their education reforms, how much those reforms resemble one another, and whether they achieved their intended outcomes. PISA results published by the OECD will be used as the yardstick in comparison of three countries, a test assessed by OECD that compares 80 countries in students' ability in creative thinking, financial literacy, and collaborative problem solving.

Theoretical Framework

This paper compares three countries from the Soviet successor states in their STEM education outcome. Two factors can influence the divergent outcome: target objectives in the science curricula and the examination system.

Having knowledge in science is not the same as having scientific skills. A student can have a great deal of knowledge in physics and yet be unable to design a scientific experiment. An education system that emphasizes knowledge will produce students that are good test takers rather than scientists.

The revised Bloom's taxonomy highlights this point. Bloom's taxonomy is a hierarchical framework used to categorize education goals into different levels, with lowest being least cognitive demanding and highest being the most. It arranges what any task requires into six levels, beginning with remembering and understanding, moving through applying and analysing, and ending at evaluating and creating. Knowledge sits at the bottom of this arrangement. Scientific reasoning sits near the top. The taxonomy lets us ask a precise question for any education system which is what level of ability it requires its students to wield.

There are two factors that are analyzed in this paper to see their impact on the STEM education outcome. The first is the science curriculum objective, which are measurable targets that define what students should be able to do and understand by the end of the course. The second is the entrance examination, which tells students what will be rewarded. Both should always complement each other in an education system.

First, the curriculum objectives set the standard. If it requires students to identify, list and describe, then teacher instruction will not go further above the demand. If its objectives demand students to analyse and evaluate, teachers have a mandate to teach reasoning and a defence when asked why they are not simply covering content.

This is why the second demand matters more than the first. Where an examination determines university admission and state funding, it becomes the practical definition of what education is for. Students study what will be tested. Teachers' lessons will be tested. Schools are judged on what will be tested. If that examination consists of multiple-choice items drawn from a known set of questions, then the rational response can be derived mostly from memorization. If the examination requires open ended investigative questions, then it requires the student to know their concept and actively use them.

According to Bloom's taxonomy, the science curriculum difficulty must match with its examination difficulty. Strong STEM outcomes come from systems where both demands sit high on the taxonomy. The curriculum tells teachers where to aim and what students should learn, on the other hand, the examination makes aiming there worth the effort and accurately measures them.

Result

Kazakhstan

Kazakhstan became part of the Russian Soviet Federal Socialist Republic in the 1920s. The Soviet period's development of education in the Kazakhstan region had strengthened tradition for Kazakhstani people's value of education. Like most of the Soviet Schools, it emphasizes state values and remains free of charge for all citizens. In the late 1980s, with the increasing demand for engineers and scientists, there was a shift towards more emphasis on STEM subjects.(Mynbayeva, 2014).

After the collapse of the Soviet Union, the government faced a series of challenges, including economic crisis and political instability, which caused state funding for education to decrease significantly. The amount of funding for education dropped from 5.7% of Gross Domestic Product (GDP) in 1990 to 3.5% in 1991–1992, and was only stabilized at 4.3% in 1996, five years after regaining independence (Mynbayeva, 2014). Due to the massive drop-out rate and the number of students who refused to attend school, the Kazakhstani government introduced a law in 1999 that made school attendance compulsory. This greatly contributed to the development of the education curriculum, thanks to a stable student population on which it could grow. Later, the new standards introduced the competence approach, with a curriculum that focused on students' ability to use the knowledge they had learned and on the concept of 'education oriented to the outcome' (Mynbayeva, 2014). This meant that the main purpose of education was oriented toward gaining knowledge, rather than toward passing exams. At the same time, specialized high schools were introduced that focused on natural sciences and other STEM subjects (Mynbayeva, 2014).

Beginning from the 2010s, the government drastically modernized its education system, focusing on improvement of teaching methodology, school infrastructure, and accessibility of education (Karabota, 2025). Furthermore, it introduced more branches of STEM education to grow the new generation of high-achieving STEM students for Kazakhstan, such as Nazarbayev

Intellectual Schools(NIS). First established in 2009, the schools with their tri-lingual curriculum (English, Russian, Kazakh) focus on developing students' skills in science, technology, engineering, and mathematics. They are recognized for their highly competitive admission and integration with America's advanced placement system, which allows students to learn college courses in high school when deemed ready. Over 20 NIS were established across Kazakhstan, building a vast alumni network of scientists who pursue higher education.

The examination system also underwent several major shifts in the post Soviet Union Era. In 2003–2004, Kazakhstan replaced traditional final examinations with the Unified National Test- a centralized test that was designed by the ministry of education aiming to standardize the examination for college entrance.. The original design combined both functions in one instrument: results counted as final assessment grades and as university entrance grades. Three subjects were compulsory, mathematics, Kazakh or Russian language, and history of Kazakhstan. One subject was optional. There were closed-ended questions with a single correct answer, three hours, scored on a 120-point scale, administered through 134 testing centres.
In 2007, a new policy shifted the school model from education for life to education throughout life and provided a new model of school education based on competencies. In 2017, the UNT was split into two procedures, a separate school-leaving assessment and a university entrance test, though the UNT remained the gateway.

Currently, the UNT now runs to 120 questions across five subjects, with a baseline threshold of 50 points rising to 75 for law and pedagogy, and minimum subject-level scores including 3 points each in reading literacy and mathematical literacy. Two attempts are permitted. Alongside it, the External Assessment of Academic Achievement is administered in grades 4, 9 and 11, and the grade 11 Final Attestation is administered at school level without Quality Control Committee oversight(Karabota, 2025).

Uzbekistan

Similar to Kazakhstan, Uzbekistan gained its independence from the Soviet Union in 1991, and had heavy influence from the Soviet Union's centralized education style and strong nationalism . It had made the attendance of school compulsory earlier than Kazakhstan in 1992, establishing a 11 year compulsory public education system.

Later in 1997, The Law on Education and the National Programme for Personnel Training were adopted, establishing compulsory secondary education including specialised and vocational education, and free choice between academic lyceum and vocational college. The eleven-year system was then replaced by nine years of compulsory secondary education followed by two compulsory three-year pathways: vocational colleges or academic lyceums. Academic lyceums were defined as schools that provided intensive development of intellectual abilities and professionally oriented training; professional colleges as providing in-depth development of professional skills in one or more specialties (Eshchanov, 2020). Later during this period, more than 2000 vocational colleges and hundreds of academic lyceums were built.

Uzbekistan made vocational or academic secondary education compulsory from 2000. Every student who finished grade 9 was routed into either a professional college or an academic lyceum. 2,500 new colleges were built between 2000 and 2015 to facilitate the students. The intended split was 90 percent into professional colleges and 10 percent into academic lyceums, which meant that the institutions built for intensive study of mathematics and science were designed for a tenth of the students. In 2017, the government reversed course, closed most vocational colleges, and restored eleven years of general schooling after pressure from parents and the public. The swing then went too far the other way. In 2025 only 182,000 of 608,000 grade 9 graduates entered technical colleges against a national capacity of 450,000.The president told regional governors to push vocational enrolment back to at least half.

Uzbekistan's College Entrance Exam also underwent many changes, from Soviet period to 1993, admission to higher education institutions was by oral or written entrance examination in three relevant subjects, administered locally by each institution. Competition in Uzbekistan was consistently tighter than the Soviet average, rising from 263 applicants per 100 places in 1980 to 342 in 1989, against a USSR figure of 228 (Eshchanov, 2020).

In 1994**.** The State Test Centre took over entrance examinations for most institutions, excluding foreign branches, academies, and performance-based programmes in arts and sport. The format was three subjects, thirty-six multiple-choice questions each, with the subjects weighted 3.1, 2.1 and 1.1 by relevance to the chosen degree, giving a maximum of 226.8 points.

By 2020, the instrument was essentially unchanged: three sections of thirty-six multiple-choice questions, two on discipline subjects and one on native language, a maximum 227 points, with most admissions falling between 65 and 90 points. Male applicants over 18 could add 60 points through an army service merit certificate. In 2026, the applicant can "test now, choose after", meaning that they don't need to decide for a designated major before finishing the examination.

Estonia

Estonia was annexed by the Soviet Union in 1940 and regained independence in 1991. The curriculum was centralized, and compared to the rest of the countries, the Soviet curriculum is considered stronger and had an emphasis on STEM education. Instruction was delivered mainly in Russian, with over half of schooling conducted in Russian as late as 1989 in which it proceeded to slowly transition into Estonian.

In the years immediately after independence, the government got rid of the centralized Soviet curriculum. Without a unified system, the schools were left to develop their curriculum independently, and most continued with the teaching inherited from the Soviet period while adding enquiry-based and active learning methods. The Education Act of 1992 legalized private schools and set out the principles of comprehensive basic education, and the Basic Schools and Upper Secondary Schools Act of 1993 established Estonians as the primary medium of instruction. The curriculum reform itself ran from 1987 to 1996 and involved over a thousand teaching faculties, making it a bottom-up rather than a ministerial one (Ruus and Sarv, 2020).

The new national curriculum was adopted in September 1996. It focused on school outputs rather than content coverage, describing competencies to be achieved at the end of each school stage and providing guidance on organising student-centred learning (Lees, 2016). The most significant change was that each school was now required to develop its own curriculum, which was a large shift for teachers accustomed to a centrally prescribed programme. Some critics argued the change happened only on paper at first and that real classroom practice did not follow until much later (Ruus and Sarv, 2020). In 2011, the national curriculum was separated into two frameworks with each framework allowing schools to develop their own curricula.

Estonia invested early in digital learning. Between 1992 and 1996 the state put about 0.2 million USD annually into school IT equipment, which was enough for roughly 40 schools to equip proper computer classrooms. Following this, the Tiger Leap programme was launched in 1996 by President Lennart Meri, Education Minister Jaak Aaviksoo, and Toomas Hendrik Ilves, then ambassador to the United States. It stood on three pillars: computers and internet access, basic teacher training, and Estonian-language electronic courseware. The Tiger Leap Foundation was established in 1997 by the Ministry of Education together with computer companies and private individuals. By 2000 all Estonian schools had computers, and by 2001 all had internet access. A follow-on programme, Tiger Leap Plus One, was launched in 2001 (Lees, 2016).

Alongside these reforms, Estonia raised the entry requirements for teaching. All teachers are now required to hold a master's degree, and teachers have autonomy in choosing study materials and in evaluation, while schools have autonomy over personnel and budget (Lees, 2016).

Estonia's examination system starts at the end of grade 9, where students take three final examinations in Estonian, mathematics, and one elective subject. The papers are administered centrally and are identical across all schools.

To graduate from upper secondary school, a student must pass three state examinations in mother tongue, mathematics and a foreign language, complete a school examination, and produce a research paper or practical work assessed at least satisfactory. The state examinations are made out of open-ended questions, prepared by working groups of subject specialists, and graded anonymously by state commissions on coded papers using a 100-point scale (Eurydice, 2026).

Since 2014, the exam is considered passed as long as the student scores at least one point out of a hundred. Therefore, students almost never fail the state examination. What can end a student's progression is coursework: a student is expelled if they receive "poor" or "weak" in more than half the courses in three or more subjects in any academic year. While the state examination produces a comparable score that universities use for admission, teachers' grading plays an important role in a student's graduation. Universities may additionally require their own tests, interviews or portfolios, with some requiring students to take a second round of exam for entrance to university(Eurydice, 2026).

Comparison of Curricula and Examination System

Science Curriculum Comparison

The first variable is the science curriculum. Each country states what its science programme is meant to produce. An objective written as an action can be sorted by the cognitive process it demands, running from remembering and understanding through applying and analysing to evaluating and creating. One country's curriculum tells us little on its own, but three curricula ranked against each other tell us whether the systems are asking their students for different things. Assuming Estonia's objectives at the analysing and evaluating levels, Kazakhstan's at applying, and Uzbekistan's at remembering, then the ranking follows the same order as the three PISA mathematics scores. The prediction is straightforward. Overall, this analysis result supports that the higher the cognitive level a science curriculum demands, the more advanced the reasoning the student should be able to demonstrate on an assessment built to test reasoning. PISA is an assessment that identifies well-rounded students, and regards students trained in analyzing to be better at STEM subjects than students who are only proficient in retrieving knowledge.

In order to compare the first variable, see table one's comparison regarding the three countries' science curriculum based on the following variables established by an existing research (Yazıcıoğlu, 2018): gain basic knowledge, environmental consciousness, values and identity, learning areas specified, solving daily problems, and who writes the curriculum.

Objective	Estonia	Kazakhstan	Uzbekistan
Gain basic knowledge	Field-related competencies shaped by subject objectives and learning outcomes; syllabi set outcomes for knowledge, skills and attitudes by end of each school year	State emphasizes the continuity in the development of knowledge through different stages, skills and abilities acquired in primary school	Content focused learning which 6 stem related subjects taught from grade 5-11
Environmental Consciousness	Delivered through cross-curricular topics rather than as a separate science objective	Education and research activities, rational nature management and environmental protection skills, necessary skills to ensure the safety of human and social life	none
Values and identity	Schools have "the responsibility to safeguard and develop the Estonian nation, language and culture"	"Development of an integrated information, value and attitude system corresponding to a multinational assembly"; patriotism consciousness stated explicitly	Handled outside the science curriculum, through separate policy on spiritual and ethical education of youths as well as patriotism values.

Learning areas specified	Formative assessment that divides learning areas into different stages and areas, described as assessment supporting learning and student development	Not included in the programme initially; expressed as "main results", "personal results" and "system activity results"	None
Solving daily problems	Competency defined as the ability "to operate creatively, entrepreneurially, flexibly and effectively in a particular area"	"To develop mental and creative skills to solve life problems, and formulate solution by speaking and writing"	Developing soft skills including independent and creative thinking, with ability to formulate ideas through writing
Who writes the curriculum	National curriculum is a framework; schools develop their own curricula from it	Centrally prescribed	Centrally prescribed

Table 1: Comparison of General Scientific Curriculum Objectives

Subject objectives: clear statements that describe what a student should know, understand, or be able to do by the end of a specific subject or lesson. According to table 1, all three curriculums have similar objectives, however, it is important to notice that only Estonia specifies learning areas, and Uzbekistan lacks environmental consciousness as well.
Gaining basic knowledge as a skill set is addressed by all three curriculums, with Estonian curriculum setting up specific field related mastery goals, Kazakhstan emphasizing on the continuity of developing this ability through different phases of school and Uzbekistan focusing on the continuity of knowledge of the crucial courses throughout the system.

Patriotism is an important theme for both Kazakhstan and Uzbekistan's curriculum, and it's centrally designed by the government. On the other hand, the Estonian curriculum focuses more on independent thinking by giving schools the right to make changes on the curriculum as well.

The second variable is the entrance examination, and it is analysed the same way by selecting 5 STEM related questions randomly from each country's college entrance exam. If Estonia's mathematics and science tested subjects sit higher on the taxonomy than Kazakhstan's, and Kazakhstan's higher than Uzbekistan's, then the same ordering appears twice and the two mechanisms reinforce each other. That result would show the curriculum and the exam asking for the same thing rather than working against each other. It would also close the gap between what a system says it wants and what it actually rewards. A country that demands reasoning in both places is teaching it twice, and its students should carry that ability into any test built to measure it, just like PISA which this paper is primarily based upon. (The questions analyzed below are obtained directly from each country's official test preparation website.)

Country	Item (abbreviated)	Question Format	Bloom level	Why it sits at this level
Estonia	Marbles drawn from two boxes, second draw dependent on the first	Open response	Analysing	The sample space changes according to an unknown prior draw, so the student must partition into cases and combine them rather than apply one formula
	Cubic function: find extrema, then determine how many real solutions an equation has and justify	Open response	Evaluating	Requires reasoning from the extrema to the number of intersections and defending the conclusion; the instruction to justify cannot be satisfied by a value
	Arithmetic sequence, then three logarithms as consecutive terms	Open response	Analysing	The student converts an additive condition on logarithms into a multiplicative one on their arguments, then tests domain validity
	Pyramid on an isosceles trapezoid, apex projected onto the shorter base	Open response	Analysing	The trapezoid must be reconstructed from partial data and the height derived across several dependent stages before any formula applies
	Pille and Kalle, 60 km ride with changing speeds	Open response	Analysing	Two time expressions are built from percentage relationships and equated; a physical constraint forces rejection of one root
Kazakhstan	Shapes puzzle: square plus triangle = 12, triangle plus circle = 9, find square minus circle	5-option multiple choice	Applying	One elimination step on a system disguised in unfamiliar notation
	Angles 3x, 5x and x on a straight line; find a ratio of two angles	5-option multiple choice	Applying	Solve 9x = 180, then form a ratio; two routine steps in sequence
	Hunter's bullets in five pockets, each between one and five and all different	5-option multiple choice	Analysing	The constraints admit exactly one assignment, so the student deduces rather than computes
	Two aeroplanes, one slowing after two hours	5-option multiple choice	Analysing	The pursuit begins from an existing lead and relative speed reverses, so the journey must be modelled in two phases

	Wage of 90,000 tenge plus a robe for 30 days, paid after 3	5-option multiple choice	Analysing	The robe's value is embedded in the total wage, so the proportional equation must be constructed
Uzbekistan	Number line: which inequality is NOT true	4-option multiple choice	Analysing	Sign and magnitude relations are extracted from a diagram, then four claims tested to find the one that fails
	Calculate (38 x 54 + 38 x 42) divided by 24	4-option multiple choice	Applying	Routine computation; the only judgment is whether to factor first
	Solve (x-7)/3 + (2x+1)/2 = 5/6	4-option multiple choice	Applying	Standard procedure for clearing denominators; no choice of method
	Two different positions among three candidates	4-option multiple choice	Applying	Recognise that order matters and count six arrangements
	Area of three planted plots forming a rectangle	4-option multiple choice	Applying	Read two dimensions from a figure and compute one area

Table 2: Comparison of cognitive demand between each country's STEM related question in the college entrance exam of Estonia, Kazakhstan, and Uzbekistan

Distribution

	Applying	**Analysing**	**Evaluating**
Estonia	0	4	1
Kazakhstan	2	3	0
Uzbekistan	4	1	0

Table 3: Distribution of items based on Bloom's taxonomy level

Bloom's taxonomy, introduced by Benjamin Bloom in 1956 and revised by Anderson and Krathwohl in 2001, is a hierarchical framework that classifies educational tasks by the type of thinking required. Its six levels run from remembering and understanding at the bottom of the cognitive demand pyramid, applying and analysing in the middle, and evaluating and creating at the peak of the pyramid. Because the levels are ordered by the complexity of the mental operation involved, the taxonomy allows curriculum objectives or examination questions to be placed on a common scale of cognitive demand. This makes it possible to compare the degree of

cognitive demand different education systems actually require from students in their examination questions.

As shown in table 2, response format varies by country. Every Estonian item requires a constructed answer. Every Kazakh and Uzbek item offers the answer among the options, with Kazakhstan offering 5 options and Uzbekistan offering 4 options.

As for the comparison of distribution of the items, the study shows that Estonia and Kazakhstan had the majority of the questions lying upon the "analyzing" category according to Bloom's taxonomy, demanding students to break down information given into small components, solving them using contrast, compare, etc.

Uzbekistan on the other hand had most of its questions falling on the "applying" category, demanding the students to solve and interpret the questions, which requires less cognitive demand than "analyzing".

Discussion

The quality of STEM education has an influence beyond just test scores. It determines how students approach unfamiliar problems and whether they can reason from evidence rather than simply recalling knowledge they have. A country's capacity to build sound infrastructure, solve public health crises, and develop competitive industries rests significantly on whether its education nurtures individuals to think critically and scientifically.

There are two factors that are consistently supported as the reason for the different STEM education quality: The objective of science curriculum and the quality of examination. These factors are drawn based on the comparison of countries' science curriculum and examination questions. The findings support the hypothesis that divergence of science education curriculum and demand of the examination has a direct correlation with students' STEM subject performances in PISA tests.

The factors are consistent with the current PISA outcome of the three countries and the existing literature of Estonia, Kazakhstan, and Uzbekistan's student STEM performance. These findings have strong implications on the importance of strengthening independent reasoning skills on STEM related subjects such as computer science, math, physics. Especially to the 15 post-Soviet countries.

There are several limitations of this paper. First, this study is based on qualitative analysis focused on the country's curriculum, historical reformation and exam demand. It does not take into account economic factors, and geopolitical relationships with its neighboring countries that could also impact its student population. Second, the number of sample exam questions analyzed in this study is small and used for only an illustrative purpose, so they might not completely represent each country's examinations. Another limitation of the study is that there are no data directly obtained from the students of each country's classroom, which would be useful to directly compare the country's teaching delivery and how much it actually aligns with the written curriculum.

The findings suggest that the ministry of education in post-Soviet countries can adopt more efficient ways of intervention than the measures they have previously implemented to improve their STEM education. Post-soviet countries can change curriculum objectives to require higher cognitive capacity and redesign the school entrance examinations to require answers that are more constructive and creative. This would cost much less than building more schools and raising teacher salaries. This would also benefit all students in the country rather than the chosen pool of students in elite institutions. For Kazakhstan, replacing the Unified National Test focused on three specific subjects with a modular instrument that assesses critical thinking and research skills as a whole will lead to a promising outcome. For Uzbekistan, where the weight of the mathematics subject in the entrance examination constitutes only eleven points out of the total of 189 points, raising that weighting would be a crucial state-wide structural change to elevate its STEM education quality.

Conclusion

This study has shown that the objective of STEM courses and exam rigor can have a direct impact on student's performance on STEM subjects. The rigor of course objectives and examinations has direct relations with PISA performance. Estonia's STEM subject curriculum and examination encourage students to think independently. It has helped the Estonian students to acquire skills needed to score higher in the PISA exam, in which Kazakhstani and Uzbekistani students struggled. In order for post-Soviet countries to improve the quality of STEM education, science curriculum objectives and STEM examination should be given more attention. As Jess Lair, children's book author said, 'Children are not things to be molded, but are people to be unfolded'. By focusing on developing creativity and reasoning skills over memorization, post-Soviet's education system will produce more innovators that can tackle many challenging problems with imaginative solutions.

Works Cited

Bokayev, Baurzhan. “The Improvement of Quality of Primary and Secondary Education in Kazakhstan.” The Improvement of Quality of Primary and Secondary Education in Kazakhstan, 1 Jan. 1970, ecommons.cornell.edu/items/546a06fd-49cb-43a5-8f42-e3a1214ec206.

Eshchanov, Ruzumboy. “The Current Core of Education Reforms in Uzbekistan: One Step Forward Two Steps Back?” European Journal of Research and Reflection in Educational Sciences , 2020, www.idpublications.org/wp-content/uploads/2020/08/Full-Paper-THE-CURRENT-CORE-OF-EDUCATION-REFORMS-IN-UZBEKISTAN-ONE-STEP-FORWARD-TWO-STEPS-BACK.pdf

Eurydice. "Estonia: Assessment in General Upper Secondary Education." European Commissions, 2025, eurydice.eacea.ec.europa.eu/.

Farxodjonqizi, Farxodjonova Nodira, and Numonjonov Shahzod Dilshodjonugli. “Innovative Processes and Trends in the Educational Process in Uzbekistan.” ACADEMICIA: An International Multidisciplinary Research Journal, South Asian Academic Research Journals, 18 May 2020, indianjournals.com/article/aca-10-4-087.

The Government of Uzbekistan. "Education Sector Plan (ESP) of Uzbekistan 2019-2023." Globalpartnership, 2019, www.globalpartnership.org/node/document/download?file=sites%2Fdefault%2Ffiles%2F2019-04-gpe-esp-uzbekistan.pdf.

Kangalakova, D. M, and D. A. Rakhmetova. "Structural Features of the Intellectual Potential of the Regions and Its Impact on the Development of the Country." Economy: Strategy and Practice, 14 Oct. 2021, doi.org/10.51176/1997-9967-2021-22-34.

Karabota, B.Sh. “History of the Development of the School Education System in the Republic Kazakhstan.” Researchgate, 2025, www.researchgate.net/publication/399451563_History_of_the_Development_of_the_School_Education_System_in_the_Republic_of_Kazakhstan.

Kozhabekov, S. "Neoliberal Reforms and Privatization of Higher Education in Kazakhstan: A Critical Discourse Analysis of Policy Documents." Social Sciences and Humanities Open, 2026, doi.org/10.1016/j.ssaho.2026.103052.

Lapping, Mark. "Education in a Restoration Democracy: The Case of Estonia." Sagepub, 2004, journals.sagepub.com/doi/pdf/10.2304/csee.2004.6.2.101.

Lees, Mihkel. "Estonian Education System 1990-2016 Reforms and Their Impact 2016." Estonian Education System 1990-2016, 2016, 4liberty.eu/wp-content/uploads/2016/08/Estonian-Education-System_1990-2016.pdf.

Mynbayeva, Aigerim, and Victoria Pogosian. "Kazakhstani School Education Development from the 1930s: History and Current Trends." Italian Journal of Sociology of Education, Padova University Press, 1 Jan. 2014, ijse.padovauniversitypress.it/2014/2/7.

Nurbekova, Zhanat. "Digitalization of the Education System in Kazakhstan: Experience, Problems and Perspectives." Central and Eastern European Online Library, Национално издателство за образование и наука „Аз-буки", 1 Jan. 2023, www.ceeol.com/search/article-detail?id=1168927.

OECD. "How Did Countries Perform in Pisa?: Pisa 2022 Results (Volume I) | OECD." PISA 2022 Results, 2023, www.oecd.org/en/publications/pisa-2022-results-volume-i_53f23881-en/full-report/how-did-countries-perform-in-pisa_dc514907.html.

Sarv, Ene. "Estonian Curriculum: Becoming Independent." Academia, 2020, www.academia.edu/43591682/Estonian_Curriculum_Becoming_Independent.

Shaturaev, Jakhongir. "Insights from the Public Education System of Uzbekistan." Academic Research in Educational Sciences, 2022, ares.uz/storage/app/uploads/public/631/61b/459/63161b45989e8375494066.pdf.

Teataja, Rigii. "National Curriculum for Basic Schools." Riigi Teataja, 2011, www.riigiteataja.ee/en/eli/529042024002/consolide.

Tompieva, Zaure. "Development of STEM-Education in the World and Kazakhstan." CyberLeninka, 2019, cyberleninka.ru/article/n/development-of-stem-education-in-the-world-and-kazakhstan.

Yazıcıoğlu, Özlem. "A Comparison of the Middle School Science Programmes in Turkey, Singapore and Kazakhstan." International Electronic Journal of Elementary Education, 2018, doi.org/10.26822/iejee.2019248588.

Yesnazar, Asel, et al. "STEM Education Impact on the Development of Primary School Students' Meta-Subject Skills: An Experimental Study in Kazakhstan." Frontiers, Frontiers, 6 Sept. 2026, doi.org/10.3389/feduc.2025.1669765.

Beyond the Music: The Engineering Secrets Behind Concert Hall Acoustics
By Jaydzia E. Khan

Abstract

Sound in a concert hall depends not only on the performers but also on the design of the space itself. This research studies how engineers design concert halls to optimize acoustics. Studies in acoustical engineering consider factors such as reverberation time, which is how long sound takes to fade after the source stops. This review finds that successful concert halls depend on a combination of multiple engineering elements rather than a single element. Understanding how engineering affects concert hall acoustics can help engineers make decisions about future performance environments. This research connects engineering with the arts, showing how interdisciplinary solutions can contribute to real-world challenges.

Keywords

Engineering, Acoustical Engineering, Concert Hall Acoustics, Reverberation Time, Acoustic Design

Introduction

A concert hall is more than a place where music is performed. It is an environment where engineering shapes how audiences experience sound. The challenge of controlling sound in a performance is not new. Ancient Greek and Roman theaters represent some of the earliest known examples of indoor places to perform built with acoustics in mind. The history of concert hall acoustics reflects a shift from trial and error toward scientific prediction. The Odeon of Pericles, built in the fifth century BCE, is an early example of an indoor performance building, showing how architecture shaped the audience's experience of sound before engineers could scientifically predict acoustic behavior.[1]

Acoustics uses physics and engineering to understand factors such as reverberation time and sound quality. Researchers have studied how factors in the environment affect the way sound reaches performers and audiences.[2] Concert hall planning allows engineers to predict how different choices will affect sound before construction.[3] About a century ago, building concert halls was becoming increasingly based on science rather than observation and trial and error. Engineers began using mathematical methods to predict acoustic conditions before construction, allowing them to choose architectural designs before a hall was built.[4] Concert hall acoustics has therefore become an interdisciplinary field combining engineering and music. Engineers create environments that support different types of performances.

This literature review traces how concert hall acoustics evolved from observation and trial and error to scientific prediction, examining the principles of sound, the role of room shape and acoustic materials, and the engineering choices seen in famous concert halls. It finds that successful concert hall acoustics depend on multiple engineering elements working together rather than on a single factor. No single element determines acoustic quality, but the combination

of these factors allows engineers to create performance environments with clear, balanced sound for performers and audiences. However, the combination of these engineering choices can create trade-offs, so engineers must balance multiple acoustic factors rather than optimize each one independently.

Understanding how engineers combine multiple acoustic factors can help improve the acoustic quality of future concert halls for performers and audiences. This paper draws on scholarly articles in acoustical engineering and related fields to explain how engineers optimize concert hall acoustics.

The paper is organised as follows: First reviewed are the basic acoustic factors that determine behavior in concert halls, including reverberation, a very important factor.[2] The next section discusses how engineers control sound quality.[3] The final section compares well-known concert hall case studies and the engineering choices used to optimize acoustics.[2]

Discussion

Sound in a Concert Hall

This section explains how acoustic factors affect sound quality in concert halls and how engineers use these factors to create acoustic conditions.

When sound waves reach walls, ceilings, floors, people, and other materials, they can reflect off surfaces or lose energy through absorption. These interactions affect sound quality and how clearly, evenly, and naturally an audience hears a performance.[3,5,6]

Reverberation refers to the continuation of sound after the original sound source stops, and occurs when sound waves repeatedly reflect from surfaces throughout a room before gradually losing energy through absorption and other processes.[5,6] One study shows that increasing absorption reduced reverberation time and increased musical clarity, without affecting sound strength much.[7] This suggests that absorption can improve clarity by lessening excessive reverberation without necessarily lessening the overall strength of the sound.

The amount of reverberation can strongly affect the way music is heard. Too much reverberation can reduce clarity, while too little can make music sound dry or lack depth.[3,6] Some research suggests that large symphony halls commonly have reverberation times around 1.8 to 2.2 seconds, although the appropriate value varies according to the size and purpose of the hall.[2] This range shows that engineers can't depend on one ideal reverberation time for every concert hall.

Other factors, such as sound level, sound clarity, and amplification, also affect how audiences experience sound. In traditional concert halls, engineers rely heavily on room shape, reflective surfaces, and reverberation to help sound reach the audience naturally. In performances that use microphones and loudspeakers, amplification can supplement the sound produced by performers and help distribute it throughout the hall. However, amplification does not replace the need for appropriate room acoustics, since reflections and reverberation continue to affect the sound listeners hear. Therefore, engineers must consider both the acoustic properties of the hall

and the use of electronic sound systems when planning for different types of performances.[2]

Sound reflection occurs when sound waves reach a surface and bounce back into the room. Hard, flat surfaces such as walls and ceilings can produce strong reflection, while the direction of a reflected wave depends on the angle at which it reaches the surfaces.[5,6]

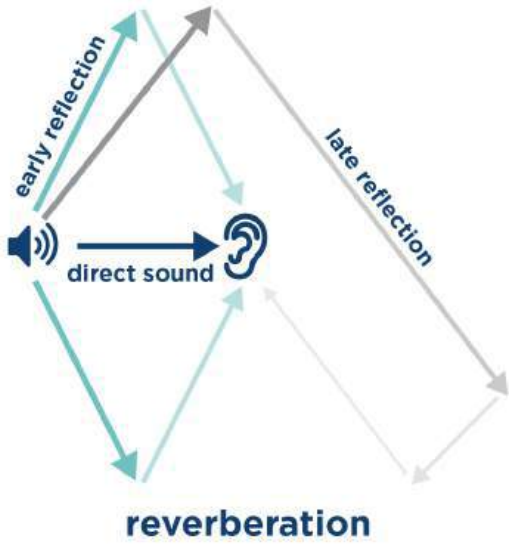

Figure 1. A diagram showing pathways of direct sound, early reflections, and late reflections. Image taken from Audiotek.[5]

Reflections can support good acoustics when they reach listeners at appropriate times and levels. As shown in Figure 1, the sound heard by a listener combines three components that arrive at their ears at different times.[5] Listeners usually receive direct sound first, followed by reflected sound that has traveled longer paths through the room.[5] Early reflections can add sound energy and help distribute music throughout a hall, while thousands of late reflections arrive so closely spaced that they blend. The difference between these reflections shows why both the timing and the strength of reflections matter, rather than simply the amount of reflected sound in a hall. However, reflections that arrive too late or are too strong can be heard as distinct echoes and reduce clarity.[5,6] The shape and location of reflective surfaces therefore affect how sound reaches different parts of the audience. Large surfaces, ceilings, balconies, and side walls can all change paths that sound waves take through a concert hall.[5]

Reflection also relates closely to reverberation because repeated reflections allow sound to persist in the room after the performer stops playing. A single reflection may be heard as an echo, while many closely spaced reflections can combine into reverberation.[6] This relationship shows why engineers must consider how sound moves throughout the entire hall rather than treating each acoustic factor separately.

Sound absorption occurs when materials take in some of the energy in sound waves instead of reflecting all of it into the room. Soft materials such as curtains, fabric, and acoustic panels can absorb sound, with different materials absorbing different amounts of sound.[3,5,6] Engineers can use these materials to reduce excessive reflections and control reverberation.

Recent research has also studied advanced sound-absorbing materials, including microperforated panels and acoustic metamaterials, which can be designed to absorb specific frequencies while providing high absorption within relatively small designs.[8]

The amount of sound absorbed by a material can be measured using an absorption coefficient, which shows how much incoming sound energy the material absorbs. It ranges from 0 (perfectly reflective) to 1 (perfectly absorptive). Materials with higher coefficients absorb more sound energy from a room.[5] Therefore, engineers can use absorption coefficients to compare materials and determine which surfaces are right for controlling reverberation. As shown in Table 1, absorption coefficients vary between materials and across different frequencies.[8] Acoustic design can therefore affect both the overall reverberation time and the clarity of a performance.

Reverberation also changes according to whether a hall is occupied. Audiences absorb sound through clothing, bodies, and seating, causing occupied halls to have shorter reverberation times than empty halls.[2] As shown in Table 1, fully occupied fabric-upholstered seats have an absorption coefficient of 0.96 at 1000 Hz, showing the absorption of occupied seating is relatively high. Brick has a significantly lower absorption coefficient of 0.04 at 1000 Hz, and this difference shows that seating and reflective materials can have very different effects on sound absorption. Reverberation therefore provides an important starting point for understanding how sound reflections and materials affect the listener.

Table 1. Absorption coefficients (α) of different materials across frequencies measured in hertz (Hz). Table adapted from Beranek 2006.[8]

Material	125 Hz	250 Hz	500 Hz	1000 Hz	2000 Hz	4000 Hz
Brick	0.03	0.03	0.03	0.04	0.05	0.07
Heavy wood paneling (2” thick over framing)	0.04	0.04	0.06	0.06	0.1	0.1
Seated audience (fabric-upholstered, per square meter of area)	0.6	0.74	0.88	0.96	0.93	0.85

Absorption must also be considered alongside reflection. A hall with too much absorption may have insufficient reverberation, while a hall with too many reflective surfaces may produce excessive reverberation or unwanted echoes.[3,6] Engineers therefore select and position materials according to the acoustic needs of the hall and the type of music performed there. The relationship between absorption and reflection leads to the broader question of how physical characteristics such as room shape and materials can be combined to produce the desired acoustic conditions.

Concert hall acoustic factors are closely related. No single factor can create ideal acoustics on its own. Instead, engineers must consider how these factors work together. These

findings raise a further question of how engineers use room shape and acoustic materials to control these acoustic factors in actual concert halls.

Use of Room Shape and Acoustic Materials

This section discusses how engineers use room shape, layout, and acoustic materials to control sound paths and create appropriate acoustic conditions in concert halls.

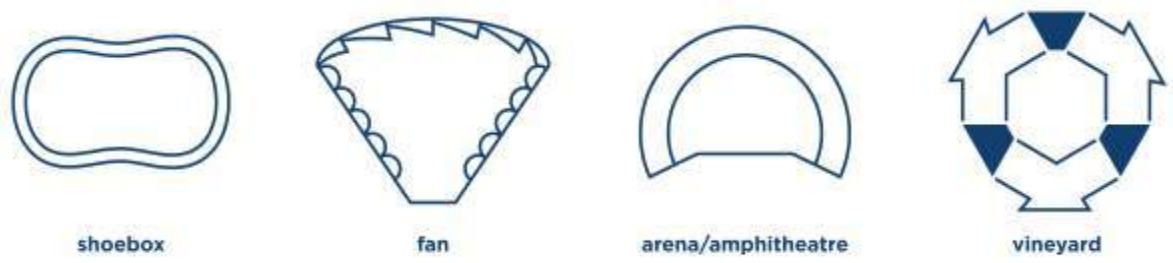

Figure 2. Diagrams showing common concert hall shapes. Image adapted from Audiotek.[5]

The physical characteristics of a concert hall affect how sound is heard by both performers and audiences. Room shape, stage placement, seating arrangement, and materials influence how sound waves travel throughout a hall.[5,6] As shown in Figure 2, research describes several common room shapes, each with different acoustic properties that affect how sound travels throughout a hall. Table 2 compares these hall types and their variables. These details are particularly important because musicians need to hear one another clearly, while audiences need consistent sound throughout the hall.[5]

Room shape affects the paths that sound waves take through a concert hall.[5] Some common shapes are shown in Figure 2. For example, comparative studies consistently found that shoebox halls, with their relatively narrow dimensions and parallel side walls, produce stronger lateral reflections and more consistent sound quality from seat to seat than many surround or vineyard layouts (surround and vineyard halls are often used interchangeably in casual discussion).[9,10] The strong lateral reflections increase orchestral dynamics and can make sound feel fuller to listeners, helping explain the long success of rectangular shoebox halls.[9] However, while side walls remain important for lateral reflections, a study found that ceiling absorption changed reverberation more than side-wall absorption.[11] This finding shows that acoustic materials can't be considered separately from where they are placed, because the same material can have different effects depending on its location within the hall. Fan-shaped halls can provide greater seating capacity and sightlines, but their wider geometry can reduce strong lateral reflections. Vineyard-style halls have seating areas around the performance space, allowing surfaces around different sections of the audience to provide additional reflections. These differences show how room geometry affects the way sound reaches an audience.[5]

Table 2. Concert hall comparison & characteristics. Table adapted from Beranek 2006,[8] Beranek 2008,[2] Beranek 2015,[12] and Jeon *et al.*[13]

Hall	Hall Type	Reverberation Time - occupied (s)	Number of Seats	Upholstered Seats	Volume (m^3)	Volume per Seat (m^3)
Boston, Symphony	Shoebox-shaped	1.9	2,625	Lightly	18,750	7.1
Munich, *Gasteig*	Fan	1.8	2,487	Medium	29,737	12
Tokyo, Suntory Hall	Arena / Surround	2.1	2,006	Lightly	21,000	10.5
Los Angeles, Disney	Vineyard / Terraced	1.85	2,265	Heavy	30,800	13.6
Carnegie Hall, New York City	Shoebox-shaped	1.8	2,804	Medium	24,268	8.7

American physics professor Wallace Clement Sabine created a mathematical formula to show that reverberation time depends on room volume and total sound absorption.[4] Volume per seat is a different measure for comparing concert halls of different sizes, but it can't independently predict reverberation time. Instead, it's context for understanding how room size relates to the acoustic conditions of each hall.

Studies on layout focused on how stage and orchestra position, canopies, and seat location change what performers and audiences hear.[14,15]

The size and position of the stage can influence how directly sound reaches listeners. Research on orchestra stages found a narrow, high stage exposed to the main hall was most beneficial for symphony players.[16] This suggests that stage geometry is not only an architectural consideration but also an important part of how musicians hear and respond to one another during a performance. Acoustic canopies suspended above the orchestra also can reflect sound back toward performers and toward nearby audience members.[5] These canopies can help musicians hear one another, supporting timing and ensemble performances.[5] This demonstrates that acoustic planning must consider performers as well as audience members, since the same hall must provide useful sound conditions on stage and throughout the seating area. Seating arrangement also matters because engineers aim to provide listeners throughout the hall with similar sound quality.[5] A vineyard hall study showed that this arrangement can work well, but seating layout strongly controls lateral energy and the relationship between the performance and the audience.[17] This demonstrates that even within the same room shape, changes in seat location can produce different acoustic conditions. These layout choices therefore work together to shape how sound moves throughout the hall.[5] The finding supports that room shape alone does not

determine the final acoustic result. Other elements can modify how sound reaches both performers and listeners.

Room geometry usually sets the acoustic pattern, while materials, stage elements, and seating are ways to correct and fine-tune the acoustic environment.[18] Materials then allow engineers to further control sound within the hall. Different materials absorb and reflect different amounts of sound energy, and their effects can also vary depending on frequency.[5] Hard, reflective materials generally produce stronger reflections, while soft, porous materials tend to absorb more sound energy.[3] Concrete, mineral wool, fabric, wood, and acoustic panels can therefore serve different acoustic purposes depending on their placement.[3,5,6]

Material studies show that the placement of materials is just as important as the material itself and how much is used.[7,11] Wood and other reflective materials can help maintain sound energy within a hall, while curtains, carpets, and acoustic panels can reduce excessive reverberation and unwanted reflections.[3,6] Certain areas may specifically require absorption, such as rear walls or balcony ceilings where reflections could create distinct echoes.[6] Curtains can also be concealed when their absorption would otherwise remove useful reflections from the stage.[3] This shows that engineers must consider both the properties and placement of materials.[3,5] A material that improves acoustics in one location may have a different effect in another location, so selecting a material is only one part of the process.

Audience seating creates another material-related challenge because people themselves absorb sound. Concert hall seats are therefore sometimes designed to provide acoustic absorption similar to that of an occupied seat, helping the hall maintain more consistent acoustics between rehearsals and performances.[2] Measurements across 10 halls quantified absorption for occupied and unoccupied seats, improving planning for stage reverberation estimates.[19] This demonstrates why engineers must think carefully about occupancy when predicting a hall's reverberation, since the audience itself becomes part of the room's acoustic behavior. As a result, acoustic properties of seating add another factor to the way engineers plan a hall.[2] Because occupancy changes sound absorption, engineers must consider the audience when predicting how a hall will sound during an actual performance rather than depending only on measurements from an empty hall.

Overall, room shape and acoustic materials are connected ways to control sound throughout a concert hall.[5] Different combinations of these elements can therefore produce different acoustic conditions for performers and audiences.[3,5] The next question is how these choices appear in actual concert halls and how different halls use different engineering solutions.

Examples from Famous Concert Halls

This section compares engineering choices in well-known concert halls and how those choices affect acoustics. The examples show how different combinations of room shape, materials, surfaces, and acoustic design can produce different results.

Concert halls provide useful examples of how engineering choices affect the way music is heard. Although successful halls can have very different shapes and materials, each must

control reverberation and other factors according to the needs of performers and audiences. Comparing several well-known halls shows that there is not one physical arrangement that produces good acoustics. Instead, different combinations of engineering choices can produce successful acoustic conditions.

Royal Albert Hall in London, completed in 1871, provides an early example of the acoustic challenges that can result from a building's shape and materials. Its unique elliptical arena design and extensive wooden paneling are important parts of its acoustic character.[20] However, the hall's large volume, highly reflective large glass and iron dome, and curved geometry used to produce a severe echo, making the hall notorious for difficult acoustics before later renovations.[4,20,21] The hall's volume is larger than usual concert halls.[20] Engineers therefore had to make changes to reduce the delayed reflections produced by the building's shape and surfaces. Since the historic geometry and volume could not be changed, engineers improved its acoustic problems using acoustic treatments and technology. These included "mushrooms" beneath the dome to reduce reflections and improve sound distribution throughout the hall. The hall demonstrates how architectural choices can create acoustic problems when a building's shape and surfaces produce excessive reflections.[20] Multiple engineering elements can work together to correct an existing hall. This example shows the importance of considering acoustics when making architectural choices, since the shape and materials can strongly affect the quality of sound within a concert hall.

Carnegie Hall in New York City, completed in 1891, provides another example of how architecture, materials, and later renovations can affect acoustics. It has a reverberation time of 1.8 seconds,[2] and a volume of 24,270 m^3.[8] main auditorium follows the traditional rectangular shoebox shape, with high ceilings and tiered seating that keeps audiences relatively close to the performers while allowing sound to reflect efficiently between parallel walls.[22] Within this overall rectangular design, the hall's architectural elements include curved surfaces and thick walls, which affect reflection patterns and create its characteristic acoustics. The hall is also known for strong early lateral reflections.[10,18,22] Carnegie Hall therefore demonstrates how a traditional shoebox design can use architecture and surfaces to create the acoustics. This provides a useful transition to a concert hall where acoustics were calculated scientifically before construction.

Boston Symphony Hall, completed in 1900, is an important example of early scientific acoustical engineering. The rectangular, shoebox-shaped hall was built for the Boston Symphony Orchestra and is widely regarded as the first acoustically designed hall.[2] Sabine created his reverberation formula before construction, and the hall achieved a target reverberation time of 1.9 to 2.1 seconds.[4] Its walls, ceiling, and floor slope inward to direct sound toward the audience, while the spacing between rows of seats was also part of the audience acoustic design.[4,12,18] This shoebox hall combines geometric choices and architectural elements such as a coffered ceiling, columns, and leather seats.[12] Rather than depending on one acoustic factor, Boston Symphony Hall combines its rectangular shape, seating, and calculated reverberation time to create its

well-known orchestral sound.[23] This early use of mathematical acoustics provides a contrast with halls that later required corrections after construction.

The Sydney Opera House, a fan-shaped hall, completed in 1973, faced poor sound until a major retrofitting was completed in 2022. As part of the renovation, 18 specially designed acoustic reflectors, known as "petals" were added above the stage to replace the previous clear reflectors. The reflectors can be raised or lowered according to the performance.[24] This change shows how engineers can modify an existing hall through specific elements rather than major architectural changes. The retrofit therefore demonstrates that acoustic problems don't always require changes to a hall's basic geometry. Other elements can provide extra control. Sydney provides a useful example of how acoustic engineering can respond to problems after construction.

Walt Disney Concert Hall in Los Angeles, completed in 2003, provides a more modern example of acoustic engineering. Designed by Frank Gehry and acoustically tuned by Yasuhisa Toyota, the hall uses a vineyard-style seating arrangement around the stage.[17,25] The exterior was designed to look like "sails" in the wind, while the interior uses warm Douglas fir.[25] The combination of its curved architecture and seating arrangement controls sound throughout the hall.[17,25] During design, the acoustical team tested a 1:10 scale model filled with nitrogen to simulate real sound conditions before construction.[25] Even the heating, ventilation, and air conditioning system was shaped with both airflow and sound clarity in mind.[25] The hall therefore combines its unusual shape with carefully selected indoor materials and acoustic reflectors.[17,25] These choices show how modern concert halls can include acoustical considerations throughout the process rather than correcting problems through later changes.

These five concert halls show that successful acoustics can result from very different combinations of engineering choices. Royal Albert Hall demonstrates the acoustic challenges that can result from a building's shape and materials.[20] Carnegie Hall demonstrates how a traditional shoebox design can use architectural features and strong lateral reflections to create acoustics.[9,10] Boston Symphony Hall was the first acoustically designed hall using mathematical calculations of reverberation time.[23] Sydney Opera House demonstrates how adding specific elements can improve an existing hall. Walt Disney Concert Hall combines vineyard-style seating, acoustic reflectors, interior materials, and acoustic modeling.[25] These examples support the idea that concert hall acoustics depend on multiple engineering choices working together rather than on a single factor.

Conclusion

Concert hall acoustics depend on multiple engineering factors rather than one element alone. The research reviewed in this paper shows that the behavior of sound in a concert hall, the use of room shape and acoustic materials, and the engineering choices made in individual concert halls all contribute to the sound experienced by performers and audiences.[5,23]

The first section showed that different elements affect how sound travels through a concert hall. Reverberation time is especially important because too much can make sound

unclear, while too little can produce a dry sound.[3] Understanding acoustic factors like reverberation time helps explain the engineering choices discussed in the second section.

The next section showed that engineers can control these acoustic factors through room shape, layout, and materials. Room geometry affects the paths that sound waves take, while reflective and absorptive materials determine how much sound energy remains in the room.[5,23] Stage placement, acoustic canopies, seating, and the placement of materials provide additional ways to control sound.[5,6,23] This suggests that successful acoustics depend on several characteristics working together rather than a single engineering element.

The examples of famous concert halls further support this conclusion. Although these halls use different shapes, materials, and technologies, each demonstrates the importance of combining multiple acoustic choices.

Overall, the research suggests that concert hall acoustics are well understood in terms of the individual factors that affect sound, including reverberation, room geometry, and materials.[5,23] However, questions continue to exist about how new technologies can be combined with acoustic principles to improve concert hall design. Future research could further look into how different new elements can work together to create more adaptable performance environments. Understanding these possibilities could help engineers make more specific acoustic choices while balancing the needs of performers and audiences.

AI Acknowledgement

During the preparation of this work, the author used Consensus to brainstorm research articles. After using this tool, the author reviewed the resulting articles and fully cited the content of each publication.

Acknowledgements

Thank you to Andrea Giusti, Marcin Kedziera, and my family.

Works Cited

(1)Ciaburro, G.; Iannace, G.; Lombardi, I.; Trematerra, A. Acoustic Design of Ancient Buildings: The Idea of Pompeii and Posillipo. Buildings 2020, 10, 224. https://doi.org/10.3390/buildings10120224

(2)Beranek, L.; Member, H. Concert Hall Acoustics—2008. 2012.

(3)Pham, M. Sound Check! Crafting Acoustics for Performance – Illumin Magazine - USC Viterbi School of Engineering. https://illumin.usc.edu/crafting-acoustics-for-performance/ (accessed 2026-07-23)

(4)Postma, B.; Katz, B. Forum—Pre-Sabine Room Acoustic Assumptions on Reverberation and Their Influence on Room Acoustic Design. J. Acoust. Soc. Am. 2020, 147, 2478–2487. https://doi.org/10.1121/10.0001082

(5)Christopher Kmiec. Concert Hall Acoustic Design: Engineering Guide. https://audiotek.com/concert-hall-acoustic-design/ (accessed 2026-07-25)

(6)Robert M. Brecht. The Basics of Concert Sound Production - TSE Entertainment. https://tseentertainment.com/understanding-the-basics-of-concert-sound-production/ (accessed 2026-07-30)

(7)Gao, J.; Zhao, Y.; Pan, L. The Effects of Sound Absorption of Stage House on the Acoustics of Auditorium in an Opera House. Buildings 2024, 14, 718. https://doi.org/10.3390/buildings14030718

(8)Beranek, L. Analysis of Sabine and Eyring Equations and Their Application to Concert Hall Audience and Chair Absorption. J. Acoust. Soc. Am. 2006, 120, 1399–1410. https://doi.org/10.1121/1.2221392

(9)Pätynen, J.; Tervo, S.; Robinson, P.; Lokki, T. Concert Halls with Strong Lateral Reflections Enhance Musical Dynamics. Proc. Natl. Acad. Sci. 2014, 111, 4409–4414. https://doi.org/10.1073/pnas.1319976111

(10)Hidaka, T.; Nishihara, N. Acoustical Quality in Concert Halls as Related to Hall Shape: Shoebox, Surround, and Other. Psychomusicology Music Mind Brain 2015, 25, 240–252. https://doi.org/10.1037/pmu0000050

(11)Yu, R.; Ma, E.; Fan, L.; Liu, J.; Cheng, B.; Jiang, Z. Effects of Different Building Materials and Treatments on Sound Field Characteristics of the Concert Hall. Buildings 2022, 12, 1613. https://doi.org/10.3390/buildings12101613

(12)Beranek, L. Concert Hall Design: Some Considerations. Psychomusicology Music Mind Brain 2015, 25, 181–186. https://doi.org/10.1037/pmu0000110

(13)Jeon, J. Y.; Jo, H. I.; Kwak, K.; Kim, H. The Design Goal of Acoustic Coefficients in Concert Halls Related to Reverberation and Loudness; 2018.

(14)Kriz, J.; Porcinai, E.; Lepa, S.; Klein, P.; Arend, J. M.; Weinzierl, S. How Early Reflections Affect the Stage Acoustic Conditions for Solo Musicians. Appl. Acoust. 2025, 235, 110644. https://doi.org/10.1016/j.apacoust.2025.110644

(15)Chen, Y.; Cabrera, D.; Alais, D. Modelling Audiovisual Seat Preference in Virtual Concert Halls. Appl. Acoust. 2023, 212, 109589. https://doi.org/10.1016/j.apacoust.2023.109589

(16)Dammerud, J. Stage Acoustics for Symphony Orchestras in Concert Halls, 2009. https://doi.org/10.13140/RG.2.2.35537.28003
(17)Jablonska, J. Architectural Acoustics in Vineyard Configuration Concert Hall. J. Archit. Eng. Technol. 2019, 07. https://doi.org/10.4172/2168-9717.1000224
(18)Beranek, L. Concert Hall Acoustics 1992. J. Acoust. Soc. Am. - J ACOUST SOC AMER 1992, 92, 1–39. https://doi.org/10.1121/1.404283
(19)Beranek, L. L.; Hidaka, T. Sound absorption in concert halls by seats, occupied and unoccupied, and by the hall's interior surfaces. https://doi.org/https://doi.org/10.1121/1.423957
(20)METKEMEIJER, R. THE ACOUSTICS OF THE AUDITORIUM OF THE ROYAL ALBERT HALL BEFORE AND AFTER REDEVELOPMENT; 2023. https://doi.org/10.25144/18331
(21)Royal Albert Hall. Grove Music Online. https://doi.org/https://doi.org/10.1093/gmo/9781561592630.article.23995
(22)Linsong, L. The Influence of Concert Hall Architecture on Acoustics (the Effect of Musical Performance). E3S Web Conf. 2023, 389. https://doi.org/10.1051/e3sconf/202338906016
(23)Prodec Group. Concert Hall Acoustic Design: Professional Standards & Solutions. https://www.prodecgroup.com/architectural-acoustic-standards/concert-hall-acoustic-design/ (accessed 2026-07-25)
(24)Taylor, L.; Claringbold, D. ACOUSTICS OF THE SYDNEY OPERA HOUSE CONCERT HALL Part One: The Client's Perspective; 2010.
(25)Toyota, Y. Acoustical Design of the Walt Disney Concert Hall in Los Angeles. J. Acoust. Soc. Am. 1999, 105 (2_Supplement), 987. https://doi.org/10.1121/1.425376

Authors

Jaydzia Khan is an analytical student planning to study physics and engineering. She combines her strengths in mathematics with her passion for music and collaborative problem-solving. A competitive athlete and student leader, she mentors younger students and advocates for equity in STEM. She hopes to solve real-world problems using engineering.

The Psychological Effects of First-Year Engineering Course Work By Mark Kassih

Abstract

First-year engineering students are often challenged with rigorous coursework and increased academic stress once transitioning to college-level coursework, playing a major role in shaping their resilience and motivation to engage with challenging coursework. This study investigated the extent to which a first-year engineering student's mindset toward failure influences their resilience and motivation to engage during rigorous hands-on coursework in an Introduction to Design course. A mixed-method approach was employed, bringing together a pre-activity survey with observational research conducted at California State Polytechnic University, Pomona, involving only first-year engineering students during a complicated design activity. Based on their pre-activity survey responses, the twenty-seven student participants were categorized as demonstrating either a growth mindset or a fixed mindset. Through the use of a structured rating scale, student engagement, persistence, resilience, and leadership were all observed and evaluated. After the design activity was completed, results indicated that students categorized as having a growth mindset consistently demonstrated higher levels of resilience, participation, and motivation to engage than students with a fixed mindset. Although limitations such as observational bias and the inability to establish causation in a correlational design persist, the findings still suggest that a mindset towards failure is associated with meaningful influences toward student engagement and resilience. These findings highlight the value of promoting growth mindset practices within first-year engineering education to better support students as they transition towards more demanding engineering coursework.

Introduction

A study conducted at the University of Toronto found that the workload of undergraduates partaking in an engineering course almost doubled after being in the class for only 3 weeks (Gerrard et al. 5). It was found that this increased workload impaired learning and inclusion for students making a shift from high school to college due to the significant change in stress levels (Gerrard et al. 5). This stress first-year engineering students are experiencing can be better explained with the rise of engineering culture, where students not only are being overworked, but also in competition with other students to become “better” than each other (Mirabelli et al.). Engineering culture has the effect of fostering high stress levels, exclusion, and a disregard for students' overall well-being, physically and mentally, through less sleep and constant stress (Mirabelli et al.). In an attempt to limit stress levels towards first-year engineering students, Cambridge University Press published a study where they included a brief mindfulness-based intervention before having first-year engineering design students partake in activities such as concept generation, concept selection and physical modeling (Nolte, et al.). With the inclusion of this short intervention, authors Hannah Nolte and others concluded that there was no significant difference in stress levels with and without the intervention, resulting in

the finding that the stress first-year engineering students are feeling is invoked from the workload and time pressure.

By understanding the stress that follows the heavy course work given to first-year engineering students, it can now be understood how this stress can impair the resilience and motivation to engage for first-year engineering students. When choosing a career path, the University of Washington concluded that 40% of freshman engineering students leave this given major by their senior year (University of Washington). By highlighting the fluctuations in resilience and engagement rates, incoming undergraduate students can better understand the situation they are getting themselves into when deciding to major in engineering. Therefore, it is in our best interest to investigate the extent to which a first-year engineering student's mindset towards failure during rigorous hands-on engineering coursework influences their resilience towards the task and motivation to engage within an Introduction to Design classroom.

Literature Review

1. Steps to Influence Resilience

When exposed to rigorous hands-on course work, first year engineering students have numerous fluctuations in their resilience during the tasks. In order to enhance resilience, Anita L. Campbell and others suggest the importance of cultivating a growth mindset to magnify an engineering undergraduates' resilience, persistence, and engagement with challenging coursework (Campbell et al. 503). When participating in challenging engineering course work, the inner ideology that the assignment is possible with perseverance greatly aids in engineering students' resilience towards the task, thus allowing the overall task to be completed more efficiently. In contrast, Jason Richard Power and others argue that self-efficacy, a student's belief about their capability to complete a task, is actually the most important factor for influencing performance, persistence, and resilience in engineering classes for undergraduates (Power et al. 2). Even though Campbell and Power have altering beliefs for the most important factor in enhancing resilience, it cannot be discredited that a growth mindset and self-efficacy do go hand-in-hand with both ideologies focusing on a positive mindset towards the task, proving how both authors do believe a persevering mindset to be linked to an increased resilience. Consequently, Madhurima Das and Maria C. Yang take an alternative approach than mindset in enhancing resilience and identify that design-based courses assist in teaching engineering students about failure and how to overcome consistent obstacles, allowing for an increase in resilience, creativity, and problem solving skills (Das and Yang 8). The discovery that the courses taken by first-year engineering students actually enhances resilience is divergent from all other forms of research and proves the magnitude teaching about failure can have on resilience. Another method to enhance resilience for first-year engineering students was unpacked by Sophia S Hall and others where they incorporated an ecological model of resilience in STEM learning and found that it actually influenced resilience towards rigorous STEM related course work (Hall et al.). Currently, there are sufficient and quantitative methods of influencing

resilience towards first-year engineering students through factors such as mindset and course selection.

2. Importance of Resilience in Engineering Education

After understanding the diverse ways a first-year engineering student's resilience can be influenced, it can lead some to wonder the importance of resilience for engineering majors. Well, Ann-Kristin Winkens and others argue for resilience thinking to be incorporated into the curriculum for engineering undergraduates due to its psychological benefits for students (Winkens et al. 522). They further discuss how teaching resilience is conceptually important for engineering students but require new teaching methods to mold resilience into students since the current form of instructional design is inadequate (Winkens et al. 529). From another perspective, Kathryn M. Bateman and Christine M. Cunningham argues that framing failure as a negative concept in engineering education is ineffective when failure should be interpreted as a step in the design process (Bateman and Cunningham 64). Perceiving failure as a positive milestone in the design process allows for students to develop resilience and problem solving skills, benefiting their education as a whole (Bateman and Cunningham 66). Even though Winken's conclusion may be different from Bateman and Cunningham's conclusion, both forms of research effectively analyze the importance of obtaining resilience in education. Additionally, the definition of resilience that everyone is familiar with today, that it is a positive attribute and contributes to success in higher education, requires a more nuanced understanding of how engineering professors conceptualize it before being able to foster it into engineering programs (Wint and Ines). Victoria Wint and Ines Direito emphasize that people who define resilience as bouncing back will differ in the ways they teach resilience than people who define resilience as adapting, proving the importance of establishing a definitive explanation for resilience and how to incorporate it in engineering education.

3. Influences on Engineering Students' Motivation to Engage

When exposed to challenging hands-on engineering coursework, a first-year engineering student's motivation to engage can vary. Meredith A. Henry and others suggest that the current tool for measuring fear of failure towards STEM students does not fully account for the unique challenges STEM students face, and should overall be modified (Henry et al. 4). This tool is classified as the Performance Failure Appraisal Inventory and currently misinterprets a STEM students fear of failure which is an influential factor for student persistence, motivation, and willingness to to take risks when working on challenging coursework (Henry et al. 17). Henry and other authors highlight the importance of understanding a student's fear of failure in order to further understand their motivation to engage, proving a correlation between fear of failure and motivation to engage. In the same manner, Kyle F. Trenshaw and others add onto the research conducted by Henry and his coauthors to argue that engineering educators should structure some of their learning around self-determination theory in order to improve engagement and motivation (Trenshaw et al.). Furthermore, Trenshaw concludes that when students have

achieved a state of competency, autonomy, and relatedness that they are more motivated to engage more thoroughly in the lesson (Trenshaw et al.). The research conducted by Trenshaw and his coauthors prove how an engineering student's motivation to engage could be influenced by the actual content being taught, as opposed to the research conducted by Henry and his coauthors who argue a mindset towards failure being influential for a motivation to engage. Additionally, Jennifer Xu and others highlight throughout their research study that students who perceive their success as a result of personal efforts and hard work tend to show more motivation when participating in coursework as opposed to those who believe success to come with natural talent or luck (Xu et al. 6). Xu and coauthors prove how motivation does in fact stem from an engineering student's overall mindset, similarly to how Henry and his coauthors concluded that an engineering student's mindset towards failure influences their motivation.

Gap in the Body of Knowledge

Current research on a first-year engineering student's resilience and motivation to engage has addressed numerous factors that influence resilience and motivation, leaving a gap in the literature on the influence of resilience and motivation among first-year engineering students during hands-on engineering coursework in an Introduction to Design course. Due to the diverse environment and potentially different levels of rigor within the coursework, the addition of hands-on coursework in an Introduction to Design course could yield results that differ from the pre-existing knowledge of first-year engineering students as a whole. To address this issue, I will conduct a mixed-methods study comprising correlational and observational research and a survey.

Hypothesis

With regards to the research question: To what extent does a first-year engineering student's mindset towards failure during rigorous hands-on engineering coursework influence their resilience towards the task and motivation to engage in an Introduction to Design course, I predict there to be a high correlation between a growth mindset to resilience and a motivation to engage in an Introduction to Design course. I predict this correlation to be greater than the correlation between a fixed mindset and resilience and motivation. I expect students who portray a growth mindset to show numerous forms of resilience, persistence, and engagement throughout the task, whereas students with a fixed mindset will exhibit characteristics of uncertainty, burnout, and disengagement, with very limited resilience and engagement.

Methodology

To answer the question: To what extent does a first-year engineering student's mindset towards failure during rigorous hands-on engineering coursework influence their resilience towards the task and motivation to engage in an Introduction to Design course, a mixed-method study has been conducted. Within this mixed-method study, I combined correlational research with observational research and a survey to best accumulate all the data required in order to

clearly determine a correlation between a student's mindset towards failure and their motivation to engage and resilience. In this study, the independent variable was each student's mindset towards failure, whether fixed or growth. The dependent variable consisted of each student's motivation to engage and resilience as I tracked these fluctuations through observational research.

Before diving into my methodology, it's important to understand what a mindset towards failure actually means and how a fixed mindset differs from a growth mindset. A mindset towards failure simply refers to a person's psychological interpretation of setbacks and challenges primarily within learning environments. In this study, I will examine two distinct mindsets towards failure: a growth mindset and a fixed mindset. A growth mindset is a psychological view of setbacks and challenges as temporary obstacles that can be overcome to accomplish a task, whereas a fixed mindset is a psychological view of setbacks and challenges as permanent obstacles that ultimately lead to failure. In the context of this study, resilience was defined as a student's ability to persevere and continue to show effort throughout the duration of the task. The students' motivation to engage was defined as each student's ability to remain active and participate thoroughly throughout the entire task.

Transitioning into my experiment, this mixed-methods study was conducted at Cal Poly Pomona, specifically within their Introduction to Design class, composed of first-year students enrolled in majors such as mechanical engineering, manufacturing engineering, and industrial engineering.

The experimental design was conducted over two days at Cal Poly Pomona, with the first day serving as an introduction to the students about what my experiment would look like and whether they would choose to participate. First, I passed out a consent form to all the students that outlined the entire experiment and their role in data collection. Upon signing the consent form, all participants provided consent prior to participating in the study, and all data collected was kept anonymous for student safety. After passing out the consent forms, I then passed out a document that contained a QR code for the students to scan and complete a short survey that contained the following questions:

Table 1

Survey questions for college students

Question	Students questioned	Students answered
What is your initial reaction when failing at a piece of coursework?	43	27
Do you interpret failure as something that is beneficial for learning or something that shows your limits in certain areas?	43	27
How did success or progress affect your motivation to continue?	43	27
How confident did you feel that you could eventually succeed, even after initial failure?	43	27

Each of these survey questions plays a key role in aiding me to determine each student's mindset towards failure. Although only 27 of the 43 present students agreed to participate in my experimental design, this granted me with more than enough students to analyze fluctuations in resilience and motivation during the task. I was able to effectively categorize each student as having a growth mindset or a fixed mindset by analyzing their responses to each question thoroughly. If a student responded to most questions in a way that indicated they could frame failure as necessary for success, I considered them to have a growth mindset. However, if the student responded to a majority of the questions in a way that indicated failure as a limitation within their ability to complete the task, I considered them to have a fixed mindset.

After passing out the document containing the QR code on the first day, I announced to the students that their participation in this study was completely up to their own discretion and that if they chose to participate, they simply had to bring the consent form signed the next day, in addition to completing the online survey the day before the actual experiment began.

On the second day at Cal Poly Pomona, I began by having each student who answered the survey questions in a growth-mindset way wear a blue penny, and each student who answered in a fixed-mindset way wear a green penny. Before beginning the task, I ensured that each team comprised 3 people with both growth and fixed mindsets. This system I created helped me keep track of students with a fixed mindset and a growth mindset. Then, the task had begun. The goal of this coursework was to build a platform composed solely of chopsticks and tape that could hold a 10-pound ball. An observational scale was replicated in my notepad to effectively record student behavior during the task. The statements are stated below.

1. This student shows participation throughout the whole length of the task.
2. This student continues to try when failing at the task once.
3. This student continues to try when failing at the task twice.
4. This student actively participates and engages with the task towards the end of the time limit.
5. This student took charge in their group of 3.

For each statement in my notebook, I used a structured scale rather than a yes/no answer to effectively measure the extent of resilience and motivation to engage. Each student would be assigned a score from 0 to 3 based on observed persistence and motivation for each of the questions listed above, with 0 showing no signs of correlation to the question and 3 showing perfectly aligned characteristics. The coursework lasted roughly 70 minutes, and within that timeframe, I was able to effectively account for each student's participation, resilience, and engagement throughout the task.

The collected data from the observational research were analyzed by comparing fluctuations in resilience and motivation among fixed-mindset and growth-mindset students throughout the task. Using percentages, I effectively identified patterns in resilience and

motivation among growth-mindset and fixed-mindset students. The survey responses were further analyzed to identify differences between growth-mindset and fixed-mindset students.

Choosing to conduct a mixed-method study containing correlational research, observational research, and survey questions was credited to an article published by the Cambridge University Press where authors Hannah Notle and others conducted an experimental design where they had "First-year engineering design students [complete] three 30-minute experimental sessions during an engineering design course, where their stress and mindfulness during three principal stages of the design process were investigated" (Notle et al.). By analyzing the effectiveness of this instance of observational research through an experimental design task, I was inspired to replicate it by incorporating observational research into my own experimental design. I further enhanced my methodology by incorporating correlational research, following the study conducted at the 2023 ASEE Annual Conference & Exposition, in which authors Abigail Nichole Lehto and Ning Fang used existing research to determine a correlation between student engagement and GPA among engineering students (Lehto and Fang 4). After closely analyzing the studies conducted in these two articles and the successes from them within the engineering field, I was able to articulate a methodology that incorporated a survey, observational research, and correlational research that combines quantitative comparison with qualitative observation, accounting for a clear correlation between the influence of a mindset towards failure on resilience and motivation within engineering students.

To reduce observer bias in the study, fluctuations in behavior were determined using predefined categories without subjective interpretation. However, since I was the only researcher, some level of observer bias will persist throughout the study.

Results

After collecting extensive data from my experiment through survey questions, observational research, and correlational research, the results of this study effectively portray how differences in a mindset towards failure can influence resilience and motivation for first-year engineering students.

To begin, the pre-activity survey was provided to all students participating in the experiment, and contained the following responses:

Table 2

Survey responses for college students

Question	Students who responded in a manner that signified a growth mindset	Students who responded in a manner that signified a fixed mindset
What is your initial reaction when failing at a piece of coursework?	15	12
Do you interpret failure as something that is beneficial for learning or something that shows your limits in certain areas?	16	11

How did success or progress affect your motivation to continue?	15	12
How confident did you feel that you could eventually succeed, even after initial failure?	17	10

Carefully examining each student's responses, I identified whether they had a growth or a fixed mindset. Although the results were not linear, and some students answered questions that signified both a growth mindset and a fixed mindset, I carefully examined each student's responses as a whole and identified them as a growth mindset student if they portrayed mainly characteristics that indicate their ability to view failure as necessary and a step in the right direction. In the same manner, I was able to identify students with a fixed mindset towards failure by observing characteristics that indicated they viewed failure as a limitation of their capabilities. At the end of the analysis, I identified 15 students with a growth mindset and 12 with a fixed mindset.

After carefully analyzing survey responses, I could begin the observational research on day 2. Throughout the experiment, I recorded a scale from 0 to 3 indicating how closely the students' characteristics during the task matched each of the 5 statements. When I noticed a student's characteristics aligned perfectly with the statement represented in my notebook, I would give them a score of 3, whereas if their characteristics did not align with the statement whatsoever, I would give them a score of 0. This system was first replicated with students who wore a blue penny and demonstrated a growth mindset.

Table 3

Growth Mindset Students During Task

Growth Mindset students	0 (No correlation)	1(little correlation)	2 (moderate correlation)	3 (highly correlated)
This student shows participation throughout the task.	1	3	6	5
This student continues to try once, failing at the task once.	0	2	7	6
This student continues to try once, failing at the task twice.	2	2	7	4
This student actively participates and engages with the task towards the end of the time limit.	1	3	6	5
This student took charge in their group of 3	0	4	5	6

As shown in the table above, a majority of students with a growth mindset were at the higher end of the spectrum, indicating a strong correlation with characteristics of engagement, resilience, persistence, and leadership. Even though a majority responded this way, a handful of students still showed little to no engagement or resilience.

When looking more closely at the data, only 4 of 15 students consistently demonstrated little to no engagement and resilience during the tasks, accounting for roughly 27% of the total group. On the other hand, when analyzing the number of students who consistently demonstrated moderate to high levels of engagement and resilience, I found that 11 out of 15 (73%) students exhibited these trends. 73% of students exhibiting a strong association to engagement and resilience levels highlights how a majority of students that contain a growth mindset contain characteristics that align with having resilience towards coursework and a strong motivation to engage during the coursework.

When looking at the questions and data more specifically, questions 1, 2, and 5 mainly assess the students' motivation to engage, while questions 2 and 3 mainly assess the students' resilience. Even though these questions may present slighty different psychological characteristics, the trend remained constant throughout the questions assessing resilience and the questions assessing motivation to engage.

After carefully gathering data on all growth-mindset students, I shifted my focus to students with a fixed mindset towards failure. These students were wearing green pennies and could be easily identified in the group. There were 12 students who contained a fixed mindset.

Table 4

Fixed Mindset Students During Task

Fixed Mindset students	0 (No correlation)	1(little correlation)	2 (moderate correlation)	3 (highly correlated)
This student shows participation throughout the task.	2	7	3	0
This student continues to try once, failing at the task once.	3	6	2	1
This student continues to try once, failing at the task twice.	6	5	1	0
This student actively participates and engages with the task towards the end of the time limit.	4	6	2	0
This student took charge in their group of 3	4	4	2	2

Of these 12 students with a fixed mindset towards failure, a majority were at the lower end of the spectrum, showing little to no correlation with statements expressing qualities such as resilience and engagement. Although the majority of students were towards the end of the

spectrum, a couple of students showed characteristics indicating they retain resilience and engagement.

When looking more closely at the data, 9 out of 12 fixed-mindset students consistently showed little to no characteristics associated with high resilience and engagement, accounting for 75% of the population. On the other side of the spectrum, 3 out of 12 fixed-mindset students consistently showed moderate to high levels of characteristics associated with resilience and engagement, accounting for only 25% of the population. When observing the data as a whole, a great majority of the population of fixed-mindset students showed little to no engagement and resilience as opposed to the minority of the fixed-mindset students who showed high levels of engagement and resilience.

When directly comparing both groups, clear differences emerged from both stages of data collection. Growth mindset students had only 27% of their population consistently score between 0 and 1 on the scale, while fixed mindset students had 75% of their population consistently score between 0 and 1, clearly highlighting the significant differences between the groups. In contrast, growth mindset students accounted for only 73% of their population who consistently scored between 2 and 3 on the scale, while fixed mindset students accounted for 25% of their population who scored between 2 and 3.

Discussion

This study sought to determine how a first-year engineering student's mindset toward failure influences their resilience and motivation to engage in an Introduction to Design course. After conducting data and observing the results, the findings suggest that the students who obtain a growth mindset tend to consistently demonstrate higher levels of persistence, engagement, and resilience, whereas fixed mindset students tend to consistently demonstrate lower levels of these traits. Overall, these results suggest the significance of a mindset towards failure in influencing resilience as well as a motivation to engage.

These results align directly with my hypothesis and further support the findings from similar research studies. For example, from an article published under *ResearchGate,* author Jaime A. Gutierrez-Monsalve and others used correlational research as well as quantitative reasoning to definitely identify a relationship among academic performance within engineering students and their demographic background and work ethic (Gutiérrez-Monsalve et al.). This experimental study correctly identifies a correlation between academic performance and work ethic, supporting the findings from my personal research, which show a correlation between a mindset towards failure and increased academic engagement and resilience. Furthermore, correlational research published in the *Institute of Electrical and Electronics Engineers* demonstrates that success in engineering is not correlated with intelligence. but with a person's persistence in overcoming challenges (Sandoval and Zapata). The results expressed in this source directly relate to how a growth mindset student in this study showed higher levels of engagement and resilience, and since a higher level of engagement and resilience is linked to greater academic performance, as expressed by Gutiérrez-Monsalve et al., it can be concluded that there

is a relationship between growth mindset students who retain qualities such as perseverance and resilience and higher. academic success.

Diverting the attention back to my personal study, some of the data collected from my methodology still remain unexplained. As previously mentioned, a majority of growth mindset students showed higher levels of resilience and engagement, whereas a majority of fixed mindset students showed lower levels of both, as expected in this study. But there were still 25% of fixed-mindset students who demonstrated high resilience and engagement, and 27% of growth-mindset students who demonstrated low levels of resilience and engagement, which may seem unusual. These unexpected results are primarily due to two factors, one of them being my method in determining whether a student portrayed a growth mindset or fixed mindset. This method was through an online survey where students could have simply submitted results that do not definitely articulate their characteristics, possibly making me categorize them as one mindset towards failure when they are the other one. Another potential reason for these unexpected results could be confounding variables I did not account for. Since this is correlational research with observational research, all confounding variables are not accounted for in this experiment and fluctuations in student resilience and engagement could be a result to outside factors and not simply their mindset towards failure.

When combining all the existing research from various forms of methodologies and experiments, it has already been established that engineering students see higher academic performance when met with an increased level of engagement and work ethic. This understanding is very closely related to the idea that growth-mindset students exhibit higher levels of engagement and resilience, which makes my research seem insignificant until I realize that it focuses only on first-year engineering students in an Introduction to Design classroom, making the study different from anything published before. Looking primarily at first-year engineering students highlights numerous differences from engineering students as a whole, since they are still adjusting to the transition and are more likely to experience fluctuations in their mental health due to stress and other factors.

In all, these findings clearly show that growth-mindset students correlate with increased resilience and motivation to engage, suggesting that future professors should implement practices that promote a growth mindset among students to ensure everyone is successful in the classroom. Although these results are only relevant to students in an Intro to Design course, this does not diminish the fact that each student's mindset towards failure is consistent and indifferent to the type of course they are in, meaning these results could potentially apply to other courses within the same field.

Limitations

Throughout the entire research process, everything went as expected and my results clearly align with the projected results from previous literature. But, there still remain some limitations that simply could not be avoided within this experiment.

For starters, this experiment had a smaller sample size: only 15 people categorized the growth-mindset students and 12 categorized the fixed-mindset students. Although this number may seem very low, the smaller sample size actually aided my observational research task, allowing me to focus more closely on a smaller group of students and gain a stronger understanding of their psychological fluctuations. Being a solo researcher made data collection much more difficult, but the smaller sample size allowed for significant time spent on each student.

Furthermore, since this study used an observational design, I could not eliminate all forms of researcher bias. Even with the inclusion of a structured scale to observe the students' psychological and physical characteristics, my discretion of their resilience and engagement levels could always be different than other researchers' discretions. In my experiment, to limit this bias as much as possible, I spent extra time with each student and viewed each student from a neutral perspective to produce the best results with the least bias.

Finally, since this mixed-method experiment also included correlational research, a definitive causation-based relationship between a mindset towards failure, resilience, and engagement rates could not be determined. Correlational research always leaves room for confounding variables to influence the results of the experiment, which could also have been a concern in my experiment, with roughly 25% of students showing characteristics not aligned with the rest. But these confounding factors are not significant enough to discredit the information and data collected from this study.

Future Directions

With the data collected from this study, there are numerous directions for future research, creating a fundamental study that articulates additional variables and real-world implications. For instance, all this data was collected within a single Introduction to Design course, allowing for future research to branch towards other engineering courses. Furthermore, this study included only first-year engineering students, allowing future research to investigate the influence of a mindset towards failure on resilience and engagement among older engineering students. This idea will further allow comparisons and differences to be drawn between first-year engineering students and future engineering students regarding resilience and motivation to engage, with respect to their mindset towards failure.

Looking beyond the engineering field, future researchers could build on this study's findings to examine the extent to which a mindset towards failure influences resilience and engagement, and compare it with that in other majors. Researchers could look at other STEM fields, such as Biology and Chemistry, which are closely related to engineering but distinct enough to yield significant results that add new insights to existing research.

Conclusion

This study suggests that, to a certain extent, a first-year engineering student's mindset towards failure influences their resilience and motivation to engage in an Introduction to Design

course. The results demonstrate a clear correlation between a growth mindset in first-year engineering students and higher levels of resilience and motivation to engage. The results also suggest a correlation between a fixed mindset among first-year engineering students and lower levels of resilience and motivation to engage in an Introduction to Design course. Overall, these findings emphasize the importance of a mindset towards failure in shaping a student's work ethic, resilience, persistence, and engagement, potentially influencing future students to adopt the characteristics of a growth mindset. The results of this research could also aid incoming engineering students making the transition from high school to college in understanding the benefits of a growth mindset and potentially developing one of their own. At a young age, a student's mindset towards failure is easily shaped by their environment and capabilities, making it beneficial for many to transition from a fixed to a growth mindset.

Works Cited

Atwood, Sarah A., et al. "Work-in-Progress: Seizing failure as an opportunity to learn: Undergraduate engineering students' conceptions of failure and iteration." American Society of Engineering Education Annual Conference, National Science Foundation, 2024, https://par.nsf.gov/biblio/10526170-work-progress-seizing-failure-opportunity-learn-undergraduate-engineering-students-conceptions-failure-iteration?utm.

Bateman, Kathryn M. and Christine M. Cunningham. "Turning Failure into Success: How Engineering Challenges Build Resilience and Problem Solving Across Contexts." Connected Science Learning, vol. 7, no. 3, Taylor and Francis, 2024, pp. 60-68, https://www.tandfonline.com/doi/full/10.1080/24758779.2025.2484352?src=exp-la&utm_source.

Campbell, Anita L., et al. "Developing growth mindsets in engineering students: a systematic literature review of interventions." European Journal of Engineering Education, vol. 46, no. 4, Taylor and Francis, 2021, pp. 503-527, https://www.tandfonline.com/doi/full/10.1080/03043797.2021.1903835?utm_source.

Das, Madhurima and Maria C. Yang. "Design Experiences as Pathways for Embracing Failure." MIT Libraries, vol. 4, American Society of Mechanical Engineers, 2021, https://dspace.mit.edu/bitstream/handle/1721.1/154884/v004t04a011-detc2021-71419.pdf?isAllowed=y&sequence=2&utm_source.

Gerrard, Darlee et al. "Are Students Overworked? Understanding the Workload Expectations and Realities of First-Year Engineering." ResearchGate, 2017, https://www.researchgate.net/publication/316736951_Are_Students_Overworked_Understanding_the_Workload_Expectations_and_Realities_of_First-Year_Engineering.

Gutiérrez-Monsalve, Jaime A., et al. "Factors related to academic performance among engineering students: a descriptive correlational research study." ResearchGate, vol. 90, no. 227, 2023, pp. 35-44, https://www.researchgate.net/publication/372372026_Factors_related_to_academic_performance_among_engineering_students_a_descriptive_correlational_research_study.

Hall, Sophia S., et al. "Resilience, science, technology, engineering, and mathematics (STEM), and anger: A linguistic inquiry into the psychological processes associated with resilience in secondary school STEM learning." PubMed Central, vol. 92, no. 3, National Library of Medicine, 2022, https://pmc.ncbi.nlm.nih.gov/articles/PMC9545978/?utm_source.

Hickey, Hannah. "Engineering overdrive: Undergrads in field overworked, not informed about their major, study finds." UW News, 2007 https://www.washington.edu/news/2007/07/05/engineering-overdrive-undergrads-in-field-overworked.

Henry, Meredith A., et al. "Quantifying fear of failure in STEM: modifying and evaluating the Performance Failure Appraisal Inventory (PFAI) for use with STEM undergraduates." International Journal of STEM Education, vol. 8, no. 2, SpringerOpen, 2021, pp. 1-20,

https://stemeducationjournal.springeropen.com/articles/10.1186/s40594-021-00300-4?utm_source

J. A. Lagos Sandoval and P. N. Zapata. “Learning styles, a correlational study in engineering students.” Institute of Electrical and Electronics Engineers, 2010, pp. 1-9, https://ieeexplore.ieee.org/document/5632894/authors#authors.

Jones, Brett D., et al. “An Analysis of Motivation Constructs with First-Year Engineering Students: Relationships Among Expectancies, Values, Achievement, and Career Plans.” Journal of Engineering Education, 2010, https://www.researchgate.net/publication/264322729_An_Analysis_of_Motivation_Constructs_with_First-Year_Engineering_Students_Relationships_Among_Expectancies_Values_Achievement_and_Career_Plans

Lehto, Abigal and Ning Fang. “Work in Progress: A Correlation Analysis of Engagement of First-Generation College Students in Engineering.” ASEE, 2023, pp. 1-8, https://peer.asee.org/work-in-progress-a-correlation-analysis-of-engagement-of-first-generation-college-students-in-engineering.

Mirabelli, Joseph F., et al. “Stressors and Normalized Stress in Undergraduate Engineering Education Culture: Development of the Engineering Stress Culture Scale and Undergraduate Engineering Stressors Questionnaire.” International Journal of STEM Education, vol. 12, no. 19, 2025, SpringerOpen, https://stemeducationjournal.springeropen.com/articles/10.1186/s40594-025-00540-8.

Nolte, Hannah, et al. “No Time for That? An Investigation of Mindfulness and Stress in First-Year Engineering Design.” Design Science, vol. 8, 2022, Cambridge University Press, https://www.cambridge.org/core/journals/design-science/article/no-time-for-that-an-investigation-of-mindfulness-and-stress-in-firstyear-engineering-design/DAA4C26F4B576CB38B04AC457FC443FC?.

Power, Jason Richard, et al. “Self-efficacy development in undergraduate engineering education.” European Journal of Engineering Education, vol. 50, no. 1, Taylor and Francis, 2024, pp. 1-25, https://www.tandfonline.com/doi/full/10.1080/03043797.2024.2368149?utm_source.

Sperling, Jessica, et al. “Effects of a first-year undergraduate engineering design course: survey study of implications for student self-efficacy and professional skills, with focus on gender/sex and race/ethnicity.” International Journal of STEM Education, vol. 11, no. 8, Springer Open, 2024, https://stemeducationjournal.springeropen.com/articles/10.1186/s40594-024-00467-6?utm_source

Trenshaw, Kyle F., et al. “Using Self Determination Theory Principles to Promote Engineering Students’ Intrinsic Motivation to Learn.” International Journal of Engineering Education, vol. 32, no. 3, 2016,

https://publish.illinois.edu/glherman/files/2016/05/2016-IJEE-IM-Relatedness.pdf?utm_source.

Winkens, Ann-Kristin, et al. “How to teach resilience thinking in engineering education.” Sustainable and Resilience Infrastructure, vol. 9, iss. 5, Taylor and Francis, 2024, pp. 513-530, https://www.tandfonline.com/doi/full/10.1080/23789689.2024.2356492?utm_source.

Wint, Victoria and Ines Direito. “(RE)DEFINING ENGINEERS’ RESILIENCE: AN EXPLORATORY STUDY INTO HOW ENGINEERING STUDENTS UNDERSTAND RESILIENCE (RESEARCH).” University College London, 2024, https://discovery.ucl.ac.uk/id/eprint/10199235/1/RESILIENCE%20SUBMITTED%20post%20review%201206.pdf.

Xu, Jennifer, et al. “Exploring Personal Beliefs that may Influence First-year Engineering Students taking Ownership of their Learning.” Proceedings of the Canadian Engineering Education Association (CEEA), 2024, pp. 1-8, https://ojs.library.queensu.ca/index.php/PCEEA/article/view/18586.

A Discussion of the Societal Impact of AI Bias in Football Scouting
By Anthony Joung Lim

ABSTRACT

Artificial intelligence is becoming increasingly popular in football scouting due to its efficiency and objectivity; however, a machine learning model trained on historical data about football players is likely to exhibit a relative age effect (RAE) pattern that the system will continue to reproduce. In this research, we investigate whether an AI-based football scouting system could be biased against young or late developers and discuss the possible ethical, social, and economic consequences of age-based discrimination. Transfermarkt data was processed in Python using a distribution comparison and a chi-square test. A positional heatmap was built to demonstrate the relative age effect (RAE) in the dataset, revealing that the bias is higher in physically developed positions, such as center back and defensive midfield, with a Q1:Q4 ratio of 1.99. The presented analysis could help create position-specific correction coefficients for future AI models, enabling algorithms to correct historical relative age bias rather than reproducing it in new assessments. Therefore, the main threat seems to be selection and visibility bias that occurs before applying any AI technologies, as well as the use of position-specific bias measures in the future.

KEYWORDS

Talent, scouting, football, Artificial Intelligence, bias, social economics

INTRODUCTION

From health-tracking wristbands that Tigerwoods wears (Sens) and stadium cameras in the Santiago Bernabeu (Magazine) to cloud servers running machine-learning models (Dachman), artificial intelligence is no longer a futurist accessory for elite sports- it's now an active partner in identifying, training, protecting, and evaluating athletes (Seçkin et al. 1). As more and more sports leagues in the world and tech firms pour billions into sensors, tracking systems, detection quality, and predictive models (Liang et al. 1), athletics are becoming increasingly data-driven. Organizations adopt these technologies to improve player performance, reduce injuries, enhance refereeing accuracy, and maximize profits through streaming services and fan engagement. The global sports market was valued at roughly $5.7 billion in 2025 and is projected to reach roughly $23 billion by 2033 (GVR). However, the athletic pipeline is shifting into something much bigger—promise and progress aside, beyond the ethical questions of fairness, privacy, sustainability, and human judgment. Failing to address these ethics adequately can have implications that go well beyond misjudging an athlete. Inaccurate scouting programs can prevent talented athletes from accessing academies, while poor privacy protocols can put confidential health and performance information at risk. Additionally, the cost implications of using such technologies, as well as environmental ones, might pose obstacles to AI use by smaller clubs that cannot afford them. More importantly, heavy dependence on AI might impair

human judgment when it treats historically biased information as facts about an athlete's talents. The most consequential application is talent identification. In contrast to performance monitoring and injury prevention, which mostly benefit athletes who have already reached an advanced level of skill, talent identification decides whether a player gets the chance to compete at a high level in the first place. Whether the AI training model used advanced deep learning or classical machine learning algorithms, the different modes of talent identification remained 87.78% accurate (95% CI: 82.66–92.90) in predicting athlete potential across multiple team sports (Pietraszewski et al.). This means the models successfully classified or predicted athletic potential in nearly 9 out of 10 cases using datasets built from movement profiles, biometric load signatures, and performance outputs. Unlike traditional scouting, which relies on subjective criteria and personal bias, machine learning models are trained under strict discipline through large-scale statistical comparisons and longitudinal performance trends. Clubs and academies around the world use these algorithms to rank upcoming prospects based on the criteria listed above, using position-specific metrics and predicted development curves. Academic reviews suggest a future of expanded ML research aimed at talent identification across team sports, suggesting measurable gains in objective evaluation. Whereas traditional scouting methods can be prone to opinion bias and observational error, machine learning seems to offer a more objective approach to ranking young athletes. The downside of such approaches, however, is that errors, whether that be systematic or due to bias in data, can be costly.

Talent identification is important because selection decisions can shape a young athlete's future. In football academies, selection can mean better coaching, facilities, competition, and exposure, while rejection can mean fewer opportunities going forward. For many young players, football is not an extracurricular activity but the center of their childhood and identity. Young athletes who train for football for most of their teenage years may miss out on time that could have been devoted to broader academic development, career exploration, and preparation for adulthood. So, if an AI-led scouting system makes a mistake and classifies a player as having low potential, it’s not just a missed football career. The player can exit the academy system without the education, qualifications, or practical skills needed to enter adult life.

It has both social and economic implications. If a player is encouraged to specialize in football at an early age but does not become a professional player, they may experience lower social mobility and fewer opportunities for stable careers. Failure to progress through elite academy systems may have broader educational and socioeconomic consequences, particularly for athletes who have specialized at an early age — not because those jobs aren’t worth doing, but because they weren’t adequately prepared for broader career options in the professional world. Clubs also risk losing money by overlooking late-developing players who could have become valuable professional athletes. (Monsees)

Bias is a big problem in youth scouting. Late-developing players may be overlooked due to relative age bias. Players born earlier in the selection year are favored because they appear more physically mature. (Musch & Grondin, 2001; Deprez et al.) If AI models are trained on past academic data, they can reproduce these patterns rather than fix them. (Kim et al.) This

feeds into the larger issue of AI fairness: algorithms can reflect biased data on body type, style of play, socioeconomic background, or how quickly a player develops. This paper, therefore, investigates the possible disadvantages of AI-based talent identification models for late-developing youth football players and the ethical, societal, and economic consequences of these systems misclassifying athletic potential (Kim et al.; Monsees).

One of the long-standing cases of bias in youth football is the Relative Age Effect (RAE), which refers to the occurrence in which young athletes born at the beginning of an annual selection cycle represent a disproportionately large share of elite development streams. Because football academies organize youth competitions according to fixed age-group cut-off dates within an annual cycle, children born at the beginning of the cycle can be up to one year older than their counterparts born at the end of the cycle. While this age gap may seem small, during childhood and adolescence it translates into significant physiological, cognitive, emotional, coordination, and athletic differences between relatively older and relatively younger players. This leads to a situation where relatively older athletes often look physically and athletically stronger than other young people in the same age group, but the reason is not their higher level of development; it is biological maturation.

The Relative Age Effect has been extensively studied in elite youth football academies worldwide and in other sports. (Cobley et al.; Helsen et al.) As has been found, there is always an over-representation of athletes born in the first quarter of an annual selection cycle relative to the birth rate in the general population, which is distributed quite evenly. It means that initially favorable conditions create a so-called Matthew Effect, whereby relatively older players receive more attention from coaches, more playing time, better opponents, and a better psychological state and greater confidence, thus further widening developmental differences throughout their youth careers. At the same time, relatively younger players experience fewer developmental opportunities, a lack of confidence, and even early dropouts from competitive sports due to perceived abilities that do not actually exist. (Delorme et al.) Despite widespread knowledge of the RAE, this bias persists, as academic recruitment processes tend to focus on immediate results rather than potential (Monsees).

Given the increasing use of AI algorithms in football talent identification (Monsees; Pietraszewski et al.), the question arises as to the extent to which well-known developmental biases are reduced or even reinforced by new technology (Kim et al.). The problem is that AI systems are trained on historical datasets generated by traditional methods of scouting and selecting young players in academy recruiting. Instead of learning how to identify true future talents, machine learning algorithms can easily pick up on the fact that features associated with relatively older players are better indicators of future success, since these young athletes have already gained many of their developmental advantages through selection bias. (Suresh & Guttag, 2021).

It is clear that such a study should be conducted to investigate whether relative-age effects exist in performance-related variables used in AI-based football talent identification systems (Pietraszewski et al.; Monsees) and to consider the implications for the ethics of

AI-assisted scouting (Kim et al.). Instead of assessing the predictive capacity of a specific machine learning model, the study seeks to determine whether performance-related variables used to assess a player's future potential correlate with an athlete's birth quarter. If this happens, the use of AI in such cases will reinforce the developmental inequalities among young players without any visible signs of bias.

The study pursues four major goals in order to answer the research question. First, there will be a review of existing literature about the use of artificial intelligence in talent identification and athlete assessment in football. Second, the study estimates the extent of relative-age bias in a football dataset from Transfermarkt. Third, the study investigates whether certain variables used to determine players' potential are linked to players' relative age. This is especially true in youth scouting programs, where performance disparities driven by temporary physical maturity could be misread as talent, causing the AI program to overlook young players. Finally, the topic to be discussed is the ethical and socioeconomic impacts of this phenomenon, including the possibility that it could perpetuate selection and visibility biases, further marginalizing relatively younger athletes before they have had an equal chance to join the academies.

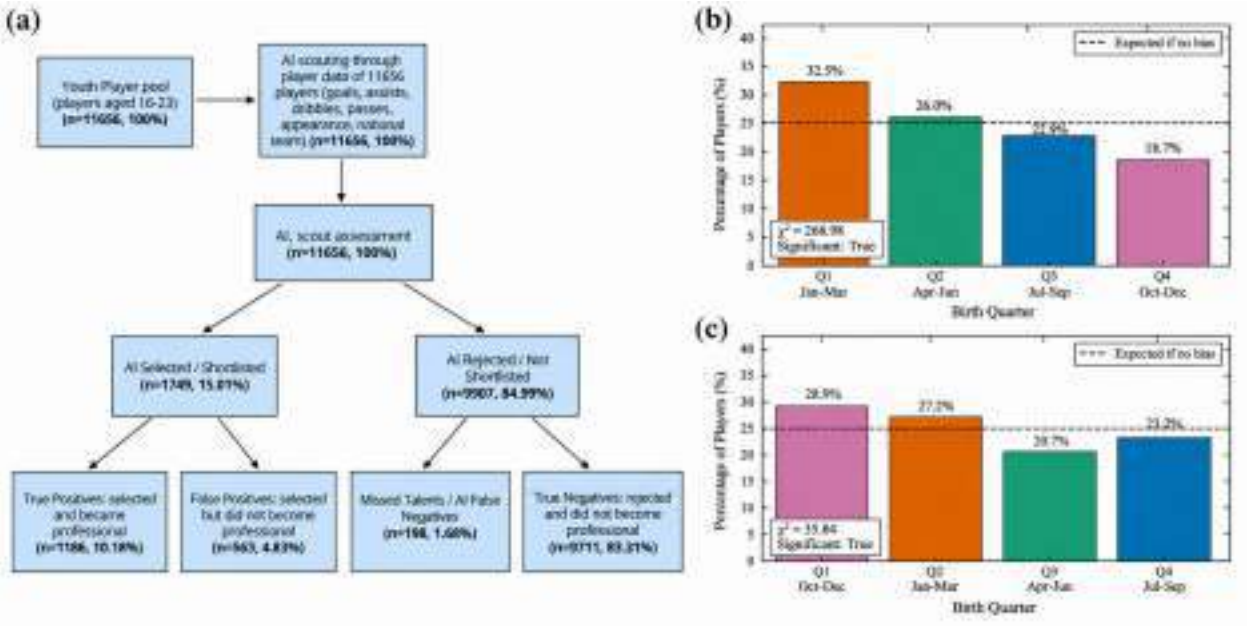

FIG 1: (a) Pipeline for identifying AI talent in the study dataset. The performance variables related to AI were evaluated for a pool of 11656 youth players. The variables were goals, assists, appearances, minutes, and market value. The model identified the top 15.01% of players as potential academy recruits or trial invitees. 84.99% of players were not selected. In the flow diagram, the results are also classified as true positives, false positives, missed talents (AI false negatives), and true negatives. This shows that the AI selection mistakes might overlook players who later become professionals. (b,c) Relative age selection bias, birth-quarter distributions. Panel (b) shows the distribution of global players according to standard calendar quarters. Panel

(c) shows the distribution of Premier League players adjusted to the British school-year system; Q1 is October–December. The dashed line is the 25% distribution expected in the absence of a bias due to the birth quarter. The panels together show how AI scouting systems could be trained on player data that already incorporates historical selection patterns and relative age effects.

METHODS AND MATERIALS

1. Data Sources and Dataset Description

In this study, we used publicly available football datasets to investigate whether AI-based scouting reproduces existing selection biases. The main dataset used was the **Transfermarkt football dataset** from Kaggle (Cariboo). The files used were players.csv, appearances.csv, and player_valuations.csv. The datasets included player information, match appearances, goals, assists, minutes played, positions, dates of birth, clubs, competitions,s and market values. The dataset used in this study contained **47,716 players**, **1,885,697 appearance records**, and **507,815 valuation records**.

To assess attacking performance, we used a second processed dataset, forward_players_2024_2025_with_birth_dates, to analyze performance in the 2024/25 season. The dataset contains variables such as name, squad, competition, age, position, date of birth, minutes played, goals, assists, expected goals, and expected assists. It was obtained from publicly available player data collected from Transfermarkt and processed using OpenAI's ChatGPT. The use of ChatGPT involved filtering the original data, categorizing players according to specific variables, and creating the final CSV file, "Forward Players 2024-2025 with Birth Dates" (OpenAI). Note: After adding the birth dates using ChatGPT, we randomly selected 20 individual players and cross-checked with Wikipedia to confirm the birth dates were correct.

This study did not directly involve any human participants. All data were secondary and publicly available;le, thus, informed consent and IRB approval were not needed. The results were presented mainly as aggregated by birth quarter, age group, cohort, and position.

2. Data Processing

The raw CSV data files were imported into Python using `pandas`. Birth dates were converted to date format and birth months extracted. The birth quarters in the global analysis were coded as: **Q1 = Jan–Mar, Q2 = Apr–Jun, Q3 = Jul–Sep, Q4 = Oct–Dec**. For the analysis of the Premier League, the quarters were adjusted to the British school-year system: Q1 = Oct–Dec; Q2 = Jan–Mar; Q3 = Apr–Jun; Q4 = Jul–Sep.

Players were filtered into a youth-player pool using the age **range 16-23** to capture players transitioning from academy football to senior professional football. Appearance data were grouped by `player_id` to calculate total minutes, goals, assists,s and appearances. Market value was aggregated by `player_id` to calculate the latest and maximum market value. Missing performance values were imputed with zeros where appropriate.

3. Data analysis

The following figures help identify the stage at which relative age bias emerges in the talent identification pathway. First, the flowchart was drawn to illustrate how a pool of youth players might be filtered via an AI-supported scouting system. In turn, it served as an illustrative example of how players' data can be divided into shortlisted and non-shortlisted categories and how false negatives can emerge when selecting potentially successful players.

The bar charts comparing the actual and expected (25%) birth-quarter distributions were also created to test whether relative age bias exists in the chosen population. At the same time, performance indicators such as goals, assists, expected goals, expected assists, and minutes played were compared between birth quarters using box plots. However, the performance figures did not indicate any advantage for early-born players. In turn, the distributions overlapped across all four quarters, and no relative age bias was apparent in players' performance figures. Therefore, since no birth-quarter advantage has been found, it is necessary to explain why relative age bias disappears among older or already selected players. The purpose was fulfilled using Figure 3, which plots the Q1:Q4 selection ratio across age and birth cohort groups. It can be observed that relative age bias is prominent only among younger players, especially those aged 17-19, and fades with age. Thus, it seems that relative age bias operates mainly during the selection and visibility stages, rather than stemming from performance differences among players who have reached professional football.

4. Stats

I used a **chi-square goodness-of-fit test** to determine whether the distribution of birth quarters differed from an equal distribution of 25% per quarter. The test had **3 degrees of freedom** and four birth quarters. The critical value for $p < 0.05$ was 7**.815**.
The **Q1:Q4 selection ratio** was also used as a measure of relative age bias:

$$Q1{:}\,Q4\ Ratio = \frac{Number\ of\ Q1\ Players}{Number\ of\ Q4\ Players}$$

A ratio greater than 1.0 indicated an overrepresentation of early-born players.
To further study the relative age effect, the Q1:Q4 ratio was calculated for each age group, birth group, and playing position. This was done to determine the reason for the reduction in the relative age effect in the performance measures once players were included in the professional football data. Box plots were constructed to compare performance measures, such as number of goals, assists, expected goals (xG), expected assists (xAG), and minutes played, across quarters of birth. Nevertheless, since there was overlap between the performance distributions and no advantage for the Q1 players, selection and visibility biases, but not the performance one, were considered.

The analysis was done in Python, in a Jupyter Notebook. Data sets were loaded, cleaned,

merged, and grouped using the `pandas` package. Calculations, z-scores, log transforms, bootstrapping,g and permutation testing were performed using `NumPy`. All figures have been generated and exported as PNG and SVG files using `Matplotlib`.
We used ChatGPT as a coding assistant to help write, debug, and format Python code. It was also used to assist in interpreting errors, arranging figures, and revising the written methodology. All final outputs were verified by executing the code in the notebook and comparing the printed output with the generated figures.

DATA

We generated all quantitative analyses and figures in Python. We cleaned, merged, and grouped the data by birth quarter, competition, playing position, age, birth cohort, and player performance.

First, I counted the number of players in each birth quarter in Python. The number of players in each birth quarter was divided by the total number of players and multiplied by 100 to obtain the percentage for each birth quarter. These percentages were then visualized in bar charts. We performed separate calculations for the overall dataset and the Premier League sample to compare the relative-age effect between the general population of footballers and a single domestic elite competition.

We also calculated performance measures for football players by birth quarter. Performance measures considered in this project include number of appearances, goals, assists, expected goals (xG), expected assists (xAG), dribbles, and minutes. Box plots were used to compare the distributions of performance measures across birth-quarter groups. Because limited playing time can skew performance comparisons between birth-quarter groups, players with fewer than 500 minutes were excluded from consideration.

Performance analysis found no evidence of advantages for early-born players over their late-born peers. While early-born players made up a larger share of the dataset, the distributions of performance measures across birth-quarter groups showed substantial overlap, and the medians for some later-born groups were even higher. Consequently, performance analysis did not address relative-age effect as an advantage for early-born players' performance. In contrast, this project examined selection and visibility biases: that is, early-born players' preference for getting into the professional datasets at all.

To study why relative age bias does not manifest in performance statistics, the Q1:Q4 selection ratio across age groups and birth cohorts was calculated. This ratio was equal to the number of early-born players divided by the number of later-born players. Then, a line graph was created from this ratio to examine how relative age bias changes with age. This way, it was possible to study whether relative age bias is highest when transitioning from youth to professional teams, and how it weakens as players grow older.

A separate football pitch visualization was developed to analyze relative age bias across playing positions. For each playing position, the Q1:Q4 selection ratio was calculated and shown on a pitch diagram. Thus, it was possible to see which playing positions have higher relative age

bias, particularly those where physical maturity can be important.

During the coding process, ChatGPT was used as a coding assistant. I described the structure of the datasets, the variables I was interested in, and the types of visualizations I needed. ChatGPT suggested and revised the Python code needed to filter the players, group observations, calculate percentages and averages, and create figures. Then, the code was run in Python; its results were compared with the printed outputs, and the code was manually revised to address errors, incorrect labels, filtering criteria, or graph layouts.

RESULTS

Birth Quarter	N Players	Goals	Assists	Expected Goals (xG)	Expected Assists (xAG)	Minutes Played
Q1	186	3	2	3.45	2	1380
Q2	151	4	2	4.1	2.2	1542
Q3	127	4	2	4.1	2.1	1603
Q4	113	3	2	3.3	1.9	1375

Table 1: Birth-quarter distribution and attacking performance indicators.

The table shows the number of players and their median performance in goals scored, assists, expected goals, expected assists, and minutes played per quarter. Birth Quarter 1 (BQ1) is made up of the most players, while Birth Quarter 4 (BQ4) is made up of the fewest players, meaning that there are too many early-born players in the database. However, the performance indicators do not show any significant superiority for players from BQ1, as players from BQ2 and BQ3 show equal or better performance in goals, expected goals, and minutes played.

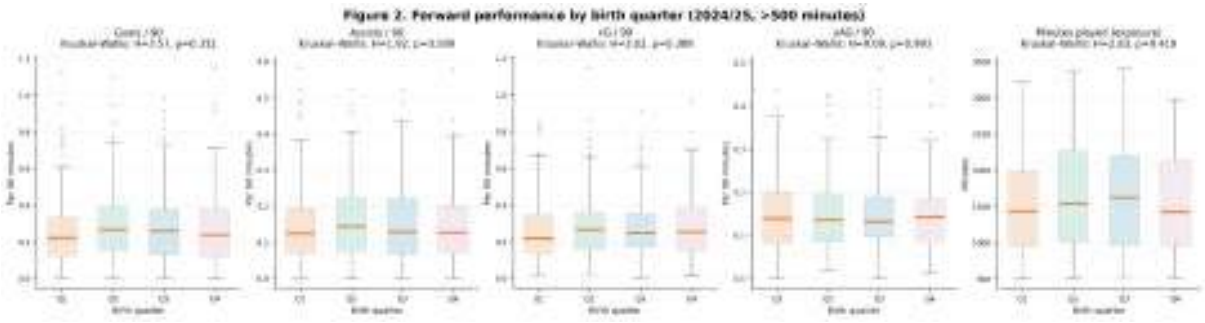

FIG 2: Forward performance metrics per 90 minutes by birth quarter.

Box plots are used to present data on goals, assists, expected goals (xG), and expected assists (xAG) per 90 minutes for forwards with more than 500 minutes of play in the 2024/25 season. The statistics were standardized by playing time, while minutes played were reported separately. Using the Kruskal-Wallis H test, no significant differences between birth quarters

were found for goals per 90 (p=0.312), assists per 90 (p=0.589), xG per 90 (p=0.389), or xAG per 90 (p=0.993). While early-born individuals continue to be over-represented in the sample, they show no superior performance in any metric, even accounting for playing time. Thus, relative age bias does not seem to influence professional attacking performance.

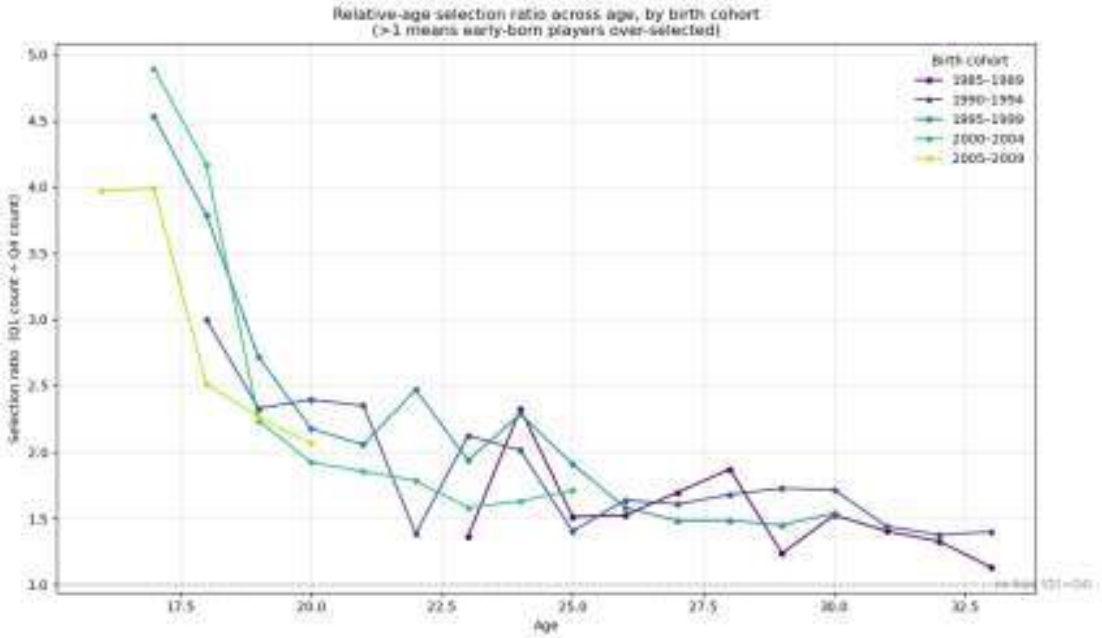

FIG 3: Relative-age selection ratio across age and birth cohort.

The line chart shows the selection ratio for Q1:Q4 across different age groups and birth cohorts, with values above 1.0 indicating overrepresentation of players born earlier. The selection ratio is maximized among young players and, more specifically, in the age range between 17 and 19, thereby suggesting that relative age bias becomes significant when there is a movement from the junior level to the professional level. With increasing age, the selection ratio tends towards 1.0, reflecting that the relative age bias decreases. (Kelly et al.).

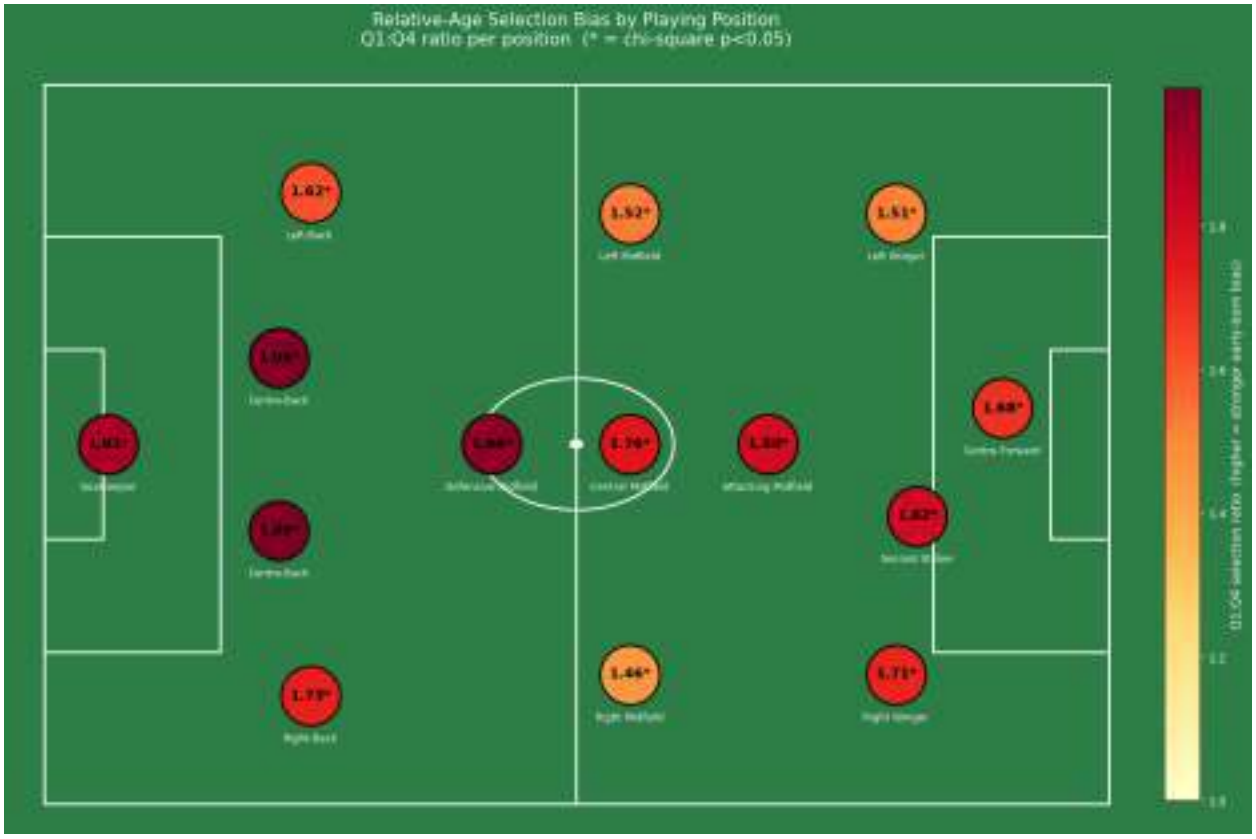

FIG 4: Relative-age selection bias by playing position.

The results first provide evidence of **relative age-selection bias** in the dataset. The number of players declines across the quarters, from **186 in Q1** to **113 in Q4** (Table 1). The ratio of Q1 is approximately **1.65**, which implies that **about 65% more players were born in Q1 relative to Q4**. For instance, Q2 and Q3 players had the same median number of goals scored (4) compared to Q1 players (3). The median expected goals (xG) of Q2 and Q3 players was also higher than that of Q1 players (4.10 versus 3.45). Q2 and Q3 players also played more minutes (1,542 and 1,603) than Q1 players (1,380). However, even after adjusting for minutes played, Q2 and Q3 players had higher median goals per 90 minutes than Q1 players (0.268 and 0.259 versus 0.223). However, according to the results of Kruskal-Wallis tests, no statistically significant differences exist in goals per 90 ($p = 0.312$) and xG per 90 ($p = 0.389$).

Figure 2 expands on this by showing goals scored, assists, expected goals, expected assists, and minutes played by birth quarter for forwards with over 500 minutes in the 2024/25 season. As shown by the box plots, the distributions overlap across all four quarters. This is an important conclusion, as it suggests that relative age bias does not manifest in performance differences when players are part of the professional sample.

Figure 3 shows how the Relative Age Effect changes across ages and birth cohorts and gives an idea of how early selection advantages affect players' careers. In particular, the selection ratio for Q1 to Q4 players is very high at ages 17-19, reaching almost 5.0 in some birth cohorts. At the same time, it drops sharply from about age 18 to age 21, falling from over 4.0 to about

2.0. After that, the decrease becomes more gradual, with the ratio varying mostly between 1.3 and 1.8 among players aged 27-33.

This trend is consistent across several birth cohorts; however, the magnitude and the age of the most pronounced decrease vary. The drastic initial decrease shows that the structure of the professional players' population is changing significantly when passing from youth to senior football. Interestingly, the ratio does not reach 1.0, even among the oldest players, which means that relatively older players usually remain overrepresented regardless of the significant decrease in bias. This suggests that early selection biases may leave lasting effects on professional players' representation, even as they lose significance with age.

The first possible reason for the overrepresentation of Q1 players is the survivor effect. Professional football datasets include information only about the players who have successfully gone through all the stages of selection, but not about all the individuals who aimed to play professional football. It means that many rejected players remain invisible. Consequently, the lack of consistent advantage in terms of performance of the selected players cannot be considered proof of the fairness of the selection process.

The possible reason for this decline is the decreasing role of physical advantages due to age increase. Players who benefited from physical advantages in childhood and adolescence may start losing their edge because younger competitors catch up. However, relative age and biological maturation are separate factors, and they cannot be used interchangeably (Hill et al.). In turn, Doncaster et al. revealed that the relative age effects were visible in Premier League youth squads, but not significant in the senior sample.

A second possible explanation for relative age effects is the underdog hypothesis. Relatively younger players who overcome selection disadvantages develop other skills (technical, tactical, psychological) to keep competing with older players. Bolckmans et al. also stated that relatively younger players in their longitudinal sample are more likely to reach senior professional status. Nevertheless, this study cannot confirm the underdog hypothesis because it did not examine psychological characteristics, biological maturation, or academy rejection history.

However, different cohort sizes and the number of active players at different ages could also affect the ratios. Because the figure reflects player representation by age rather than each player's career progression, it is impossible to say who was released from teams and why.

Figure 4 shows the bias of the relative age effect across playing positions, measured by the Q1 selection ratio. All positions under consideration have ratios greater than 1.0, suggesting that relatively older-born players prevail on the playing field. However, the strength of the bias varies across positions, indicating that the Relative Age Effect (RAE) is position-dependent rather than uniformly distributed. (Schorer et al.) The highest ratios were observed for positions such as center-back, defensive midfielder, goalkeeper, and center-forward, where features such as physical development, height, strength, confidence, and tactical skills may be considered valuable at a young age during the player selection process. For example, the center-back and defensive midfielder positions have ratios of 1.9 or higher, suggesting that relatively older

players are represented nearly twice as often as relatively younger players. Methodologically, this means an AI-based scouting system cannot use the same evaluation standard for all positions (Monsees; Pietraszewski et al.). On the contrary, this information can help develop specific correction factors to account for the RAE in specific playing positions. In cases where there is an observable overrepresentation of relatively older players, AI-based evaluation can compensate for the disadvantage of temporarily higher physical development and treat it as a manifestation of skill rather than talent. It does not, however, mean the acceptance of relatively young players as a whole; it provides additional context for interpreting the players' performance data. Thus, Figure 4 shows that combating AI scouting bias goes far beyond simple fairness and requires accounting for interactions among birth quarter, age, and playing position (Kim et al.).

These results have important implications for football clubs and AI-based talent identification. If an AI system is trained only on professional players, it will learn from a biased sample based on previous selection outcomes. Thus, clubs may repeatedly select the same types of players who benefited from early exposure to the sport while overlooking younger or late-developing players who did not get enough chances to produce performance data and are therefore unknown. Some of these unknown players might become very valuable professional players, thus causing economic and sporting losses for the clubs.

Therefore, AI scouting systems should account for birth quarter, age, biological maturation, and playing position, while also including information about rejected players.

Overall, the results suggest that **relative age bias is largely a selection and visibility issue and not a clear performance advantage.**

DISCUSSION

Figure 5. How Relative Age Bias Can Affect AI Scouting

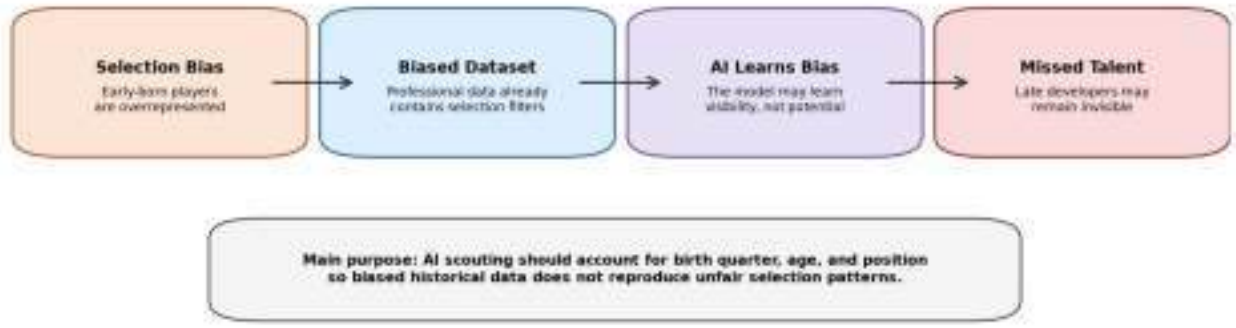

Figure 5. Conceptual model of how relative age bias can affect AI scouting.

As the figure displays, there is a chance that initial selection bias will be present in football data and thus influence AI-based scouting systems. The reason for this problem is that a

high percentage of players in the professional data were born in the first quarter. In turn, when an AI system is trained using such data, it learns visibility biases rather than the true potential of those players. Thus, late-born players are likely to go unnoticed.

This is indicative of a relative age effect in the process of football talent identification, particularly at the levels of selection and exposure, rather than at the performance level. The distribution of the birth quarter shows an excess of people born earlier than the expected 25% in the selection year, indicating that players born earlier are more likely to appear in professional football databases. However, this does not appear to confer any performance advantage to early-born players, as evidenced by indicators such as goals, assists,

xG, xAG, and minutes. The latter show no difference between birth-quarter categories, and sometimes even players born in Q2 and Q3 perform better than those born in Q1. Thus, the problem does not appear to be linked to the general talent of players born earlier, but rather to how they are selected, developed, and exposed. As for the social consequences of the problem, it should be noted that visibility-related bias is being transferred through artificial intelligence (AI) scouting systems trained on professional football databases. Thus, late-developing players who are rejected from further participation early on may lack the data required for objective evaluation and, as a result, be overlooked by the algorithm. This is true not only for football, but for many other activities for young athletes.

The findings are consistent with existing literature on the Relative Age Effect in football. Helsen, Van Winckel, and Williams showed that among European youth soccer players, those born earlier in the selection year are overrepresented, suggesting that birth timing influences the likelihood of selection for elite development (Helsen et al.). Kelly and colleagues found significant birth-quartile deviations in Premier League U18 and U21 squads, whereas senior squads show more balanced distributions, indicating that relative age bias is most pronounced during the youth development stage and fades afterward. (Kelly et al.) The present study confirms this hypothesis. Figure 3 shows why the bias seems to disappear in performance statistics: the Q1:Q4 ratio is highest among younger players, especially during the transition from youth to professional status, and decreases with age. This means that relatively older players receive certain advantages during the selection process, and only the best relatively younger players survive to enter the pool of professional players. Thus, when performance metrics are evaluated, the sample has already been filtered. Therefore, these findings confirm that AI scouting systems should not be judged solely by their accuracy in evaluating already visible players, but also by whether the evaluation dataset excluded certain individuals from the selection process.

The above findings are illustrated in Figure 5, which shows how a bias in relative age can extend from youth selection to AI-powered scouting systems. In this figure, the pathway starts with the overrepresentation of early-born individuals due to selection bias, thereby creating a professional database pre-filtered by historical selection. When the AI system is trained on such databases, it will learn more about visibility than about those players' potential. Therefore, late-born or late-maturing players will not be available to the algorithm before a fair analysis.

The goal of Figure 5 is to connect the findings of this study to the ethical question of using AI in scouting: although the algorithm itself is objective, the database it uses can be biased. (Suresh & Guttag, 2021; Kim et al.)

There are several limitations to the study. First, the database used consisted mainly of players who had made it into professional football databases, which means it was impossible to directly measure the most important missing group of players – those who were released, overlooked, or had never appeared in elite football datasets. Second, the list of performance metrics was limited to publicly available football statistics, including goals, assists, minutes, xG, xAG, appearances, and market value. These variables do not cover certain factors such as psychological development, tactical skills, coaching quality, injury history, family background, physical maturation, and educational outcomes. Third, the analysis relied solely on publicly available information rather than private data from academies or scouts. Future research should consider longitudinal data collected at the academy level, include players who did not become professionals, and follow both their football and educational careers after they leave the academy. Such research could test whether late-born or late-developing players are actually being excluded from the selection system before they have fully developed their talents.

Also, it is important to mention the technical secrecy around AI-based scouting systems used in Premier League clubs. They can purchase information about players provided by commercial suppliers like Hudl, StatsBomb, and Wyscout. For instance, StatsBomb collects 3,400 events per match, including pressures,

passes, and positions (Hudl). However, the competitive advantage of Premier League clubs does not lie solely in the ability to obtain this data, but in having their algorithms analyze statistics differently. For example, Brighton & Hove Albion relies on its exclusive dataset and unique approach to uncover undervalued players (Anderson). Also, Brentford incorporates proprietary data and specific requirements for position-related scouting into its recruitment strategy (Goodwin). Because other teams can find the same transfer targets if the details of the approach become available, protecting scouting algorithms and evaluations has major commercial significance. Technical measures such as database access control, encryption, and confidentiality agreements may be used to safeguard this kind of intellectual property. Nonetheless, there is also an ethical issue related to this problem since the lack of transparency does not allow independent researchers to check whether these models create a relative age effect and exclude late developers.

This aspect concerns the evaluation of the performance and commercial efficiency of these systems as well. For instance, a good practice is the assessment of data-driven recruitment carried out by human evaluators at Brentford. According to the club's technical director, Lee Dykes, the scouting process starts by reducing approximately 85,500 players to 16 positional categories, followed by assessment by scouts, analysts, coaches, and other decision-makers (Goodwin). It is also worth noting that, according to the club owner, Matthew Benham, the model did not approve the recruitment of Bryan Mbeumo due to a lack of data; however, positive reviews from human scouts led to his signing (BBC Sport). From the commercial perspective, it

is known that data-driven recruitment led to Brighton's success in recruiting and selling players. It earned the club £122.8 million in after-tax profit in the 2022/23 financial year, partly due to significant transfer income (Brighton & Hove Albion). However, this result does not prove the independent effectiveness of these algorithms, because transfer income also depends on coaches' work and the club's negotiating skills. Moreover, clubs usually do not disclose the predictive accuracy and fairness of their proprietary scouting models. Therefore, future research should consider not only financial gains and successful transfers, but also prediction errors and missed talents.

CONCLUSION

In summary, the main purpose of the current study was to investigate whether the use of artificial intelligence in football scouting would replicate preexisting bias, using RAE as an example. As a result, the research demonstrated an overrepresentation of early-born players regardless of performance differences in the dataset and across various positions. Thus, relative age bias appears to function primarily through selection and visibility, and not differences in long-term footballing abilities. However, there are additional risks of introducing this kind of bias due to the use of AI. In particular, **aggregation bias** can arise if a single model evaluates players across different positions or development phases, despite differences in performance-potential relationships. While developing position-specific models could solve this issue, it might create new inequalities. Moreover, **deployment bias** can emerge if a prediction designed for scouting is taken as sufficient proof for either recruiting or rejecting players (Suresh and Guttag). Thus, an AI system trained on historically selected footballers will continue to maintain existing inequality, even if its predictions are statistically objective

Beyond the possibility of becoming professional footballers, individual players face numerous societal consequences. First, athletes overlooked by academies will not have access to professional training, educational support, mentoring, and job opportunities. Players who invest many years of their adolescence in developing footballing skills may struggle to find another way to make a living after leaving the sport because of limited academic and vocational education. Furthermore, rejection by the academy can affect young people's psychological well-being, identity, and confidence because of their connection to football (Scandle et al.). Physical development also represents another aspect where individual differences might play a role since players of the same chronological age can demonstrate substantial variation in biological maturation. Different growth patterns, physical demands, and training loads might influence injury profiles, suggesting that academy selection and development should be based not only on chronological age but also biological maturation (Wik).

These inequalities might be amplified by socioeconomic, geographical, ethnic, and nationality disadvantages. In particular, children from richer families may have more opportunities for private coaching, specialized training, transportation, and overall development than poor athletes, despite similar potential. For instance, Kelly et al. note the role of socioeconomic background and psychological characteristics in academy talent development.

Geographical factors may also introduce inequality

because players living near academies and development centers have more opportunities to be scouted than those living far from them. Finnegan et al. demonstrated that Irish youth players born in counties with an elite development center were approximately 50% more likely to get selected compared to those who were born in counties without one. The intersection of ethnicity and nationality can create additional inequality if certain populations are disproportionately represented in the recruitment network or the historical dataset. Although this study does not directly consider socioeconomic, ethnic, and nationality bias, these aspects should be addressed when assessing the fairness of AI scouting systems. If these inequalities go unaddressed, AI might amplify differences in social mobility by advancing players with more opportunities to develop and be noticed.

However, **AI scouting also has economic implications for clubs and the transfer market.** Overselecting early-maturing players whose initial advantages do not translate into performance would mean inefficient spending on academy players, including coaches, facilities, training, and development costs. At the same time, rejection of late-maturing players would mean a missed sporting and financial opportunity. If recruitment algorithms favor highly visible players because of greater historical exposure than equally talented individuals from less recognized pathways, this could create inefficiencies in the transfer market and distort transfer value, even though this effect was not measured in the current study. For smaller clubs, this would mean an opportunity to discover talented players at relatively low transfer cost and compete with richer clubs. However, such opportunities reflect a more serious issue: the transfer market would reward clubs that can recognize talent overlooked by other scouting systems. Thus, improving fairness could benefit both players and clubs.

One **solution is to** develop maturity-adjusted, position-specific scouting systems that combine AI and human assessment. By estimating biological maturation instead of chronological age, clubs can assess players' maturity through repeated measurements of height, growth velocity, and proximity to peak height velocity, while accounting for uncertainties in these estimates. Psychological characteristics, such as adaptability, decision-making, resilience, and responses to competitive pressure, can also be evaluated using validated instruments and longitudinal observations by experts. These characteristics should inform player development, not serve as a permanent label that defines their future. AI models can then evaluate performance in relation to birth quarter, biological maturation, position, and available developmental opportunities. Most importantly, **AI should complement rather than replace human scouts**, whose observations can catch what performance figures miss. However, human assessments should also be structured and independently reviewed in order to prevent biased judgment from replicating the same inequality.

Some **policy changes at the league level** could make these solutions feasible and consistent. Football organizations can mandate the review of selection and rejection rates of players by birth quarter, age, position, and other relevant characteristics, while protecting players' sensitive information. They can also promote the practice of periodic re-evaluation of rejected

players, geographically accessible trials, and special support for disadvantaged areas. In particular, the Elite Player Performance Plan of the Premier League includes biological maturation screening and bio-banding, grouping players by developmental stage instead of chronological age (Premier League). These practices could form the foundation for fairness norms in professional leagues. Independent audits of AI scouting systems could check prediction accuracy, false-negative rates, and unequal selection. Education, career counseling, and psychological support should also remain available for players rejected by the academy system. The **goal of using AI in football scouting should not be to eliminate human judgment, but rather to combine technological efficiency with developmental expertise and institutional responsibility. Thus, the measure of success of AI scouting should not only be its ability to find talent among already known players but also its capacity to discover previously unnoticed talent.

DATA AVAILABILITY

The processed dataset used to analyze birth-quarter distributions and attacking performance in the 2024/25 season is publicly available on Kaggle: https://www.kaggle.com/datasets/anthonyjounglim/forward-players-2024-2025-with-birth-dates The original Transfermarket football dataset used in this study is also publicly available on Kaggle

ACKNOWLEDGEMENTS (Optional)

I would like to express my sincere gratitude to my mentor, Kishwar Iqbal, for their guidance, supervision, and continued support throughout this project, as well as for the many valuable discussions that helped shape this research.

Works Cited

1. Anderson, Barry. “Inside Tony Bloom’s Algorithms: How It Works & What Hearts Can Expect from Brighton’s Maths Genius.” Edinburgh News, 17 Sept. 2024, https://www.edinburghnews.scotsman.com/sport/football/hearts/inside-tony-blooms-algorithms-how-it-works-what-hearts-can-expect-from-brightons-maths-genius-4786037. 2. “Athlete Monitoring System.” Catapult, Catapult Sports,

2.www.catapult.com/solutions/athlete-monitoring. Accessed 16 Dec. 2025. 3. Baer, Benjamin. “MLS Partners with Second Spectrum on Advanced Tracking Data System.” MLSsoccer.com, Major League Soccer, 26 Feb. 2020,

3.www.mlssoccer.com/news/mls-partners-second-spectrum-advanced-tracking-data-system . Accessed 16 Dec. 2025.

4. Brighton & Hove Albion. “Albion’s Annual Accounts.” Brighton & Hove Albion FC, 2 Apr. 2024,
https://www.brightonandhovealbion.com/media-article/Albion%27s-annual-accounts.

5. Cariboo, David. “Football Data from Transfermarkt.” Kaggle, https://www.kaggle.com/datasets/davidcariboo/player-scores. Accessed 7 Sept. 2026. 6. Cobley, Stephen, Joseph Baker, Nick Wattie, and Jim McKenna. “Annual Age-Grouping and Athlete Development: A Meta-Analytical Review of Relative Age Effects in Sport.” Sports Medicine, vol. 39, no. 3, 2009, pp. 235–256. https://doi.org/10.2165/00007256-200939030-00005.

7. Delorme, Nicolas, Julie Boiché, and Michel Raspaud. “Relative Age and Dropout in French Male Soccer.” Journal of Sports Sciences, vol. 28, no. 7, 2010, pp. 717–722. https://doi.org/10.1080/02640411003663276.

8. Doncaster, Greg, et al. “Relative Age Effects and the Premier League’s Elite Player Performance Plan (EPPP): A Comparison of Birthdate Distributions Within and Between Age Groups.” Journal of Science in Sport and Exercise, vol. 8, 2026, pp. 207–216. https://doi.org/10.1007/s42978-024-00285-w.

9. Dubow, Josh. “NFL Uses AI to Predict Injuries, Aiming to Keep Players Healthier.” AP News, Associated Press, 14 Oct. 2025, apnews.com/article/nfl-ai-injury-prevention-c345d5f16205379e2029057cd97ca08f. Accessed 17 Dec. 2025.

10. Feinberg, Doug. “WNBA to Use Optical Tracking to Enhance Player Analysis.” AP News, Associated Press, 6 Mar. 2024, apnews.com/article/wnba-3d-tracking-c7c7c51c94fde0f4ed6cb2d12eef84e8. Accessed 16 Dec. 2025.

11. https://chatgpt.com3. Goodwin, Adam. “Lee Dykes: Brentford’s Seven Stages of Recruitment.” Brentford FC, 18 Dec. 2023, https://www.brentfordfc.com/en/news/the-long-read-brentford-technical-director-lee-dykes-interview-transfer-window-moneyball./c/6a353870-c3c0-83ee-81f5-ae5265f539ab

12. GVR Next Generation Technologies Research Team. “Sports Analytics Market Size And

Share Report, 2026-2033." GrandViewInc, Grand View Research, 22 June 2026, https://www.grandviewresearch.com/industry-analysis/sports-analytics-market.
13. "Have Questions about Hudl Statsbomb? We Have the Answers." Hudl, https://www.hudl.com/products/statsbomb/faq. Accessed 23 Oct. 2026.
14. Helsen, Werner F., Jan Van Winckel, and A. Mark Williams. "The Relative Age Effect in Youth Soccer across Europe." Journal of Sports Sciences, vol. 23, no. 6, 2005, pp. 629–636. https://doi.org/10.1080/02640410400021310.
15. "How Brentford's Recruitment Model Works." BBC Sport, 11 Mar. 2026, https://www.bbc.co.uk/sport/football/articles/cx2lwk275v3o.
16. Raffell, Joseph. "UK School Holidays and School Days Explained." Engage Education, 7 July 2026, https://engage-education.com/blog/uk-school-holidays-and-school-days-explained-a-guide-for-relocators/.
17. Kim, Jae-Hak, Janghyeon Kim, Hoseok Kang, and Bo-Young Youn. "Ethical Implications of Artificial Intelligence in Sport: A Systematic Scoping Review." Journal of Sport and Health Science, vol. 14, 2025, p. 101047. https://doi.org/10.1016/j.jshs.2025.101047.
18. Lovell, Ric, et al. "Soccer Player Characteristics in English Lower-League Development Programmes: The Relationships between Relative Age, Maturation, Anthropometry and Physical Fitness." PLOS ONE, vol. 10, no. 9, 2015, e0137238. https://doi.org/10.1371/journal.pone.0137238.
19. "MLS Taps Second Spectrum for AI-Powered Data-Tracking." SportsPro, 26 Feb. 2020, www.sportspro.com/news/mls-second-spectrum-artificial-intelligence-player-data-tracking/. Accessed 17 Dec. 2025.
20. Monsees, Leah M. "'There Is a Lot More Potential': Practitioner Perspectives on Technology and Data-Driven Talent Identification, Selection, and Development in a German Bundesliga Academy." International Journal of Sports Science & Coaching, vol. 20, no. 2, 2025, pp. 628–638. https://doi.org/10.1177/17479541241308519.
21. Musch, Jochen, and Simon Grondin. "Unequal Competition as an Impediment to Personal Development: A Review of the Relative Age Effect in Sport." Developmental Review, vol. 21, no. 2, 2001, pp. 147–167. https://doi.org/10.1006/drev.2000.0516.
22. Ocak, Mustafa Furkan, et al. "Validity and Reliability of Catapult Vector S7 GNSS Units in Distance Measurement." BMC Sports Science, Medicine and Rehabilitation, vol. 17, no. 1, 2025. https://doi.org/10.1186/s13102-025-01335-8.
23. OpenAI. Forward Players 2024–2025 with Birth Dates. Generated using ChatGPT, 7 Sept. 2026. CSV file.
24. Pappalardo, Luca, et al. "PlayeRank: Data-Driven Performance Evaluation and Player Ranking in Soccer via a Machine Learning Approach." ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 5, 2019, pp. 1–27. https://doi.org/10.1145/3343172.
25. Pietraszewski, Przemysław, et al. "The Role of Artificial Intelligence in Sports Analytics: A

Systematic Review and Meta-Analysis of Performance Trends." Applied Sciences, vol. 15, no. 13, 2025, p. 7254. https://doi.org/10.3390/app15137254.
26. Radanliev, Petar. "Privacy, Ethics, Transparency, and Accountability in AI Systems for Wearable Devices." Frontiers in Digital Health, vol. 7, 2025. https://doi.org/10.3389/fdgth.2025.1431246.
27. Schorer, Jörg, Stephen Cobley, Dirk Büsch, Harald Bräutigam, and Joseph Baker. "Influences of Competition Level, Gender, Player Nationality, Career Stage and Playing Position on Relative Age Effects." Scandinavian Journal of Medicine & Science in Sports, vol. 19, no. 5, 2009, pp. 720–730. https://doi.org/10.1111/j.1600-0838.2008.00838.x.
28. Souaifi, Marouen, et al. "Artificial Intelligence in Sports Biomechanics: A Scoping Review on Wearable Technology, Motion Analysis, and Injury Prevention." Bioengineering, vol. 12, no. 8, 2025, p. 887. https://doi.org/10.3390/bioengineering12080887.
29. Suresh, Harini, and John Guttag. "A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle." Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, Association for Computing Machinery, 2021. https://doi.org/10.1145/3465416.3483305.
30. Vaeyens, Roel, Renaat M. Philippaerts, and Robert M. Malina. "The Relative Age Effect in Soccer: A Match-Related Perspective." Journal of Sports Sciences, vol. 23, no. 7, 2005, pp. 747–756. https://doi.org/10.1080/02640410400022052.

The Global Superconducting Quantum Computer Race: A Review of the Achievements and Remaining Challenges that Superconducting Quantum Processors Still Face by Comparing Competitors By Max Chen

Abstract

Despite quantum computers not being a direct replacement for the ubiquitous classical computers in our society, they can leverage the unique properties of the quantum realm to substantially enhance computational power. The incalculable potential of this industry jump-started the worldwide quantum computing race. Amongst this race, the technological platform currently gaining the most traction in the industry is superconducting quantum computers. In particular, we will discuss three prominent contenders: Google, IBM and Rigetti. This paper argues that the superconducting quantum computer industry has made steady progress in raising the scale, improving connectivity and improving the general fidelity of quantum computers; however, the immaturity of Quantum Error Correction (QEC) continues to be problematic, limiting our ability to further scale quantum computers without trade-offs. Furthermore, this article shows that the absence of a universal benchmark, the lack of accessible data on the latest quantum computers and the inconsistencies in the measurement of metrics suggest that there is no conclusive leader in this field. Ultimately, identifying which tech powerhouse is leading this industry will reveal the most vital obstacles that quantum computing currently faces, as well as the future direction that quantum computing is headed towards.

Keywords

Engineering, Quantum Computers, Superconducting Quantum Computers, Quantum Error Correction, Scalability

Introduction

In the history of humankind, there have been four industrial revolutions so far: the First Industrial Revolution, which began the adoption of machinery and steam power; the Second Industrial Revolution, primarily focused on the adoption of mass production; the Third Industrial Revolution, which is the digitisation of various aspects of life; and the Fourth Industrial Revolution, which is primarily about the maturation of artificial intelligence. (Groumpos) Recently, the field of quantum computing has been at the forefront of the IT industry and its superior capabilities, in comparison to its classical counterpart, have led experts to hypothesise the possibility of a "5th industrial revolution"(Rietsche et al.). Despite the imperfection of current quantum computers, they have endless potential applications, such as creating new materials, discovering new drugs (Santagati et al.), enhancing the efficiency of supply chains (Phillipson), aiding scientific research (AbuGhanem and Eleuch) and cultivating financial modelling. (Orus et al.) Once quantum computers are able to consistently and maturely conduct these applications, it will be a monumental transformation for the world. For instance, Domingo et al. strongly suggest that quantum computers coupled with machine learning will have 'the

ability to predict the binding affinity between a potential drug and its target protein' with 40% time and cost savings plus 20% less complexity compared to traditional methods. (Domingo et al.) Even though this practice is not currently standardised in the pharmaceutical industry, it is an irrefutable indicator of the pivotal effects quantum computing can have after evolving over time. Thus, the aim of this article is to determine the achievements and remaining challenges that superconducting quantum computers/processors still face by comparing the leading competitors in the field of quantum computing: Google, IBM and Rigetti. This will allow us to see how close we are to a fault-tolerant quantum computing era and highlight the current most dire issues at hand that are hindering the pervasive adoption of quantum computers into our society.

There are various types of technological platforms of quantum computers, which usually consist of superconducting qubits(Ladd et al.; Wu et al.; Arute et al.; AbuGhanem, "Superconducting Quantum Computers"), photonic qubits(Ladd et al.; Zhong et al.), trapped-ion qubits(Zhang et al.; Debnath et al.), nuclear magnetic resonance qubits(Ladd et al.; Czincila et al.) and cold atom qubits(Acuaviva et al.; Bernien et al.; Ebadi et al.) . This paper solely focuses on analysing superconducting quantum computers due to the fact that superconducting quantum computers are the most popular type. (Rossi et al.) Additionally, by only utilising superconducting computers, the comparison process becomes substantially more feasible.

Furthermore, the justification for selecting the companies Google, IBM and Rigetti is due to the widely accessible data on their quantum computers, as well as their well-known names, indicating their success in the field of quantum computing. Google, IBM and Rigetti are also three companies in the quantum computing field with the most experience (Google Quantum AI division founded in 2012("We're Scaling Quantum Computing Even Faster with Atlantic Quantum."), IBM built superconducting qubits in the early 2000s(Five Years Ago Today, We Put the First Quantum Computer on the Cloud. Here's How We Did It. | IBM Quantum Computing Blog), and Rigetti was founded in 2013 (Rigetti et al.), which signifies that their quantum computers are the most advanced and refined.

Firstly, this paper will explain the fundamental concepts of quantum computing. Secondly, this paper will discuss how the quantum computers of each company will be compared. Additionally, the subsequent section will analyse each of the three companies: 1) current progress/breakthroughs and 2) Issues. Following that, there will be a summary of all the metrics from each company. Lastly, the paper will discuss the common obstacles that everyone is facing and what is required for evolution into the fault-tolerant era to conclude.

Fundamentals of Quantum Computing

The major difference between quantum computers and classical computers is how data is stored. On one hand, classical computers consist of minuscule switches that store data in the unit "bits", and these bits can either be in a state of 0 or a state of 1. Whereas quantum computers are able to leverage the principles of quantum mechanics to store and compute data units called qubits — which stands for "quantum bit". The superposition property of qubits allows information to be stored as a combination of both the 0 basis state and the 1 basis state.

For instance, whilst classical computers are either 100% in the 0 state or 100% in the 1 state, quantum computers can be 70% in the 0 state and 30% in the 1 state or 100% in the 0 state and 0% in the 1 state. These percentages represent the probability that the qubit is in either the 0 or 1 state.

Furthermore, this superpositional state is mathematically described by a state vector that consists of complex numbers in the Hilbert space. Since the components of the state vector mathematically relate to the probability of each basis state, they are also known as quantum probability amplitudes. These probability amplitudes evolve according to Schrodinger's Equation, allowing wave-like interference between probabilities. (Shirts and Welch) As a result, complex numbers are required to be able to represent both the phase and the magnitude properties.

Additionally, the state vector components do not directly translate to the percentages of the superpositional states. The percentages are equivalent to the modulus squared of the probability amplitudes in the state vector — this is known as the Born rule. (Neumaier) Consequently, a probability distribution of all the possible bitstrings that the qubit could represent is formed from squaring each state vector component (which is why each component of the state vector is also called a probability amplitude).

The number of components in the state vector corresponds to the possible number of bit strings, and the number of possible bit strings is determined by the number of qubits. For example, if a quantum computer has 2 qubits, the possible bit strings are $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$. Thus, this is the relation between the number of Qubits and the number of possible Bit Strings: if a quantum computer has K Qubits, the number of possible bit strings is 2^K. Thus, the state vector would have 2^K number of components.

To sum up, the property of superposition enables qubits to represent all the possible bitstrings simultaneously. Hence, in quantum systems the computational space can ‘increase exponentially with the size of the system’(Rieffel and Polak), ‘which enables exponential parallelism’(Rieffel and Polak), significantly exceeding the computational power of classical computers in specific tasks. (Shor; Bauer et al.; Herman et al.)

Regardless, once the data being stored in a qubit is measured/read, the probability distribution “collapses” (Shrapnel et al.) and all the probability concentrates at one basis state/bit string, so the qubit loses its quantum nature and conforms to the classical world. Mathematically, in the Hilbert space, the state vector rotates to completely align itself with the axis of one of the basis states (this property of quantum mechanics is famously modelled by Schrödinger’s Cat analogy). As a result, the whole point of quantum algorithms is to attempt to shift the state vector as close as possible to the correct axis of the basis state/bit string, so that the probability of the correct basis state is maximised. This process of “shifting” the state vector is essentially using quantum algorithms to cause interference between the probability amplitudes of the wavefunction. When waves constructively interfere, the probability amplitude increases and when the wave destructively interferes, the probability amplitude decreases. This is what allows the probability to be gradually concentrated towards 1 basis state during quantum computation.

To conclude, the main concept of a quantum computer is that qubits in superposition states are manipulated until the probability of measuring a certain basis state is maximised.

Another property that quantum computers exploit is entanglement, meaning that when "the state of one qubit is dependent on the state of another qubit"(Rietsche et al.) The benefit that entanglement provides is that it permits each qubit to coordinate in sync to reach a final output, in a way that bits in classical computers are not able to do.(Rietsche et al.) This property applies regardless of the distance between the two qubits, which is what ignited the theory of quantum teleportation. (Steane)

How to Compare the Performance of Quantum Computers?

For classical computers (and the STEM realm in general), the implementation of standard benchmarks has yielded positive results and contributed heavily to the evolution of science(Sim et al.): it elucidates a clear set of objective criteria that aid experts to compare their accomplishments and verify empirical progress. (Acuaviva et al.) Problematically, as the number of competitors in the nascent industry of quantum computing continues to grow, the absence of a universal method to compare the performance of quantum computers may trigger a manifestation of Goodhart's law. (Goodhart) Nevertheless, the massively distinct approaches to the creation of quantum computers between major players of the industry significantly complicate the comparison process. The two primary differences between the quantum computers that hinder our ability to compare quantum computers are the technological platforms/qubit hardware (as noted above) and the computational paradigms. (Acuaviva et al.) Whereas the type of quantum computing paradigm usually consists of these 3: annealing-based quantum computing (adiabatic theorem), analogue quantum computing, digital quantum computing. This is a crucial reason as to why this paper isolates one type of technological platform to compare (which is superconducting quantum computers). Moreover, the immaturity of the current field of quantum mechanics also explains why a straightforward way to pursue the innovation of quantum computers is nonexistent. Therefore, despite exclusively comparing superconducting quantum computers, the lack of a universal benchmark between different types of quantum computers still forces us to use simple metrics instead.

Benchmarks and metrics are very similar terms, but their definitions portray how a lack of benchmarks can be an issue. A benchmark is 'a test, or set of tests, that aims to measure or evaluate the performance, efficiency, or other properties of a quantum processor or hardware component for a certain task.'(Acuaviva et al.) Whereas a metric is 'a value derived from the evaluation of a function that maps a considered set of objects or events to another set of values, characterising a given performance-related property of a system.'(Acuaviva et al.) Essentially, a benchmark is intended to test the overall performance of a quantum computer, whilst the majority of metrics give a singular value of a specific aspect of the quantum computer, which could potentially be misleading. For example, a quantum computer with superior qubit connectivity but less number of qubits could potentially have better performance than a quantum computer with inferior qubit connectivity despite having more qubits. As a result, in addition to

comparing the five quantum computers with metrics, this paper will be aided by an evaluation to determine whether this comparison could be unreliable.

The metrics that will be utilised in this paper to analyse and compare the performance of quantum computers are separated into three categories: scalability metrics, quality metrics and speed metrics. (Acuaviva et al.; AbuGhanem, "IBM Quantum Computers") The metrics used in this work are summarised in Table 1.

Table 1. Metrics used to compare the superconducting quantum computers/processors in 3 categories.

Metric (Quantitative)	
Scale	Number of Qubits
Quality	Readout Fidelity (%)
	Single-Qubit Gate Fidelity (%)
	Two-Qubit Gate Fidelity (%)
	Coherence Time T1 (μs)
	Coherence Time T2 (μs)
Speed	1 Qubit Gate Time (ns)
	2 Qubit Gate Time (ns)
	Measurement Time (ns)

To summarise, these are the criteria that we will base our analysis, comparison and evaluation of quantum computers on in this paper.

Latest Quantum Computers/Processors of the Companies

In this section, we will discuss 1) the Current Progress/Breakthroughs and 2) the Issues of the latest superconducting quantum computer/processor of each company: Google's "Willow", IBM's "Nighthawk" and Rigetti's "Cepheus-1-108Q"

Google's "Willow"

Current Progress/Breakthroughs

The newly developed quantum processor produced by google, called the "Willow", has made considerable improvements in comparison to its previous model, named the "Sycamore". The Willow has reached a staggering 105 qubits in comparison to the previous upgraded version of the Sycamore that had 70 Qubits. Additionally, the coherence time of Willow has also increased by a factor of 5 relative to the Sycamore.

Google has also tested their new quantum processor using an industry-standard benchmark called Random Circuit Sampling (RCS): this benchmark essentially samples 'from the probability distributions of randomly chosen quantum circuits'(Bouland et al.) (a benchmark

invented to be able to demonstrate quantum supremacy). Impressively, the RCS result of the Willow showed that it was able to complete a computation task that takes the most advanced supercomputer 10^{25} years in less than 5 minutes. (Jin)

The Willow has also displayed a massive progression in Quantum Error Correction (QEC) through the incorporation of Surface Code, which is a 'type of topological error correction code'(Jin) that enhances the quantum computer's tolerance to local errors. Moreover, it is important to note that the Surface Code only needs nearest-neighbour connectivity between qubits, which is especially advantageous for a superconducting quantum processor due to its 2D circuit architecture where qubits are only connected to other nearby qubits. The surface code has actually existed for quite a few years, but experts have struggled with lowering the probability of error below the surface code threshold. If the error rate is above the threshold, every additional qubit would cause an increase in the error rate, which is not ideal. In 2024, Google was able to conduct quantum error correction below the surface code threshold(Google Quantum AI and Collaborators et al.), 'exponentially suppressing the logical error rate by a factor of 2.14+- 0.02' (Jin). Simply put, Google is now able to scale up the number of qubits whilst exponentially decreasing their error rate, signifying an incredibly large step towards fault-tolerant quantum computers.

To conclude, Google has substantially increased the number of qubits and coherence time T1 and T2 (which led to higher gate fidelity and readout fidelity) in their superconducting quantum processor, accomplished revolutionary results in RCS and successfully conducted a scalable QEC method via the surface code.

Issues

On the other hand, regardless of how well Google performs in the RCS benchmark, it is simply an indication of performance, unable to show how the Willow will hold up when being utilised for practical computational tasks. For example, artificial intelligence, materials searching and chemical dynamics, logistics optimisation, financial activities, decryption, etc (Jin). Secondly, Willow still struggles with error correction of the noise floor despite some evident improvements. The term 'noise floor' is the residual noise that affects the computation process of quantum systems due to high-energy events (for instance, cosmic rays(McEwen et al.)). In addition, the Willow has burst errors, which is when a failing qubit contagiously affects its neighbouring qubits(Google Quantum AI and Collaborators et al.), spanning across 30 qubits due to unknown events that Google has yet to discover. Despite the faster decay rates, Google is currently still attempting to resolve this issue. Ultimately, Google's Willow processor has three main issues: the lack of experimentation with practical computation tasks, error correction of the noise floor and burst errors.

IBM's "Nighthawk"

Current Progress/Breakthroughs

The most recent model of quantum processor that IBM has produced is the Quantum "Nighthawk". In fact, this model was just introduced in November 2025; thus, there is a lack of literature review, and most of the information obtained is from news article sources published by IBM. The Nighthawk is a superconducting quantum processor that has achieved the impressive scale of 120 qubits. In terms of connectivity, the qubits are connected by 218 next-generation tunable couplers to their four nearest neighbouring qubits in a square lattice architecture. There are two major advantages that tunable couplers provide, which is that it allows faster two-qubit gate speeds and reduces the gate errors induced from undesired 'always-on' coupling between qubits. (Stehlik et al.) Previously, the "Heron" quantum processor model only had 176 tunable couplers, but had 156 qubits. By sacrificing the number of qubits to improve connectivity, IBM's Nighthawk is able to operate circuits that are 30% more complex in comparison to IBM's previous generation. Additionally, the execution of these computations can be done accurately with high fidelity. Furthermore, the improved architecture will permit users to compute tasks with more complex entangling operations. In fact, the Nighthawk can handle computational problems that need up to 5000 two-qubit gates.

Issues

Nighthawk is evidence that IBM has continued to refine their quantum processors, but IBM still faces a few obstacles. Firstly, the goal that IBM is still unable to achieve, as well as the field of quantum computing in general, is to produce a fault-tolerant quantum computer, which is a quantum computer that is able to operate with accuracy despite the unavoidable disturbances in quantum systems. Mainly, this is because of inadequate quantum error correction that fails to deal with noise, qubit connectivity and other hardware limitations. (Preskill) Secondly, IBM still struggles with scaling its quantum processors. This is because 'As quantum systems grow in size, the complexity of managing qubits, minimising noise, and maintaining coherence increases exponentially.'(AbuGhanem, "IBM Quantum Computers") Furthermore, other factors such as temperature, precision control and error correction become a greater challenge to resolve when scaling. (AbuGhanem, "IBM Quantum Computers") Specifically for the temperature, quantum systems require an ultra-low temperature environment to conduct computation; thus heat dissipation becomes much more troublesome. Connectivity between qubits also becomes problematic when scaling up. The increase in interactions between qubits is very likely to also increase the error rates and induce cross-talk (which is unintentional interactions between neighbouring qubits). (AbuGhanem, "IBM Quantum Computers") Lastly, as a superconducting quantum processor, the Nighthawk naturally has some sacrifices. Most notably, when superconducting quantum processors scale, they lack the ability to preserve high fidelity.

Rigetti's Cepheus-1-108Q

Current Progress/Breakthroughs

Rigetti is also a strong competitor in the quantum computing industry, known for its innovative methods of error reduction in quantum computation(Arrasmith et al.) and its focus on producing high-quality superconducting quantum processors by incorporating superconducting microwave devices. (Sete et al.; Vahidpour et al.; Reagor et al.) One of their most recent quantum processors is the Cepheus-1-108Q model, which is a superconducting gate-based quantum processor that has 108 qubits: a system of twelve interconnected 9-qubit chiplets. (Rigetti Announces General Availability of 108-Qubit System) Due to the fact that the Cepheus-1-108Q model was released in April 2026, there is a lack of literature review; the vast majority of this information here is sourced from news articles published by Rigetti. This model is an incredibly massive milestone since it is the largest modular quantum computing system in the entire field of quantum computing, based on Rigetti's self-developed chiplet-based architecture. By refining the qubit and coupler design, this processor had faster two-qubit gate speeds and higher fidelity than previous generations. In addition, the implementation of CZ gates for error correction allowed the Cepheus-1-108Q to reach a two-qubit gate fidelity rate of 99.9% at 60 nanoseconds, which is a noticeable improvement in comparison to the two-qubit gate fidelity of ANKAA-3 of 99.5% (previous-generation processor). (Rigetti Computing Launches 84-Qubit Ankaa™-3 System; Achieves 99.5% Median Two-Qubit Gate Fidelity Milestone) Moreover, the signal-to-noise ratio during qubit readout has been upgraded with the aid of the latest engineered control system. Lastly, a technique Rigetti introduces called the Alternating-Bias Assisted Annealing technique in the fabrication process is able to maintain qubit consistency better. Ultimately, this allows better control, thus reducing the error rate, and contributes to higher fidelity. To sum up, in Rigetti's latest processor, they made advancements in qubit and coupler design, utilised CZ gates for error correction, improved control electronics and used the ABAA technique in the fabrication process (in general, these improvements contribute to higher fidelity).

Issues

Despite Rigetti's accomplishments in the quantum computing field, like its rivals, there are still hardships Rigetti needs to overcome. One of the main problems is that the two-qubit gate fidelity of Cepheus-1-108Q is inadequate, currently sitting at approximately 99.1%. Hence, Rigetti is urgently attempting to improve the accuracy of their quantum computer to have a two-qubit gate fidelity of 99.5% in the remaining months of this year. Other than that, due to the new designs, a new problem regarding the coherence times has occurred. Additionally, Rigetti's Cepheus-1-108Q, unlike Google's "Willow", has yet to achieve quantum advantage. (Jin) Nevertheless, Rigetti is aiming to accomplish this feat in the next 3 years. To conclude, Rigetti's newest model of quantum computer struggles with accuracy, coherence times and achieving quantum advantage.

Metric Comparison

Table 2. Comparison of the key performance metrics of each superconducting quantum computer/processor.

Metric (Quantitative)		Superconducting Quantum Computers/Processors		
		Nighthawk	**Willow**	**Cepheus-1-108Q**
Scale	Number of Qubits	120	105	108
Quality	Readout Fidelity (%)	97.76*	99.50	NR
	Single-Qubit Gate Fidelity (%)	NR	99.97	99.90
	Two-Qubit Gate Fidelity (%)	99.89*	99.88	99.10
	Coherence Time T1 (μs)	350	68	25
	Coherence Time T2 (μs)	NR	89	10
Speed	1 Qubit Gate Time (ns)	NR	~25 #	NR
	2 Qubit Gate Time (ns)	NR	~40 #	~60
	Measurement Time (ns)	NR	~500 #	NR

Data compiled from Refs(Jin; Google Quantum AI and Collaborators et al.; “Building Scalable, Innovative Quantum Systems”; “IBM Quantum Compute Service Changelog”; Compute Resources | IBM Quantum Platform; Google LLC (via Public) / Our Quantum Hardware: The Engine for Verifiable Quantum Advantage)

The values in the Nighthawk column marked with * are the current calibration values. These will alter depending on how IBM recalibrates the Nighthawk. (properties like T1, T2, and operation errors can alter). The values in the Willow column marked with # are from Kong et al.(Kong et al.) Spaces labelled NR mean “Not Reported”.

Despite the presence of “not reported”values in Table 2, from the metrics available, we are able to deduce that the “Willow” (from Google) is likely the superconducting quantum computer/processor with the greatest computational power. Firstly, although the Willow has the least amount of total qubits (so it places last in the scale criteria), it has very similar fidelity values in comparison to the Nighthawk and exceeds the values of Cepheus-1-108Q. Evidently, this suggests it is one of the if not the most accurate quantum computer/processor out of the line up, hence it places first in the criteria of quality, which is arguably the most important criteria considering all other criteria of quantum computers being useless without accurate computation. Additionally, the Willow is the sole quantum computer that has achieved the feat of quantum advantage out of the three, further suggesting its computational superiority. Nonetheless, there are certain notable inconsistencies that render it nearly impossible to fairly compare the metrics. The metrics of each quantum computer have slight differences during the measurement process. For example, the coherence times of Google’s Willow utilise the mean times, whilst Rigetti uses the medians for their coherence times. Thus, this is technically an unfair comparison. Additionally, due to these quantum computers/processors being the latest generations, there is a lack of literature review and available information accessible to the public; hence there are numerous “not reported” values. Lastly, there is a difference in evidence quality. The Nighthawk

and Cepheus-1-108Q's data were obtained from company announcements, whereas the Willow's data is retrieved from a peer-reviewed paper. This also makes the comparison process much more difficult and unfair.

Conclusion

To conclude, the comparison of the top three contenders in the superconducting quantum computer industry elucidates how substantial achievements have been made, but it remains a flourishing industry that lacks sophistication due to its novelty. Notably, Google's Random Circuit Sampling (RCS) has yielded positive results that prove quantum advantage. Additionally, all competitors have displayed significant scaling improvements in comparison to their previous quantum computer/processor generations. Other than purely the increase in the number of qubits, the connectivity methods have also been upgraded to prevent crosstalk and reduce error rates. More importantly, despite the exponential increase in interactions between qubits when scaling, refined connectivity across all 3 companies has allowed higher two-qubit gate fidelity. Thus, enabling more complex entangling operations. On the other hand, Quantum Error Correction (QEC) still struggles to deal with the disruptive noise that interrupts quantum systems. All three companies have attempted various methods to deal with the noise: Surface code, tunable couplers and CZ gates are some examples. Indeed, these methods have raised the coherence times of their quantum computers/processors, which leads to higher gate and readout fidelity, meaning fewer errors. Nevertheless, reports across all three companies state that this level of QEC is simply not powerful enough, as they aim for higher fidelity rates in future years. To sum up, our recent advances are yet to reach the fault-tolerant era of quantum computers.

Moreover, by comparing the achievements and issues of each company, the metrics likely suggest that Google's "Willow" is the most powerful quantum computer/processor out of the three. Nonetheless, due to the lack of a universal benchmark, the lack of accessible metrics for the latest quantum technology and inconsistencies in how the metrics are measured, this conclusion about the "Willow" is unreliable and unfair. In order to allow the evolution of this field to be tracked coherently and pragmatically, a standard benchmark that is capable of testing all quantum computers, regardless of their technological platform and architecture, and proving practical use is incredibly crucial.

Work Cited

AbuGhanem, Muhammad. “IBM Quantum Computers: Evolution, Performance, and Future Directions.” The Journal of Supercomputing, vol. 81, no. 5, Apr. 2025, p. 687. Springer Link, https://doi.org/10.1007/s11227-025-07047-7.

———. “Superconducting Quantum Computers: Who Is Leading the Future?” EPJ Quantum Technology, vol. 12, no. 1, Dec. 2025, p. 102. DOI.org (Crossref), https://doi.org/10.1140/epjqt/s40507-025-00405-7.

AbuGhanem, Muhammad, and Hichem Eleuch. “NISQ Computers: A Path to Quantum Supremacy.” IEEE Access, vol. 12, 2024, pp. 102941–61. arXiv.org, https://doi.org/10.1109/ACCESS.2024.3432330.

Acuaviva, Arturo, et al. “Benchmarking Quantum Computers: Towards a Standard Performance Evaluation Approach.” Quantum Science and Technology, vol. 11, no. 2, Feb. 2026, p. 025004. Institute of Physics, https://doi.org/10.1088/2058-9565/ae4269.

Arrasmith, Andrew, et al. Development and Demonstration of an Efficient Readout Error Mitigation Technique for Use in NISQ Algorithms.

Arute, Frank, et al. “Quantum Supremacy Using a Programmable Superconducting Processor.” Nature, 2019. sci-hub.box, https://doi.org/10.1038/s41586-019-1666-5.

Bauer, Bela, et al. “Quantum Algorithms for Quantum Chemistry and Quantum Materials Science.” Chemical Reviews, 2020. sci-hub.box, https://doi.org/10.1021/acs.chemrev.9b00829.

Bernien, Hannes, et al. “Probing Many-Body Dynamics on a 51-Atom Quantum Simulator.” Nature, 2017. sci-hub.box, https://doi.org/10.1038/nature24622.

Bouland, Adam, et al. “On the Complexity and Verification of Quantum Random Circuit Sampling.” Nature Physics, 2018. sci-hub.box, https://doi.org/10.1038/s41567-018-0318-2.

“Building Scalable, Innovative Quantum Systems.” Rigetti Computing, https://www.rigetti.com/what-we-build. Accessed 6 Sept. 2026.

Compute Resources | IBM Quantum Platform. https://quantum.cloud.ibm.com/computers?processorType=Nighthawk&. Accessed 6 Sept. 2026.

Czincila, Connor, et al. Nuclear Magnetic Resonance (NMR) Qubits.

Debnath, S., et al. “Demonstration of a Small Programmable Quantum Computer with Atomic Qubits.” Nature, 2016. sci-hub.box, https://doi.org/10.1038/nature18648.

Domingo, L., et al. “Binding Affinity Predictions with Hybrid Quantum-Classical Convolutional Neural Networks.” Scientific Reports, vol. 13, Oct. 2023, p. 17951. PubMed Central, https://doi.org/10.1038/s41598-023-45269-y.

Ebadi, Sepehr, et al. “Quantum Phases of Matter on a 256-Atom Programmable Quantum Simulator.” Nature, 2021. sci-hub.box, https://doi.org/10.1038/s41586-021-03582-4.

Five Years Ago Today, We Put the First Quantum Computer on the Cloud. Here's How We Did It. | IBM Quantum Computing Blog. https://www.ibm.com/quantum/blog/quantum-five-years. Accessed 11 Aug. 2026.

Goodhart, C. A. E. (Charles Albert Eric). Monetary Theory and Practice : The UK Experience. With Internet Archive, London : Macmillan, 1984. Internet Archive, http://archive.org/details/monetarytheorypr0000good.

Google LLC (via Public) / Our Quantum Hardware: The Engine for Verifiable Quantum Advantage. https://www.publicnow.com/view/FCB712F89E85303E02CA0BFC5E662A48CF475140. Accessed 6 Sept. 2026.

Google Quantum AI and Collaborators, et al. "Quantum Error Correction below the Surface Code Threshold." Nature, vol. 638, no. 8052, Feb. 2025, pp. 920–26. DOI.org (Crossref), https://doi.org/10.1038/s41586-024-08449-y.

Groumpos, Peter P. "A Critical Historical and Scientific Overview of All Industrial Revolutions." IFAC-PapersOnLine, vol. 54, no. 13, 2021, pp. 464–71. DOI.org (Crossref), https://doi.org/10.1016/j.ifacol.2021.10.492.

Herman, Dylan, et al. "Quantum Computing for Finance." Nature Reviews Physics, vol. 5, no. 8, July 2023, pp. 450–65. DOI.org (Crossref), https://doi.org/10.1038/s42254-023-00603-1.

"IBM Quantum Compute Service Changelog." IBM Quantum Documentation, https://quantum.cloud.ibm.com/docs/en/guides/quantum.cloud.ibm.com/docs/en/guides/changelog-quantum-compute-service. Accessed 6 Sept. 2026.

Jin, Yirong. "Google's 'Willow' Quantum Processor: New RCS Record and First Error Correction below the Surface Code Threshold." The Innovation, vol. 6, no. 7, July 2025. www.cell.com, https://doi.org/10.1016/j.xinn.2025.100942.

Kong, Linghang, et al. "Benchmarking Fault-Tolerant Quantum Computing Hardware via QLOPS." ACM Transactions on Quantum Computing, vol. 7, no. 2, Apr. 2026, p. 14:1-14:14. ACM Digital Library, https://doi.org/10.1145/3797968.

Ladd, T. D., et al. "Quantum Computers." Nature, vol. 464, no. 7285, Mar. 2010, pp. 45–53. DOI.org (Crossref), https://doi.org/10.1038/nature08812.

McEwen, Matt, et al. "Resolving Catastrophic Error Bursts from Cosmic Rays in Large Arrays of Superconducting Qubits." Nature Physics, 2021. sci-hub.box, https://doi.org/10.1038/s41567-021-01432-8.

Neumaier, Arnold. "The Born Rule—100 Years Ago and Today." Entropy, vol. 27, no. 4, Apr. 2025, p. 415. www.mdpi.com, https://doi.org/10.3390/e27040415.

Orus, Roman, et al. "Quantum Computing for Finance: Overview and Prospects." Reviews in Physics, vol. 4, Nov. 2019, p. 100028. arXiv.org, https://doi.org/10.1016/j.revip.2019.100028.

Phillipson, Frank. "Quantum Computing in Logistics and Supply Chain Management an Overview." arXiv:2402.17520, arXiv, 2 Apr. 2025. arXiv.org, https://doi.org/10.48550/arXiv.2402.17520.

Preskill, John. “Quantum Computing in the NISQ Era and Beyond.” Quantum, vol. 2, Aug. 2018, p. 79. arXiv.org, https://doi.org/10.22331/q-2018-08-06-79.

Reagor, Matthew J., et al. Hardware Optimized Parity Check Gates for Superconducting Surface Codes.

Rieffel, Eleanor G., and Wolfgang Polak. “An Introduction to Quantum Computing for Non-Physicists.” arXiv:quant-ph/9809016, arXiv, 19 Jan. 2000. arXiv.org, https://doi.org/10.48550/arXiv.quant-ph/9809016.

Rietsche, Roman, et al. “Quantum Computing.” Electronic Markets, vol. 32, no. 4, Dec. 2022, pp. 2525–36. DOI.org (Crossref), https://doi.org/10.1007/s12525-022-00570-y.

Rigetti Announces General Availability of 108-Qubit System.

Rigetti, Chad, et al. Superconducting Qubit in Waveguide Cavity with Coherence Time Approaching 0.1ms.

Rigetti Computing Launches 84-Qubit Ankaa™-3 System; Achieves 99.5% Median Two-Qubit Gate Fidelity Milestone.

Rossi, Alessandro, et al. “Quantum Computing Hardware in the Cloud: Should a Computational Chemist Care?” International Journal of Quantum Chemistry, vol. 121, no. 14, 2021, p. e26688. Wiley Online Library, https://doi.org/10.1002/qua.26688.

Santagati, Raffaele, et al. “Drug Design on Quantum Computers.” Nature Physics, vol. 20, no. 4, Apr. 2024, pp. 549–57. arXiv.org, https://doi.org/10.1038/s41567-024-02411-5.

Sete, Eyob A., et al. “A Functional Architecture for Scalable Quantum Computing.” 2016 IEEE International Conference on Rebooting Computing (ICRC), 2016. sci-hub.box, https://doi.org/10.1109/ICRC.2016.7738703.

Shirts, Randall B., and John S. Welch. “Quantum Interstate Phase Differences and Multiphoton Processes: Quantum Jumps or Dynamic Beats?” ACS Omega, vol. 7, no. 34, Aug. 2022, pp. 30632–41. DOI.org (Crossref), https://doi.org/10.1021/acsomega.2c04554.

Shor, P. W. “Algorithms for Quantum Computation: Discrete Logarithms and Factoring.” Proceedings 35th Annual Symposium on Foundations of Computer Science. sci-hub.box, https://doi.org/10.1109/SFCS.1994.365700. Accessed 6 Aug. 2026.

Shrapnel, Sally, et al. “Updating the Born Rule.” New Journal of Physics, vol. 20, no. 5, May 2018, p. 053010. Institute of Physics, https://doi.org/10.1088/1367-2630/aabe12.

Sim, S. E., et al. “Using Benchmarking to Advance Research: A Challenge to Software Engineering.” 25th International Conference on Software Engineering, 2003. Proceedings., 2003. sci-hub.box, https://doi.org/10.1109/ICSE.2003.1201189.

Steane, Andrew. “Quantum Computing.” Reports on Progress in Physics, vol. 61, no. 2, Feb. 1998, p. 117. Institute of Physics, https://doi.org/10.1088/0034-4885/61/2/002.

Stehlik, J., et al. “Tunable Coupling Architecture for Fixed-Frequency Transmon Superconducting Qubits.” Physical Review Letters, vol. 127, no. 8, Aug. 2021, p. 080505. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.127.080505.

Vahidpour, Mehrnoosh, et al. Superconducting Through-Silicon Vias for Quantum Integrated Circuits.

"We're Scaling Quantum Computing Even Faster with Atlantic Quantum." Google, 2 Oct. 2025, https://blog.google/innovation-and-ai/technology/research/scaling-quantum-computing-even-faster-with-atlantic-quantum/.

Wu, Yulin, et al. "Strong Quantum Computational Advantage Using a Superconducting Quantum Processor." Physical Review Letters, vol. 127, no. 18, Oct. 2021, p. 180501. DOI.org (Crossref), https://doi.org/10.1103/PhysRevLett.127.180501.

Zhang, J., et al. "Observation of a Many-Body Dynamical Phase Transition with a 53-Qubit Quantum Simulator." Nature, 2017. sci-hub.box, https://doi.org/10.1038/nature24654.

Zhong, Han-Sen, et al. "Quantum Computational Advantage Using Photons." Science, 2020. sci-hub.box, https://doi.org/10.1126/science.abe8770.

Can a Vegan Diet Meet the Nutritional Needs of Children to Promote Health? A Narrative Review of Current Evidence By Carlos Orfanidis Cordoba

Abstract

Plant based diets have grown in popularity in recent years, even decades, for many reasons. The vegan diet is among these, and there has been considerable research done on vegan diets followed by adults. On the other hand, children following vegan diets are both harder to find and thus study. Hence, the research and literature regarding vegan diets for child populations is limited. By studying the dietary recommendations for children of all ages, the existing literature on vegan children, and the importance of family planning for healthy eating habits of children, it is easier to conclude whether or not the vegan diet can meet the nutritional needs of children to thus promote health. This included the consulting of health authorities like the WHO, National Institutes of Health, and Johns Hopkins Medicine group for recommended dietary practices and micronutrient allowances. It also involved a literature review on micronutrient bioavailabilities in specific contexts and child eating habits under family households. Lastly, it involved the careful examination of three major vegan-children cohort studies. Ultimately, all these investigations showed that the vegan diet can promote health when common pitfalls are foreseen and addressed in dietary planning, specifically as the responsibility of parents. This especially included supplementation and/or fortification of micronutrients vitamin B12, iron, and calcium. Moreover, they pointed to the potential for vegan diets to promote good to excellent cardiovascular health through biomarkers like cholesterol metrics and certain anthropometrics.

Introduction

Vegans and veganism. Nowadays, most people have heard of these words, but perhaps not what they mean. *Veganism* is the set of beliefs that *vegans* live by, and it rests on the simple idea that humans should try their best to avoid exploiting animals for "food, clothing, or any other purpose" (The Vegan Society). For many reasons, people tend to clash with this doctrine, particularly its outlook on diet and nutrition: the upkeep of the vegan, or *fully* plant-based diet.

The question of the adequacy of the vegan diet is an important one, because of the ethical and environmental implications it entails. More importantly, a growing number of people around the world are taking on the vegan diet, and events like "Veganuary" have also grown in popularity in different parts of the world, ranging from Brazil to Malaysia (Chelsea). More children and adults in high-income regions are uptaking this type of lifestyle, and the dietary pattern it encompasses (Desmond et al.). The scientific method and rigorous research can be utilised to determine whether or not this dietary pattern is healthful. Oftentimes, however, it can be tricky to navigate the nuances of nutritional information, all the while grappling with digital tribal wars between dietary camps like the vegan, carnivorous, and omnivorous ones that, very unfortunately, tend to stall this advance of scientific discussion and research. The aim of this paper is to answer the question: *can the vegan diet actually meet the nutritional needs of a growing child to promote health?*

This question is attempted to be answered with a clear structure: discussing crucial micro and macro nutrient concerns for infants and toddlers, examining three major cohort studies of vegan children populations, and discussing healthy eating habits promoted in households. This provided a sound and broad review of what factors determine the healthiness of the vegan diet, involving both social feasibility, the hypothetical capacity for plant based foods to deliver micronutrients, and real-life cohorts of vegan children with measurable health outcomes and biomarkers. In the initial discussion of infant nutrition, there is little reference to whole food groups and an emphasis on micronutrients of concern because infants are nourished largely by human milk and/or infant formula rather than an array of whole foods. In the discussion of toddler nutrition, this opens up to discuss said food groups, considering them more deeply. Discussion of age groups beyond toddlers is excluded because it would involve largely the same facts on micronutrients of concern and the whole food groups.

Methodology

The primary means of research herein was investigating publications from internationally recognized health authorities such as the World Health Organization, the (American) National Institutes of Health, and others like Johns Hopkins Medicine Group. The publicly available information on micronutrients, recommended daily allowances, and food groups these organizations offer form the backbone of the research done. No exclusion dates-of-publication were employed in selection of articles, but peer reviewed publications were used to investigate and communicate micronutrient absorption rates (particularly for zinc, calcium, and iron). Use of databases PubMed and journal MDPI nutrients also provided resources for micronutrient discussion among both infant and toddler nutrition consideration. Keywords involved each micronutrient and “intake”, “recommendations” and “dietary allowances”. All this based the review of child micronutrient intake, possibilities for deficiencies, and recommendations pertaining to each micronutrient explored.

On the other hand, the review of literature on vegan-vegetarian (child) populations was backed by three major studies: the VeChi Diet study of children 1-3 years of age, the VeChi Youth study of children 6-18 years of age, and a cross-sectional study of Polish children aged 5-10 from University College London and collaborators. All of these are peer reviewed and published in renowned journals (MDPI nutrients, MDPI nutrients, and The American Journal of Clinical Nutrition respectively).

Lastly, a peer reviewed literature review was selected to inform research on the social, family, and habitual considerations of not only the vegan diet, but children eating habits overall. This publication, peer reviewed in the MDPI nutrients journal, informs the final discussion on the vegan diet from a different but highly relevant perspective, which is its feasibility from a social, human perspective.

Regarding nutrition, the two lenses through which the adequacy of the vegan (VN) diet was examined were: dietary recommendations for varying age groups (considering mostly micro nutrient needs), assessing how adequately a VN diet can take care of those needs. Secondly, the

above outlined studies examine the rather limited but known populations of children on a VN diet and observe the health outcomes associated with their diet. To gauge the deepest insight on this topic, an overlapping of both of these perspectives was employed, then paired with a human habits discussion.

Diagrams in the discussion of cohort studies of vegan children populations are AI created by ChatGPT model 5.6 Luna.

A Brief Initial Discussion - Vitamin B12

A vegan diet is known for lacking any whole food source naturally containing vitamin B12. The following sections on the potential for a vegan diet to meet infant and toddler nutrition thus do not discuss it. However, discussion on vitamin B12 is extensive in the population studies of vegan and vegetarian children.

Any well planned vegan diet should include fortified foods or supplementation, and the range of supplements commercially available facilitates this planning. It would be very unwise for parents not to at least consider fortified foods, like breakfast cereals, dairy alternatives, and other options for vitamin B12 sourcing.

Infant Nutrition Considerations and Micronutrients

Infants, or children aged 0 to 12 months, have nutritional needs which boil down to: exclusive breastfeeding until 6 months of age (or infant formula, which exists in plant based forms), and breastfeeding (or the same infant formula) paired with complementary foods until even after 12 months of age (Sadeeh et al.). Good complementary foods help bridge the gap between, after 6 months of age onwards, the nutritional needs of infants and the nutrients provided by breastmilk (which cannot fulfil said needs). Detailed discussion on the 5 whole food groups encouraged after 6 months is reserved until toddler nutrition is discussed, where omega 3 fatty acids are also introduced. These are also important for complementarily feeding infants. The foods important for complementarily feeding infants are, though, “rich in energy, protein and micronutrients (particularly iron, zinc, calcium, vitamin A, vitamin C and folate)”, according to an informal World Health Organization document.

It is indicated that the following micronutrients are of particular importance for infant nutrition: iron, zinc, calcium, vitamin A, vitamin C, folate. Energy density and protein density are also important.

Because of the increased bioavailability of iron in animal foods, these are recommended as a daily staple in the mixture of complementary foods needed to ensure the health of growing infants. Infants have iron stores which deplete after about 6 months, at which point these complementary foods consequently become necessary (Sadeeh et al). Indeed, it is emphasized that iron is the most difficult nutrient to source. So can plant-based foods, besides providing energy and protein, also provide hard-to-source micronutrients like iron?

Iron

The WHO has identified three levels of average bioavailability of iron in different diets in 1988. The diet with 'generous' amounts of 'flesh food' and vitamin C had the highest absorption rate (15%) while "very restricted" vegetarian diets had the lowest absorption rate (5%) (National Academies, Iron). Plant-based foods rich in iron include certain legumes, and leafy greens (Johns Hopkins Medicine, Iron Deficiency). Their absorption, however, varies greatly due to other dietary factors, whereas heme iron is "always well absorbed" (National Academies, Iron). That said, the pairing of non heme iron with ascorbic acid - vitamin c - enhances absorption (National Academies, Iron). This presents a better alternative for iron absorption than plain consumption of the aforementioned foods. The question, however, regarding a 'better alternative' for iron absorption and fulfillment for adequate nutrition remains difficult to answer, mostly because the bioavailability of iron across plant and animal foods varies even across the foods examined, thus making it challenging to universally declare rates of absorption among these two food groups. It is advisable that parents decide which food groups are more suitable for the health of their child and iron fulfillment. This, however, is a choice somewhat simplified by the existence, in most parts of the world, of iron fortified foods. Nowadays, with the advent of renewed and better nutritional information, manufacturing companies of food products can fortify their products with micronutrients like iron for their consumers. Infant formulas, iron-fortified cereals, flours, juices, etc. are popular and more inexpensive choices for families, presenting an easy fix to the difficulty that many children worldwide face in consuming enough iron (Baltussen et al.**)** Therefore, it seems judicious to say that iron sourcing is feasible through a vegan diet when vitamin C is paired with iron rich foods and iron fortified foods are regularly consumed. Nevertheless, careful consideration of this micronutrient and its provision in daily food products is of important concern for infants consuming a vegan diet.

Zinc

The case for zinc is very similar to that of iron. Plant based zinc is less readily absorbed by the body because the chemical compound phytate, which is found in many plant based foods, binds to zinc and inhibits its absorption. This means zinc from animal sources has higher bioavailability because of lower phytate content (National Academies, Zinc). Thus, it is listed by the WHO document that good sources of zinc include "liver and offal of all kinds, foods prepared with blood, flesh of animals, birds and fish, shellfish, egg yolk" (Sadeeh et al.). Plant based foods with high zinc content, on the other hand, include certain lentils, cereals - particularly breakfast fortified ones - pumpkin seeds, and oats (National Institutes of Health, Zinc). As vitamin C boosts the absorption of iron in plant foods, foods can be soaked, fermented, malted, or germinated to reduce their phytate content (Hall and King). The biggest concern, regarding both zinc and iron, that parents face in feeding their infants, is that without fortified infant formulas or foods, it is challenging to guarantee adequate zinc absorption rates of proper intakes. This, however, is dependent on which food products are chosen as sources of iron and zinc, and the satiety levels experienced by the child consuming them - again, it is wise for parents to explore options safely and smartly, and certainly those that are fortified. Experts

recommend up to 50% greater intakes of dietary zinc content to be sure of adequate intake levels, so parents should be conscious of the challenge that arises from sourcing it (National Academies, Zinc).

Calcium

The next micronutrient of concern is calcium. The WHO document lists milk products and fish consumed bone-in as good sources of calcium (Sadeeh et al.). Dairy products are commonly consumed sources of calcium, but they are not the only ones. Plant based sources, unlisted there, include calcium set tofu, some dark leafy greens (kale and bok choy), some seeds (chia, sesame, poppy), almonds, and some legumes (soybeans, chickpeas, goa beans, etc) (Shkembi and Huppertz). Calcium absorption in these plant sources varies, where some sources actually have increased absorption rates compared to dairy, whereas others have lower absorption rates (Shkembi and Huppertz). Cruciferous vegetables like broccoli, cauliflower, and kale have high rates of absorption similar to dairy products but their calcium density varies (Shkembi and Huppertz). Moreover, fortified foods like cereals, juices and most importantly, infant formulas, containing calcium do too exist. Because a growing infant needs about 270 milligrams of calcium when complementary feeding starts at the 6 month mark, a healthful combination of infant formula and complementary foods is likely capable of meeting this need (National Institutes of Health, Calcium). Assuming an infant formula rich in nutrients like calcium is being fed, which is suggested to address the nutrient needs of the child, and that its calcium content is at a minimum of 50 milligrams per 100 grams of infant formula, calcium requirements can likely be met. Approximately 260 milligrams of calcium can be consumed daily *assuming* 500 to 700 milliliters of this formula are consumed. Moreover, this does not even account for the complementary foods aforementioned. Thus, infant formula alone, rich in calcium, covers the majority of calcium consumption needed for a healthful infant from 6 to 12 months of age, where complementary foods like cruciferous vegetables and soy derived products, among others, bolster this consumption.

Other Micronutrients

To briefly comment on the other micronutrients that the informal WHO document mentions of importance in the complementary foods that infants should begin to consume at 6 months of age, which are vitamin A, vitamin C, and folate, these are accessible in plant based sources. Vitamin A is formed in the body through the consumption of pro vitamin A carotenoids which are compounds found in certain plant foods like carrots and sweet potatoes (National Institutes of Health, Vitamin A and Carotenoids). However, genetic variability may affect some people's ability to convert these carotenoids to vitamin A (National Institutes of Health, Vitamin A and Carotenoids). Vitamin A is fat soluble, meaning it is best absorbed when eaten with "healthy fats" (Cleveland Clinic, Fat Soluble Vitamins).

Vitamin C is readily found in fruits and vegetables, so for the plant based diet, it is of much lower concern. Kiwis, strawberries, oranges, broccoli, and peppers are accessible foods

rich in vitamin C (National Institutes of Health, Vitamin C).

Similarly, folate is readily found in plant foods too. Legumes, rice, spinach, asparagus, and brussels sprouts, to name but a few, are good plant sources of folates - of which infants past 6 months need about 80 micrograms (National Institutes of Health, Folate). 40 grams of boiled spinach alone, for instance, provides over 50 micrograms of folate (National Institutes of Health, Folate).

Toddler Nutrition Considerations and Micronutrients

Once infants reach a year of age, they are toddlers, and it is no longer necessary for them to consume infant formula - though breastfeeding, if "mutually desired", is recommended beyond even two years of age (American Academy of Pediatrics). Among the Johns Hopkins medical association, Nemours children hospital, and Raising Children Network, there is a consensus of feeding toddlers 5 whole food groups in defined ranges of servings everyday. These are: whole grains, fruits, vegetables, proteins, and dairy. Considering these food groups allows us to zoom out from the discussion of micronutrients alone, and into the elimination of certain foods entailed by a vegan diet. This discussion is not only relevant to toddlers, but infants (past the 6 month mark) and children of all ages. It also introduces the importance of omega 3 fatty acids.

Protein and Dairy

Because the first three of the five whole food groups consist of plant foods, only protein and dairy require scrutiny. The protein food group, first of all, consists of a diverse range of proteins. The aforementioned organizations all recommend beans, with Johns Hopkins and Raising Children Network also encouraging nuts and seeds (in paste form to be easily consumed). Is a risk posed by removing fish and meat from this food group, such as a micronutrient deficiency?

Iron

We have covered heme and non heme iron already, in how they differ between animal and plant products. Again, vitamin C pairing and methods of cooking like boiling help boost lower absorption rates, and fortified foods are adequate iron sources too. Livers, fish, chicken, turkey, among other meats offer higher absorption rates or this important micronutrient, of which toddlers need 7 mg (National Institutes of Health, Iron). Ultimately, it is the job of parents to make informed decisions on which foods make up for cutting out heme iron from meat of the protein group. When iron-rich whole foods like some legumes are paired with fortified foods, there is reason to believe iron is largely obtainable from plant based sources.

Omega-3 fatty acids, For both Toddlers and Infants.

Another incredibly important nutrient is the omega 3 fatty acids group. Omega 3 fatty acids are long chained molecules that the body cannot produce, thus making them essential (Cleveland Clinic). Though many of these molecules exist, only three of them are of major

importance in scientific study: EPA (eicosapentaenoic acid), DHA (docosahexaenoic acid), and ALA (Alpha-Linolenic Acid) (National Institutes of Health, Omega 3 Fatty Acids). ALA is readily available in several plant foods, like flaxseeds, flaxseed oil, chia seeds, and certain walnuts (National Institutes of Health, Omega 3 Fatty Acids). That said, ALA is actually converted into DHA and EPA by the human body at limited rates which may not support the body's needs (National Institutes of Health, Omega 3 Fatty Acids). The only true plant sources of DHA and EPA are microalgae, which are consumable in supplements like soft gels (Bailey et al.). Should parents consider powders or pills for omega 3 sourcing? It may be an adequate choice, but the microalgae species being consumed matters. Perhaps fish is more convenient for most, but parents may make the choice of using supplemented or fortified foods (like plant based butters, juices, or infant formulas) to source DHA and EPA. Prices would vary between options of sources, from supplements to fish. All in all, omega 3 fatty acids, particularly DHA and EPA, are much harder to source from whole plant foods than the other micronutrients discussed. Hence, a plant based diet not including fortified foods or supplementation has the potential to increase the risk of omega 3 deficiency.

Zinc

The discussion of zinc remains similar as above. Toddlers need about 3 mg of zinc a day (CDC). Fortified foods like cereals, certain legumes already mentioned, and certain seeds (like pumpkin seeds) are all rich sources of zinc. However, their absorption rates are lower than those of animal foods, like beef and certain seafood**.** Even for cereals, the rate of zinc absorption varies hugely based on the content levels of phytates, and so dietary absorption may range from 5 to 50% based on phytate content consumed daily - so experts recommend up to 50% greater intakes of dietary zinc content to be sure of adequate intake levels (National Academies, Zinc). This means upwards of 4.5mg of zinc is suggested, which may be achievable with mixed fortified grains, legumes and certain seeds - but it may pose a challenge to parents whose children are picky eaters, so this may not be a convenient alternative. Ultimately, parents should decide a dietary course of action - and for the protein food group, switching fully to plant sources may pose an inconvenience, particularly in trying to reach adequate zinc intakes.

The Dairy Food Group

Let us now consider the dairy food group. What micronutrients may dairy offer that are hard to obtain from plant sources? While dairy is rich in calcium, high protein density, and other notable micronutrients, many fortified plant based alternatives are meant to match these nutrient profiles (Ramsing et al.). That said, "retail unit price" for these alternatives is generally higher (Ramsing et al.). Importantly, the American Academy of Pediatrics only recognizes soy based milk alternatives as safe to be consumed for children under 24 months of age (Ramsing et al.). There are countless other plant foods rich in calcium, and this discussion may be found in the above section on calcium intake for toddlers - the facts on plant foods and bioavailability remain the same, it is just that toddlers require about 700 mg of calcium instead of 260 (NIH). Plant

based yoghurts, like soy based ones, are widespread in availability too (Grand View Research). These can also be found in fortified forms, with calcium and B12 added in particular, to match the nutritional profile of dairy products (D'Andrea et al.). It may, however, be to the detriment of many that these products are slightly more costly, thus posing an inconvenience. That said, there are seemingly available commercial vegan alternatives to most dairy products, particularly because calcium and vitamin B12 are obtainable in many of these alternatives. Parents who opt for these should look for said fortified options, and if desired, those with higher protein density.

By and large, the VN diet remains mostly adequate for toddler nutrition, but it has a few drawbacks. These are correctible with proper parent planning and consideration, yet may hinder the feasibility of the diet. To reiterate, these are the needs to consider: omega 3 sourcing, adequate zinc intake because of phytate-caused inhibition of zinc absorption, and higher costs of vegan alternatives to dairy products.

Summarizing Micronutrients of Special Concern

Micronutrients that require special attention on a vegan diet are iron, zinc, calcium, and omega 3 fatty acids. To ensure adequate levels of regular intake, parents are encouraged to find fortified foods, employ means to increase absorption levels with respective micronutrients, in some cases, increase daily intakes, and if necessary, supplement the diet as seen fit. It is highly advisable that a nutritionist be formally consulted for this final suggestion.

Considering Studied Vegan (Child) Populations

VeChi Youth Study

The first VeChi study being discussed, based in Germany, examined children and adolescents aged 6-18 across a 15 month period starting October 2017 and ending in January 2019. It collected data in a cross sectional design on "anthropometrics, diet, lifestyle and nutritional status". 401 children were studied and extensive blood sampling was completed.

Its dietary assessment of foods consumed by each diet group (VN, Vegetarian (VG), Omnivore (OM)) was measured with data from a 3 day period where participating families measured and listed all foods consumed in that period. Some semi quantitative measuring was used when these families could not use kitchen scales to weigh foods. The VeChi Youth Study lists extensive procedural standards for both its blood sampling methods and the use of statistics in its design. To best understand what is meant by "statistically significant", consult table 3 (below).

Supplementation

"Slightly more than half of the children and adolescents reported the use of dietary supplements". The group with highest supplementation rate was VN. 52% of VG (vegetarians) used dietary supplements, 95.5% of VN, and only 16.8% of OM (omnivores). 88% of VN supplemented vitamin B12, an important indication of a health conscious sample.

Dietary Energy and Macronutrient status

- On average, there was no statistically significant difference in protein intakes between diet groups, although the median intake for each group differed, with OM being highest (1.36 g/kg of bodyweight daily) and VG being lowest (1.14) - VN at 1.16.
- All groups exceeded the German reference intake for protein a day of 0.9 g/kg of bodyweight daily.
- Carbohydrate intake was, on average, higher in VG and VN than OM. OM ate the lowest fibre daily, whereas VN ate the most.
- Moreover, it was statistically significantly shown that VG and OM ate more fat than VN, though to a lesser degree of statistical significance.
- All in all, measured macronutrient intakes suggested all groups ate mostly healthily. This does not consider micronutrients.

Micronutrient Status among Dietary Groups

Measured values for micronutrient intake among the three dietary groups varied for most micronutrients. Vitamin B12 was particularly low for the VN group. Vitamin A intakes were similar across all three groups. The following diagram breaks down micronutrient status for all groups:

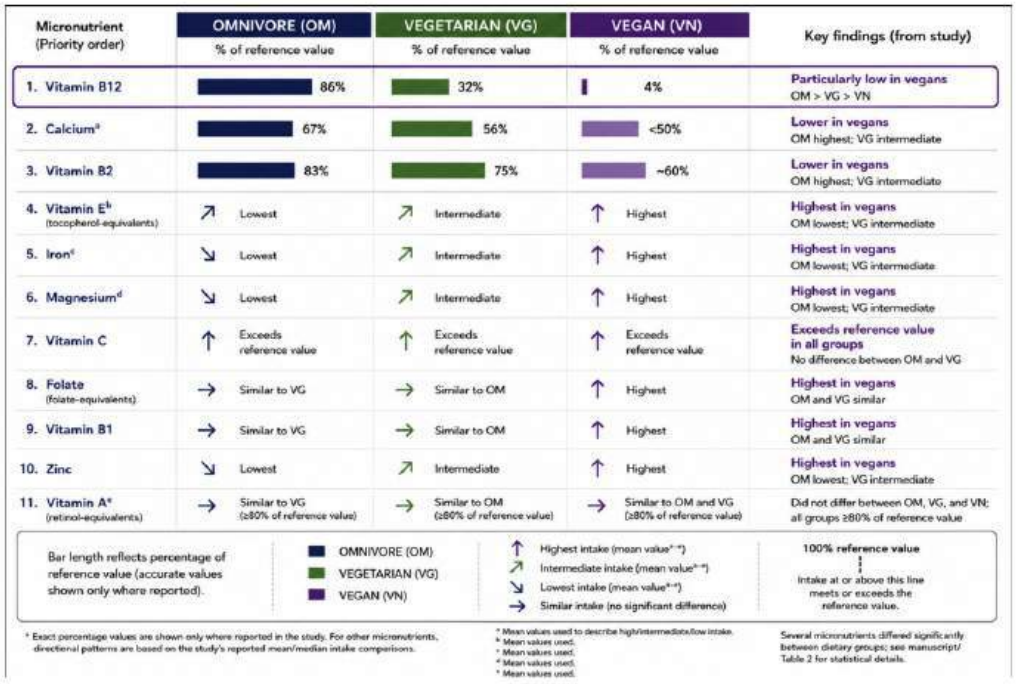

Figure 1: Micronutrient Intakes of Dietary Groups in the VeChi Youth Study
Data Compiled from Alexy, Ute et al, 2021, section 3.2.3 'Micronutrient Intake'. Created by ChatGPT, model 5.6 Luna, OpenAI,

Above are the key findings of measured micronutrient intakes across the groups. While Calcium, Vitamin E, Iron, Magnesium, and Vitamin A were measured with mean intakes across the groups as a % of reference intake, all other micronutrients were measured as % of reference intakes with median intakes across the groups.

The VN group was lowest in intake of calcium, Vitamin B12, and vitamin B2. It is worth noticing that calcium intake was lower than reference value (it is unspecified what reference value is taken) for all groups, as is the case with vitamins B12 and B2. However, vitamin B12 intake for the VN group was remarkably low. These findings reinforce the importance of careful planning and deliberation in dietary practices for VN diets. It is important, therefore, to note that the vegan diet must be well planned to ensure micronutrient intakes are at adequate levels to promote health. On the other hand, it is worth noting that zinc intake was highest among VN, which is in contrast to the mechanistic review of infant and toddler nutrition above. This finding indicates the complexity of food matrices and the limitations inherent in examining micronutrients alone. Other micronutrients like magnesium and vitamin B1 were also highest among VN, suggesting a somewhat increased capacity for this diet to provide other micronutrients like these.

Blood Lipids and Nutrient Biomarker results

Although the results for micronutrient intakes tell us a story, the blood testing for participants in this study tells us another. They add a layer of evidence, nuance, and complexity to the health these participants display in tested biomarkers. For example, while the VN group had the highest intake of iron, they had the lowest levels of blood ferritin (the protein storing iron in blood), statistically significantly lower than the OM group. Moreover, while vegans had the lowest measured intakes of vitamin B2, their blood levels did not significantly differ compared to the other groups. The same is true for methylmalonic acid, suggesting vitamin B12 deficiency should be tested for rather than asserted as a consequence of low B12 intake. These findings largely emphasize that micronutrient deficiencies should be assessed beyond intake levels, because the relation between intake levels and health outcomes is not so clear. For the vegan diet, it remains a fact that proper planning is crucial to success (avoiding pitfalls like B12 deficiency), but to be sure of health outcomes clinical practices should be employed. Moreover, the blood lipid metrics suggest a healthier cardiovascular profile of vegans because of lower total cholesterol and LDL cholesterol in particular. The below diagram shows the results of the blood lipids and nutrient biomarkers more simply:

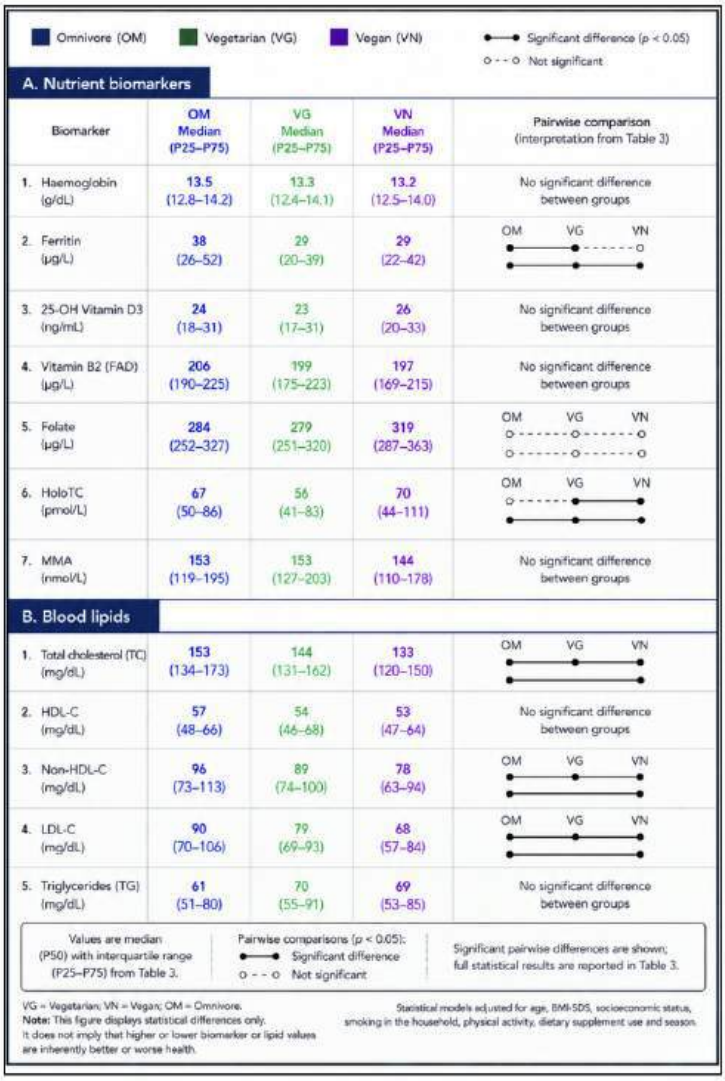

Figure 2: Blood Testing Results (blood lipids and nutrient biomarkers)
Data Compiled from Alexy, Ute, et al, 2021, section 3.3 'Nutrient Status and Blood Lipids', table 3. Created by ChatGPT, model 5.6 Luna, OpenAI,

Above are the key findings of the blood tests from the VeChi Youth Study. The bottom right footnote mentions 'Table 3', which is found in the VeChi Youth Study's publication itself, while this paper's 'figure 3' is distinguished by an asterisk. Both Figure 2 and Figure 3* should be read together for the best insight. Figure 3* is found directly below.*

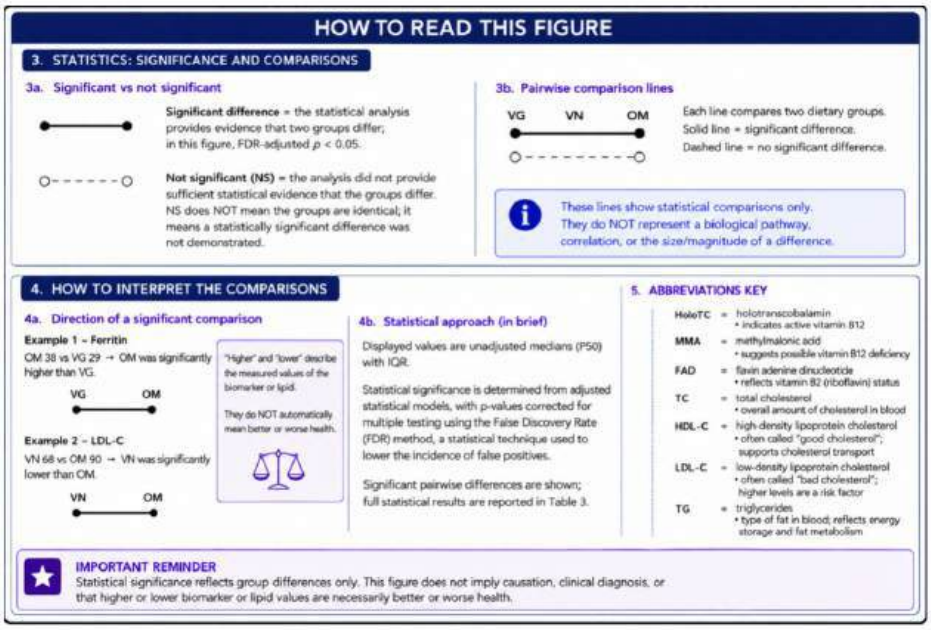

Figure 3*: A supporting diagram to interpret table 2. ***Created by ChatGPT, model 5.6 Luna.***

The following diagram (table 4) helps understand the percent of each dietary group being 'deficient' or below healthy references values from the VeChi Youth Study (or above, in the case of MMA)

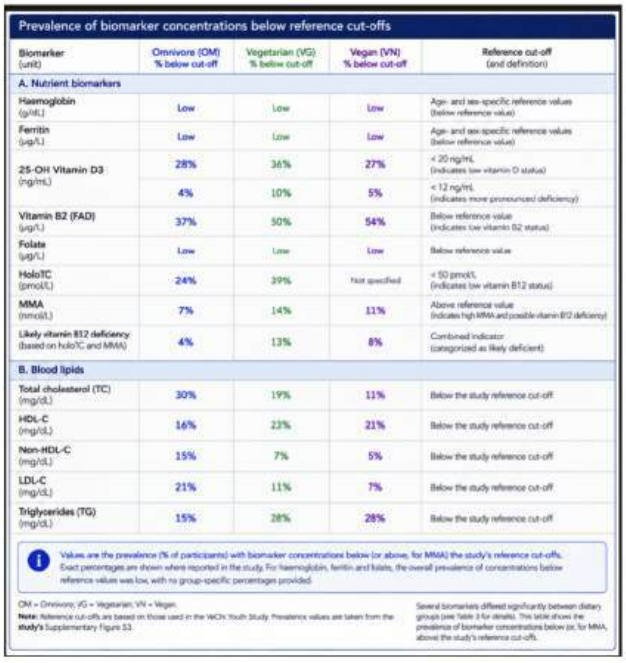

Prevalence of biomarker concentrations below reference cut-offs

Biomarker (unit)	Omnivore (OM) % below cut-off	Vegetarian (VG) % below cut-off	Vegan (VN) % below cut-off	Reference cut-off (and definition)
A. Nutrient biomarkers				
Haemoglobin (g/dL)	Low	Low	Low	Age- and sex-specific reference values (below reference value)
Ferritin (µg/L)	Low	Low	Low	Age- and sex-specific reference values (below reference value)
25-OH Vitamin D3 (ng/mL)	28%	36%	27%	< 20 ng/mL (indicates low vitamin D status)
	4%	10%	5%	< 12 ng/mL (indicates more pronounced deficiency)
Vitamin B2 (FAD) (µg/L)	37%	50%	54%	Below reference value (indicates low vitamin B2 status)
Folate (µg/L)	Low	Low	Low	Below reference value
HoloTC (pmol/L)	24%	39%	Not specified	< 50 pmol/L (indicates low vitamin B12 status)
MMA (nmol/L)	7%	14%	11%	Above reference value (indicates high MMA and possible vitamin B12 deficiency)
Likely vitamin B12 deficiency (based on holoTC and MMA)	4%	13%	8%	Combined indicator (categorized as likely deficient)
B. Blood lipids				
Total cholesterol (TC) (mg/dL)	30%	19%	11%	Below the study reference cut-off
HDL-C (mg/dL)	16%	23%	21%	Below the study reference cut-off
Non-HDL-C (mg/dL)	15%	7%	5%	Below the study reference cut-off
LDL-C (mg/dL)	21%	11%	7%	Below the study reference cut-off
Triglycerides (TG) (mg/dL)	15%	28%	28%	Below the study reference cut-off

Values are the prevalence (% of participants) with biomarker concentrations below (or above, for MMA) the study's reference cut-offs. Exact percentages are shown where reported in the study. For haemoglobin, ferritin and folate, the overall prevalence of concentrations below reference values was low, with no group-specific percentages provided.

OM = Omnivore; VG = Vegetarian; VN = Vegan.
Note: Reference cut-offs are based on those used in the VeChi Youth Study. Prevalence values are taken from the **study's** Supplementary Figure S3.

Several biomarkers differed significantly between dietary groups (see Table 3 for details). This table shows the prevalence of biomarker concentrations below (or, for MMA, above) the study's reference cut-offs.

Figure 4: Prevalence of Biomarker Concentrations Below Reference Cut Offs for the VeChi Youth Study Data Compiled from Alexy, Ute, et al, 2021, section 3.3 'Nutrient Status and Blood Lipids', Created by ChatGPT, model 5.6 Luna.

Study Limitations and Conclusions

While this study is insightful because it combines both measured micronutrient intakes and blood biomarkers for a fuller story on VG, VN, and OM health in its population sample across a relatively long period, it does not capture nutrient status for other micronutrients like omega 3 fatty acids. Moreover, it accounts for only 6-18 year olds. It also only provides, according to the study authors, a "glimpse" into the healthiness of plant based diets, due to its cross sectional design.

Relevance of the Study to Inform Nutritional Quality of the Vegan Diet

All in all, the authors suggest that the results of the study suggest that vegetarian diets, "including" vegan diets, can "meet the recommended nutrient requirements" from childhood to adolescence. Blood tests show that measuring consumed intakes is not sufficient to suggest micronutrient deficiencies. It can be interpreted that while proper planning is essential to avoid what may otherwise be common pitfalls of the VN diet, like insufficient calcium and vitamin B12 intake, there is a greater margin for underconsuming micronutrient intakes and remaining healthy than was previously highlighted in earlier discussions.

VeChi Diet Study

Of the three studies mentioned in the methodology, only one studies a population of children in the infants to toddlers age range. It is the VeChi Diet Study. It is a cross sectional study collecting data from a three day period on:

- anthropometrics (body height and weight)
- food consumption: weighted food and beverages with kitchen scales and approximate measures, estimated breast milk consumption, and simulated energy and nutrient contents of commercial food products by their labeled ingredients and content.

It includes data on covariates in the study, such as socioeconomic status, urbanicity, and physical activity of children. The study authors adjust their results on data such as total energy and protein intake with these variables to ensure the vegan diet remains fairly measured without strong confounding factors. It examines 127 vegetarian children, 139 vegan children, and 164 omnivorous children. The VeChi Diet Study's, after adjusting for confounders in covariates, key findings are:

Dietary Energy and Macronutrient Status

- Total energy intake and dietary energy density "did not differ significantly" between the three diet group samples
- OM (omnivorous) children had the "highest intake of total protein, total fat, and added sugars", and VN (vegan) children had the "highest intake of carbohydrates and fibre".
- The consumption of protein, total fat, and added sugars increased with the "inclusion" of

more animal food products, highest in omnivorous children and lowest in vegan children.
- The consumption of carbohydrates and fibre rose with the lowered inclusion of animal products, highest in vegan children and lowest in omnivorous children.
- On average, all groups still consumed 2.3 to 2.5 times more the reference value of protein in Germany of 1g of protein per kilogram of bodyweight a day.
- On average, all groups "met the German reference value" for daily fat intake.
- On average, all groups "met the reference intake" for carbohydrates.

Study limitations

The VeChi diet study was cross-sectional in design, providing, as put by its authors, only a "glance" into the "food intake" and "anthropometrics" of its sample populations. It planned to, in the future, assess micronutrient intakes with blood and urine tests. It has a risk of bias in some of its anthropometrics, because of proxy reporting of body weight and height. It also has a risk of systematic error in its method for breastmilk intake estimation, by multiplying the reported number of breast meals with median intakes of breastmilk according to the age of children. All in all, it suggests the vegan diet is capable of providing a strong macronutrient profile for vegan children, though to a lesser extent than its 'Youth Study' counterpart. Thus, it remains significant for parents to consider micronutrient risks of greater importance than the capacity to provide energy and protein density in a vegan diet.

The UCL, Child's Memorial Health Institute Study

In a cross sectional study called "Growth, body composition, and cardiovascular and nutritional risk of 5- to 10-y-old children consuming vegetarian, vegan, or omnivore diets", study collaborators in Poland examined healthy Polish children aged 5 to 10 years old. The sample populations for VN, VG, and OM groups were all of white European ethnicity, with total populations of 52, 63, 72 respectively. The study uses clear and relevant cut offs to categorize children in each of these dietary groups fairly. From September 2014 to July 2017, these 187 children recruited through advertisements and by word of mouth provided just a day's worth of data for blood testing and anthropometrics. This data was combined with personal food diaries which tracked all foods eaten and drunk over a 4 day period, which were obtained by the study authors two weeks prior to the blood tests and anthropometric measuring of the children. This methodology provided outcomes on:

- Anthropometry
- Body composition
- Bone health
- Cardiovascular disease risk markers
- Micronutrient Status

With strong statistical methods of analyses and extensive consideration of procedures used for the measuring of outcomes, the UCL and collaborators study provides an overall robust

methodology.

Results of Relevant Measured Outcomes

Macronutrient status and Vitamin B12

- From the VG and VN groups, nearly a third were not consuming either vitamin B12 supplements or vitamin B12 fortified foods. Median intake for the VG and VN groups for vitamin B12 (including supplements) was 2.4 and 1.2 micrograms respectively, compared to 2.7 on the OM diet. This study hence indicates the importance of sourcing B12 on a vegan diet in particular.
- Median protein intake was statistically significantly higher for OM than both VG and VN, at 57.3 grams versus 45.1 and 42.4 grams respectively.
- Median carbohydrate intake was statistically significantly higher for VN than both VG and OM, at 266 grams versus 243 and 216 grams respectively.
- Median fibre intake was statistically significantly higher for VN than both VG and OM, at 31.7 grams versus 24.7 and 14.9 grams respectively.
- Median fat intake was statistically significantly higher for OM than VG and VN, at 65.7 grams versus 53.5 and 51.7 grams respectively.
- Overall, DEI (dietary energy intake) did not statistically significantly differ between diet groups.

Bone Health, Cardiovascular Risk, and body Iron Status

- The total body bone mineral content excluding skulls for VN children was statistically significantly lower than the OM mean value, at -3.7% less.
- Ferritin levels for VN children were 25% lower than the OM mean, while they were 14% lower than the OM mean for VG children.
- VN children had statistically significantly lower ferritin levels than the OM mean value, at 25% lower. VG children were 14% lower but not statistically significantly lower.

Vitamin B12 Status and Markers

- All in all, VN children had lower levels of vitamin B12 levels measured in blood. Those without diet supplements and fortified foods in their diet had a mean 217.6 picomoles per liter lower than the OM mean. VN children without diet supplements but with fortified foods had a mean value of vitamin B12 blood levels 139.8 picomoles per liter lower than the OM mean value. It is not specified whether these levels meet reference age group intakes, despite being lower than the OM means.

Estimated Deficiencies among Diet Groups in Percentages for Iron and Mean Differences in Vitamin B12 Blood Levels

(data compiled from Desmond, Małgorzata A, et al. section “Prevalence of abnormal vitamin

B-12, hemoglobin, depleted iron stores, and LDL- and HDL cholesterol status", table 6)

Table 1: Iron, Ferritin Levels, less than 15 μg / L

OM: 12.8%	VG: 18.3%	VN: 30.2%

Table 2: Mean Differences of Vitamin B12 among VG and VN compared to OM mean in picomoles per liter (measured blood levels)

VG: no supplementation or fortification	VG: supplementation and fortification	VN: no supplementation or fortification	VN: fortification only
-90.9	+68.1	-217.6	-139.8

These results indicate the importance, again, of vitamin B12 supplementation or fortification on the vegan diet. It is reinforced that careful deliberation and planning is essential to avoid B12 deficiency. Higher estimated deficiencies of iron through lower ferritin levels in the VN diet are also of important concern.

Study Limitations, Conclusions, and Relevance to Informing Nutritional Quality Of VN Diet

The UCL and collaborators study provided robust methods to measure its outcomes, particularly through the use of statistical techniques to account for confounding factors. Moreover, these outcomes were broader than both VeChi studies, including more unique measures in anthropometrics and bone health. Unfortunately, its measuring of blood levels was performed across a single day period for all participants, and its overall cross sectional design limited its insightfulness, particularly in helping establish causal relations between dietary factors and health outcomes. It also ran the risk of selection bias, by using convenience sampling rather than random selection of participants. Overall, it still reinforced the importance of supplementation and fortification for the VN diet, particularly concerning vitamin B12 and potentially iron, and introduced cardiovascular biomarkers to actually highlight the association between the VN diet and positive metrics concerning the risk for cardiovascular disease. However, these findings on vitamin B12 should be viewed in the context of a study with convenience sampling: worthy of noting, but not necessarily generalizable to broader vegan child populations. The study authors also communicate the importance of further deeper study on the VN diet for children, particularly to understand the relation between plant based diets and bone health.

Preschoolers' Nutrition and Family Habits

Children aged 3 to 5 are commonly referred to as preschoolers, and guidelines from reputable organizations on their nutrition remain largely unchanged compared to toddler nutrition

(CDC) (Johns Hopkins Medicine). The 5 whole food groups previously referred to remain the staples of their healthy diet. Thus, the same 'rules' and facts aforementioned in the toddler nutrition discussion largely apply to preschooler nutrition. Outside of this discussion, however, there are some new considerations. These, however, largely do not pertain to health and nutrition facts. They are psychological, human considerations that should help parents develop healthy habits with or among their children. Scaglioni et al. (2018) describe several factors that may influence the eating behaviours of children, which include: parenting style, family structure and meal structure, among others. Outside of sourcing certain micronutrients on the vegan diet, and outside of other considerations such as energy and protein density, there are important considerations on developing healthy eating habits that will promote lifelong health. These start at home, encouraged and facilitated by parents, and are especially relevant to the vegan diet because it may be tricky to encourage picky children to eat healthily this way.

The "family environment" is known to be important for the development of children's health behaviours, though the reasons why remain unclear (Scaglioni et al.). A suspected reason is that children have the ability to observe and imitate the actions of important others, like parents or role models (Scaglioni et al.). The vegan diet, because of its restriction in food groups, should be followed and promoted in the household with positive habits, with parents taking charge of this positive 'health environment'.

Scaglioni et al. (2018) propose several influences and methods to promote this environment. One is the parental influence style, particularly the authoritative one, which is described as "demanding and responsive" but characterized by "high levels of warmth and control". This style is said to garner the "highest levels of self-efficacy, self-discipline" and "emotional maturity" of children, highly relevant characteristics for the maintenance of a healthful vegan diet. It is also suggested that common mealtimes are important in establishing healthy eating habits, where it was suggested that children and adolescents who join in fewer family meals have a greater tendency to "consume more unhealthy food" (Scaglioni et al.). In fact, there is a "positive relation between frequent family meals and greater consumption of healthy foods", which were characterized as fruits, vegetables, and calcium rich foods (Scaglioni et al.).

All in all, nutrition encompasses more than just eating the right foods. For parents, especially, it becomes a task to encourage healthy eating habits, at the right times and as a family, making sure it is not done in an excessively authoritative or strict manner.

Conclusions and Limitations

Despite handling a wide body of evidence, the data considered here still come from only cross sectional studies of cohorts of a relatively small, population-restricted nature (prone to selection bias). The UCL study assesses the biomarkers and nutrient intake of only Polish children, in the specific age group of 5 to 10 years old. The VeChi Youth Study examines children mostly belonging to middle to high socioeconomic status in Germany. Moreover, while certain trends recurred through these studies, like the highlighted importance of vitamin B12

supplementation, other seemingly statistically insignificant factors cannot be deemed entirely insignificant because of the small populations being studied. A simple example is dietary energy: while the studies suggested dietary energy needs were largely met by vegan cohorts, this does not mean that *all* vegan children are guaranteed to meet this need. Statistical insignificance cannot guarantee the absence of an outcome. Hence, the results of these studies are important, but not telling of veganism at a much larger scale - they cannot be smoothly and validly generalized to other populations of vegan children. Lastly, they largely do not provide information on causal relations between nutrition and health outcomes, instead suggesting *possibilities* of outcomes related to nutrition.

Ultimately, the vegan diet can promote health by meeting nutritional needs through deliberate planning and awareness. Without proper education on the importance of supplementation, fortification, and techniques to increase micronutrient bioavailability, families likely cannot properly plan a vegan diet to ensure the health of their children. In particular, children following a vegan diet should be backed by vitamin B12, iron, omega 3 fatty acids, zinc, and potentially calcium either supplementation, fortification, or micronutrient rich whole foods. Moreover, household habits matter, because they promote healthy behaviours among growing children to ensure these are carried into later life. A simple example is having common meals as a family.

In essence, the vegan diet can be healthful for children so long as it is promoted by educated, aware parents, ensuring common pitfalls like B12 deficiency are avoided through adequate planning and family habits.

Works Cited

Alexy, Ute, et al. "Nutrient Intake and Status of German Children and Adolescents Consuming Vegetarian, Vegan or Omnivore Diets: Results of the VeChi Youth Study." Nutrients, vol. 13, no. 5, 18 May 2021, p. 1707, 10.3390/nu13051707.

Bailey, Eileen, et al. "Comparative Bioavailability of DHA and EPA From Microalgal and Fish Oil in Adults." International Journal of Molecular Sciences, vol. 26, no. 19, 24 Sept. 2025, https://www.mdpi.com/1422-0067/26/19/9343, 10.3390/ijms26199343. Accessed 29 Aug. 2026.

Baltussen, Rob, et al. "Iron Fortification and Iron Supplementation Are Cost-Effective Interventions to Reduce Iron Deficiency in Four Subregions of the World." The Journal of Nutrition, vol. 134, no. 10, 1 Oct. 2004, pp. 2678–2684, https://academic.oup.com/jn/article/134/10/2678/4688367, 10.1093/jn/134.10.2678. Accessed 22 Aug. 2026.

CDC. "Zinc." 11 Dec. 2024, https://www.cdc.gov/infant-toddler-nutrition/vitamins-minerals/zinc.html

Chelsea. "Veganuary 2025 Results Explained." Veganuary, 3 Feb. 2025, https://veganuary.com/veganuary-2025-results/.

Claydon, Rachel. "Are Sustainable Foods Accessible for Everyone?" 2025, https://foodfoundation.org.uk/news/are-sustainable-foods-accessible-everyone

Cleveland Clinic. "Omega-3 Fatty Acids." 17 Nov. 2022, https://my.clevelandclinic.org/health/articles/17290-omega-3-fatty-acids

Clinic, Cleveland. "Fat-Soluble Vitamins: What They Are and How to Get the Most out of Them." Cleveland Clinic, 17 May 2023, https://health.clevelandclinic.org/fat-soluble-vitamins. Accessed 24 Aug. 2026.

Dahdouh, Sergio, et al. "Development of the FAO/INFOODS/IZINCG Global Food Composition Database for Phytate." Journal of Food Composition and Analysis, vol. 78, 2 Feb. 2019, pp. 42–48, https://www.sciencedirect.com/science/article/pii/S0889157518304496, 10.1016/j.jfca.2019.01.023. Accessed 22 Aug. 2026.

D'Andrea, Astrid E., et al. "A Comparison of the Nutritional Profile and Nutrient Density of Commercially Available Plant-Based and Dairy Yogurts in the United States." Frontiers in Nutrition, vol. 10, 25 May 2023, https://pmc.ncbi.nlm.nih.gov/articles/PMC10248066/#sec21, 10.3389/fnut.2023.1195045. Accessed 28 July 2026.

Desmond, Malgorzata A., et al. "Plant-Based Diets in Children: Secular Trends, Health Outcomes, and a Roadmap for Urgent Practice Recommendations and Research—a Systematic Review." Nutrients, vol. 16, no. 5, 1 Mar. 2024, p. 723, https://pmc.ncbi.nlm.nih.gov/articles/PMC10934552/#sec2-nutrients-16-00723, 10.3390/nu16050723. Accessed 14 Aug. 2026.

Desmond, Małgorzata A, et al., 33740036. "Growth, Body Composition, and Cardiovascular and Nutritional Risk of 5- to 10-Y-Old Children Consuming Vegetarian, Vegan, or Omnivore

Diets.” The American Journal of Clinical Nutrition, vol. 113, no. 6, 1 June 2021, pp. 1565–1577. PubMed Central, https://pubmed.ncbi.nlm.nih.gov/33740036/, 10.1093/ajcn/nqaa445. Accessed 2 Sept. 2026.

Grand View Research. “Plant-Based Yogurt Market Size to Reach $7.53Bn By 2030.” July 2024, https://www.grandviewresearch.com/press-release/global-plant-based-yogurt-market

Hall, Andrew G., and Janet C. King. “Zinc Fortification: Current Trends and Strategies.” Nutrients, vol. 14, no. 19, 21 Sept. 2022, p. 3895, 10.3390/nu14193895, https://www.mdpi.com/2072-6643/14/19/3895

Johns Hopkins Medicine. “Iron-Deficiency Anemia.” Iron-Deficiency Anemia, 2019, https://www.hopkinsmedicine.org/health/conditions-and-diseases/irondeficiency-anemia. Accessed 22 Aug. 2026.

National Academies. “Dietary Reference Intakes for Vitamin A, Vitamin K, Arsenic, Boron, Chromium, Copper, Iodine, Iron, Manganese, Molybdenum, Nickel, Silicon, Vanadium, and Zinc (2001) .” Nationalacademies.Org, National Academy Press, 2025, https://www.nationalacademies.org/read/10026/chapter/1. Accessed 21 Aug. 2026.

National Institutes of Health. “Office of Dietary Supplements - Calcium.” 6 Oct. 2022, https://ods.od.nih.gov/factsheets/Calcium-HealthProfessional/

“Office of Dietary Supplements - Vitamin A.” Ods.Od.Nih.Gov, 10 Mar. 2025, https://ods.od.nih.gov/factsheets/VitaminA-HealthProfessional/#h2. Accessed 24 Aug. 2026.

“Office of Dietary Supplements - Zinc.” Nih.Gov, 2017, https://ods.od.nih.gov/factsheets/Zinc-HealthProfessional/#h4. Accessed 22 Aug. 2026.

“Vitamin C.” National Institutes of Health, National Institutes of Health, 26 Mar. 2021, https://ods.od.nih.gov/factsheets/VitaminC-HealthProfessional/. Accessed 9 July 2026.

National Institutes of Health. “Office of Dietary Supplements - Folate.” Nih.Gov, National Institutes of Health, 2022, https://ods.od.nih.gov/factsheets/Folate-HealthProfessional/. Accessed 29 Aug. 2026.

“Office of Dietary Supplements - Iron.” Ods.Od.Nih.Gov, 15 June 2023, https://ods.od.nih.gov/factsheets/Iron-HealthProfessional/#h3. Accessed 2 Sept. 2026.

“Omega-3 Fatty Acids.” Nih.Gov, 15 Feb. 2023, https://ods.od.nih.gov/factsheets/Omega3FattyAcids-HealthProfessional/. Accessed 26 Aug. 2026.

Ramsing, Rebecca, et al. “Dairy and Plant‐Based Milks: Implications for Nutrition and Planetary Health.” Springer Nature, vol. 10, 2023, pp. 291–302, 10.1007/s40572-023-00400-z.

Sadeeh, Randa, and José Martines. Complementary Feeding, Family Foods for Breastfed Children. 2000, https://iris.who.int/server/api/core/bitstreams/10a51d17-3d8d-47a6-8780-c5faed918274/content. Accessed 21 Aug. 2026.

Shkembi, Blerina , and Thom Huppertz. “Calcium Absorption from Food Products: Food Matrix Effects.” Nih.Gov, 2026,

https://pmc.ncbi.nlm.nih.gov/articles/PMC8746734/#sec7-nutrients-14-00180. Accessed 17 Aug. 2026.

The Vegan Society. “Definition of Veganism.” The Vegan Society, 2025, https://www.vegansociety.com/go-vegan/definition-veganism.

Weder, Stine, et al. “Energy, Macronutrient Intake, and Anthropometrics of Vegetarian, Vegan, and Omnivorous Children (1–3 Years) in Germany (VeChi Diet Study).” Nutrients, vol. 11, no. 4, 12 Apr. 2019, p. 832, 10.3390/nu11040832.

World Health Organization. “Infant and Young Child Feeding.” World Health Organization, 20 Dec. 2023, https://www.who.int/news-room/fact-sheets/detail/infant-and-young-child-feeding.

Exploring ADHD Neurochemistry, Pharmacological, and Non-Pharmacological Interventions: A Literature Review By Gaven Z Li

Abstract

Attention Deficit Hyperactivity Disorder (ADHD) is a neurodevelopmental condition that can affect attention, activity level, and impulse control. The neurochemical basis of ADHD is often associated with an altered regulation of the catecholamines (e.g., dopamine and norepinephrine) in brain regions responsible for executive functioning. Factors contributing to these neurochemical differences may include genetics, exposure to toxic chemicals, and nutritional factors. This article discusses two categories of medications: stimulant and non-stimulant medications. These drugs help regulate the levels of dopamine and norepinephrine in the brain, providing an effective treatment option for ADHD. In addition, nonpharmacological approaches such as physical activity, nutritional interventions, and behavioral therapy are also discussed in this paper. These approaches may help improve executive functioning and reduce the symptoms of ADHD. In this review, I compare the benefits and limitations of different treatment options. However, for each individual with ADHD, the best treatment approach is one that provides sufficient symptom control for the needed length of time while causing the fewest troublesome side effects.

Keywords Attention Deficit Hyperactivity Disorder, Neurochemical Imbalance, Pharmacological Intervention, Non-pharmacological Intervention.

1. Introduction

Inattention, hyperactivity, and impulsivity are commonly associated with Attention Deficit Hyperactivity Disorder (ADHD), a neurodevelopmental condition (Casella and Casella). Inattention involves difficulty sustaining focus, staying organized, and resisting distractions, while hyperactivity refers to excessive movement or restlessness. Impulsivity, on the other hand, involves acting or responding without fully considering possible consequences. Such behavior can result in weakened performance in areas requiring mental focus, like academics, driving, socializing, and drug use (Childress and Berry; Dopheide).

The condition known as ADHD has a high global and national prevalence rate in children and young adults. For example, ADHD affects about 5.9% to 7.1% of children and 1.2% to 7.3% of adults globally (Huss et al.). The Centers for Disease Control and Prevention estimates that about 11.7% of children aged 3-17 and 6% of adults in the United States of America (US) have ADHD (CDC). For children specifically, more boys (13%) than girls (7%) are frequently diagnosed with ADHD, with substantial variations in racial diagnosis: Asians (3%), Native Hawaiians/Pacific Islanders (6%), Hispanics and Native Americans (8%), Blacks and Whites (11%) (CDC).

While more people are being diagnosed with ADHD and ADHD medications are becoming increasingly common, the causes of ADHD remain complex, with factors such as

genetics, environment, and certain childhood experiences or injuries playing pivotal roles in the condition's development. These factors heighten the risk of ADHD symptomatology, including decreased absorption of important neurotransmitters like catecholamines, dopamine, and norepinephrine (Sharma and Couture). The condition is also linked to slower development in the brain's prefrontal areas (Shaw et al.), leading to deficits in advanced executive functions such as perception, planning, organization, behavioral control, motor control, and cognition (Castellanos and Proal).

As for treatment, the general methods used involve drug therapy, behavioral therapy, or a combination of the two (Sharma and Couture). Drug therapy appears to be much more effective than behavioral therapy, however (Sharma and Couture), and as a result, ADHD treatment is heavily focused on medication. The type of treatment that is best for someone with ADHD is highly subjective since each drug has its benefits and setbacks. In this article, I will discuss the benefits and setbacks of these medications.

I completed this review article through a literature search on Google Scholar. The main keywords used included “neurochemical imbalance ADHD,” “pharmacological intervention ADHD,” and “Adderall ADHD.” I identified the articles most relevant to my topic from the Google Scholar database. I downloaded and analyzed these articles to better understand ADHD’s neurochemical basis and common medication treatments.

One major motivation for this work stems from my personal experience with the conditions examined in this review. These experiences helped me identify several gaps in the existing literature and highlighted important factors to consider when interpreting and comparing findings across previous studies. For example, most previous studies on ADHD have been written from a medical or clinical perspective, with an emphasis on understanding the causes of ADHD, improving diagnosis, and identifying effective treatments for particular groups of patients. However, ADHD also needs to be considered from the perspectives of patients, families, educators, and policymakers. In real-life situations, diagnosis and treatment can be more complicated than they may appear in controlled studies. Some individuals may be misdiagnosed, while others may live with ADHD for many years before receiving a diagnosis. Of all the adults diagnosed with ADHD, about 55.9% were diagnosed after the age of 18. Many people with undiagnosed ADHD believe that their struggles are normal until they are aware of the condition (Staley et al.). ADHD can also occur alongside other conditions, such as autism, anxiety, or depression. About 30% of kids and 50% of adults with ADHD may also have an anxiety disorder, and the percentage rises to about 67% for other conditions (CHADD).

Furthermore, about 38% to 40% of individuals on the autism spectrum meet the requirements for receiving an ADHD diagnosis (Rong et al.). These factors can complicate both diagnosis and treatment. These issues suggest that reported ADHD data may not always fully reflect the complexity of patients’ experiences. More research is also needed that considers long-term treatment experiences, co-occurring conditions, and how patients respond to different medications and interventions.

This comprehensive literature review aims to help address these gaps by examining

ADHD neurochemistry, pharmacological, and non-pharmacological treatments from a clinical perspective, while drawing on personal experience with ADHD to inform the review. Rather than identifying one medication as “best” for everyone, the review compares the effectiveness, benefits, limitations, and potential side effects of major pharmaceutical treatments while considering differences among individuals and patient groups. It also examines challenges related to diagnosis, co-occurring conditions, and treatment monitoring. By bringing together clinical research with issues that directly affect people living with ADHD, this review seeks to make the existing evidence more understandable and useful to a broader audience. Ultimately, the goal is to help patients and families become better informed about ADHD and its treatment so that they can participate more effectively in discussions and decisions with qualified medical professionals.

2.1 Neurochemical Imbalance of ADHD

As stated before, those affected by ADHD have a decreased absorption of certain neurotransmitters in the prefrontal areas of the brain. Specifically, the prefrontal cortex, corpus striatum, and cerebellum struggle to receive dopamine and norepinephrine, two types of neurotransmitters (Arnsten and Pliszka; Del Campo et al.). These chemicals and their receptors all contribute to the condition. Research suggests that the underlying neurochemical malfunctioning in ADHD is not limited to the role of one chemical in the brain but also involves how different chemicals interact with each other (Arnsten and Pliszka). It is also known that for those with ADHD, the brain areas mentioned mature more slowly (Shaw et al.) and have less volume than for neurotypicals (Berquin et al.; Castellanos et al.; Seidman et al.).
ADHD symptoms prominently implicate the prefrontal cortex, corpus striatum, and cerebellum. These brain regions support executive functions such as attention, emotional regulation, and time management, so people with ADHD have difficulty controlling these functions (Arnsten and Pliszka; Kesner and Churchwell).

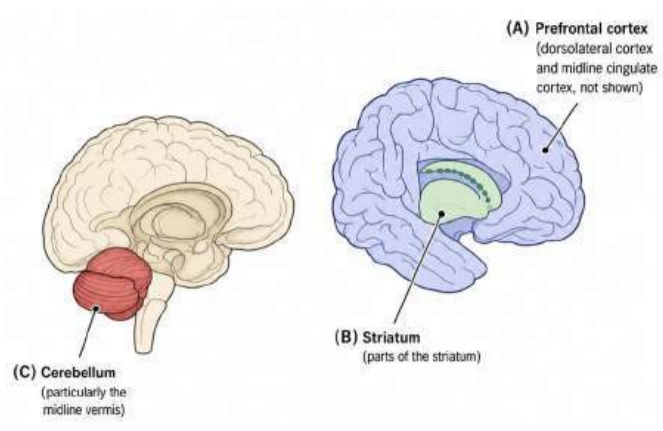

Figure 1. Simplified schematic of three brain hubs implicated in ADHD: prefrontal cortex, striatum/corpus striatum, and cerebellum. These interconnected regions contribute to neural

networks that regulate attention, executive control, and action. The prefrontal component includes regions such as the dorsolateral prefrontal cortex and midline cingulate cortex (not shown). In contrast, the striatal and cerebellar components include selected regions of the striatum and the cerebellar vermis, respectively. Created by author with assistance from OpenArt AI, based on (Shaw).

ADHD is associated with dysregulation of three key neurotransmitters: dopamine, norepinephrine, and serotonin. Dopamine supports focus and motivation and is implicated in reward, while norepinephrine helps with alertness, cognitive processing, and working memory. Serotonin, on the other hand, regulates mood, emotion, anxiety, sleep, and circadian rhythm (Del Campo et al.). Of the three neurotransmitters mentioned, dopamine and norepinephrine are the main catecholamines, a class of hormones defined by their molecular structure (Cleveland Clinic, "Dopamine"). High or low levels of dopamine can cause issues with attention and movement, such as ADHD, Parkinson's disease, and restless leg syndrome (Cleveland Clinic, "Dopamine"). Additionally, low levels of norepinephrine and serotonin can contribute to ADHD, as well as depression, anxiety, and other conditions related to mental functioning (Cleveland Clinic, "Serotonin"; Cleveland Clinic, "Norepinephrine")

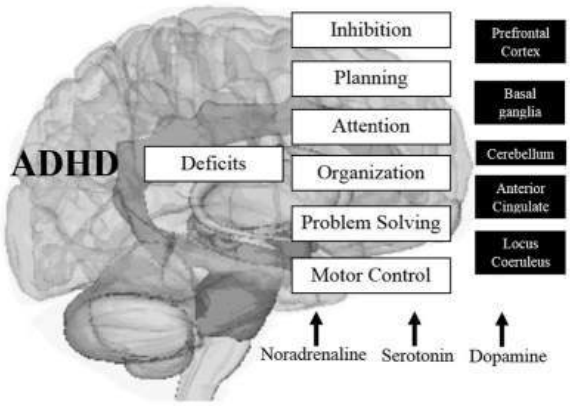

Figure 2. This figure illustrates key brain regions linked to prefrontal cortical activity and the executive processes they support. Reproduced from (Martínez et al.), Frontiers in Psychology, licensed under CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/).

ADHD has various causes, and there isn't really one factor that could be pinned as the sole cause of the condition. However, there is reason to believe that genetics play a role in the condition's existence. In twins, for example, if one twin has ADHD, then the chance of the other twin sharing the same condition is approximately 76% (Larsson et al.). Moreover, around 40% of parents with ADHD have children who show ADHD-related behavior (Uchida et al.). These statistics provide robust evidence for genetic involvement in the condition. Specifically, several genes contribute to behavior consistent with ADHD diagnostic criteria, particularly dopamine-related genes such as DRD4, DRD5, and DAT-1, which have an anomalous shape

(Faraone and Larsson). These “flawed” genes could cause the dopamine receptors in the brain to be less effective, ultimately resulting in ADHD traits like inattention and hyperactivity (Faraone and Larsson). In addition, norepinephrine also contributes to the condition if genes that control this hormone are impaired, with impaired function of the alpha-2A adrenergic receptor (ADRA2A), a receptor responsible for stabilizing norepinephrine levels, being particularly implicated in ADHD (Berridge and Devilbiss; Vanicek et al.). Irregularities in genes responsible for serotonin and even the nicotinic acetylcholine system contribute to behavior consistent with ADHD’s diagnostic criteria (Oades et al.; Potter et al.). Because many genes control these factors, the genetic roots of this condition remain an esoteric topic (Faraone and Larsson; Tistarelli et al.).

The environment someone grew up in can also change or exacerbate their neurochemistry to the point that it mimics symptoms of inattention and impulsivity. For example, lead poisoning has been shown to increase the likelihood of children developing ADHD (Harshitha et al.) due to its ability to mimic calcium in the body (Harshitha et al.). Calcium is important for many functions, including neurological development, and lead can compete with and replace calcium in proteins like kinase C. Lead replacement in kinase C, specifically, alters how endothelial brain cells behave, disrupting the blood-brain barrier (Collin et al.). This phenomenon ultimately contributes to disrupted neural connection formation (Mousa et al.), which could lead to irreversible damage to a child's nervous system, since they form millions of these connections per second (Tau and Peterson). Moreover, childhood injuries such as head trauma can cause permanent damage to one’s neurochemistry. This type of injury, especially when it affects the prefrontal cortex, can immediately disrupt how the brain produces and uses dopamine and norepinephrine, the primary chemical messengers responsible for motivation, focus, and impulse control (Asarnow et al.). There is statistical proof of this phenomenon: according to a 2018 study published in JAMA Pediatrics, children who have a history of traumatic head injuries met the diagnostic criteria of ADHD more often than those who were much safer (Asarnow et al.).

Along with that, nutritional deficiencies (e.g., iron, zinc, magnesium, omega-3 fatty acids) and dietary and metabolic factors (e. g., blood sugar swings, gut-brain axis) can cause neurochemical imbalances that alter brain function and worsen or contribute to ADHD symptoms (Hunter et al.). For example, iron is a vital cofactor required to synthesize dopamine. Low ferritin levels (the amount of iron stored in the blood) reduce dopamine production, impairing attention and focus (Hunter et al.). Additionally, zinc helps regulate dopamine transporter function and protects against oxidative stress. A copper-zinc imbalance is frequently observed in individuals with ADHD, disrupting both dopamine and norepinephrine pathways (Hunter et al.). Furthermore, emerging research indicates that gut health and intestinal permeability influence systemic inflammation and neurotransmitter production (Hunter et al.).

2.2 Pharmacological treatment for ADHD

In medicine, there are two types of medications approved for the treatment of ADHD: stimulants and non-stimulants (FDA). Stimulant drugs act immediately and directly in the brain,

while non-stimulant drugs usually have slower effects on ADHD treatment (Greenhill et al.; Wernicke and Kratochvil). Stimulants directly increase dopamine and norepinephrine activity, while non-stimulants increase dopamine indirectly by affecting other neurotransmitters and receptors (Cortese, Adamo, et al.). Due to these differences, these two types of medications have their own unique side effects that heavily affect one's preference for the choice of ADHD treatment (Nutt et al.). In this section, I will explore how effective these medications are and what types of people normally take them.

2.2.1 Stimulant Drugs

Adderall is a mixture of dextroamphetamine and levoamphetamine salts in a 3:1 ratio (DailyMed-Adderall). Adderall was introduced by Richwood Pharmaceuticals, which received FDA re-approval to market it for ADHD in 1996; Richwood was later acquired by Shire Pharmaceuticals in 1997 (Swanson and Volkow). Amphetamines are thought to block the reuptake of norepinephrine and dopamine and increase the release of these neurotransmitters into the extraneuronal space (DailyMed-Adderall). Amphetamines are highly effective for reducing ADHD symptoms in the short term, with evidence supporting them as a preferred first-choice medication in adults. In contrast, methylphenidate is preferred in children and adolescents (Cortese, Adamo, et al.). Adderall can cause side effects, including decreased appetite, weight loss, gastrointestinal disturbances, irritability, and increases in heart rate or blood pressure (DailyMed-Adderall). Adderall also has a high potential for abuse and misuse, which can lead to addiction, and it may produce physical dependence (DailyMed-Adderall).

Ritalin is the brand name for methylphenidate hydrochloride, also called methyl 2-phenyl-2-piperidin-2-ylacetate hydrochloride in International Union of Pure and Applied Chemistry (IUPAC) notation (DailyMed-Ritalin; PubChem). Ciba developed methylphenidate, and Ritalin received its initial U.S. approval in 1955 (DailyMed-Ritalin). It primarily blocks the reuptake of dopamine and norepinephrine, allowing for increased signaling between brain regions responsible for executive functioning when these neurotransmitters are involved (DailyMed-Ritalin). Methylphenidate is effective for reducing ADHD symptoms in the short term, although amphetamines showed greater efficacy in the comparative analysis by Cortese et al. (Cortese, Adamo, et al.). Side effects include decreased appetite, weight loss, insomnia, headache, anxiety, nausea, abdominal pain, increases in heart rate or blood pressure, psychiatric reactions, cardiovascular effects, growth suppression, misuse, and dependence (DailyMed-Ritalin).

Although Adderall and Ritalin are both stimulant medications that increase dopamine and norepinephrine signaling, they do so somewhat differently. Amphetamines such as Adderall both block the reuptake and increase the release of dopamine and norepinephrine, while methylphenidate primarily blocks their reuptake (DailyMed-Adderall; DailyMed-Ritalin). Both are effective for ADHD, although short-term clinical evidence generally shows greater efficacy for amphetamines than for methylphenidate (Cortese, Adamo, et al.).

2.2.2 Non-Stimulant Drugs

Several non-stimulant drugs have been approved as treatment options for ADHD, with one of the first ones being Strattera. Strattera is the brand name for atomoxetine hydrochloride, developed by Eli Lilly and FDA-approved for ADHD in 2002 (FDA Pediatric Labeling–Strattera). It is thought to work by selectively inhibiting the presynaptic norepinephrine transporter, although its precise mechanism in ADHD is unknown (DailyMed-Strattera). Unlike stimulant medications, atomoxetine is not associated with stimulant or euphoric properties (DailyMed-Strattera). Clinical trials found that once-daily administration of atomoxetine is an effective treatment for children and adolescents with ADHD (Michelson et al.). Common side effects include nausea, vomiting, fatigue, decreased appetite, abdominal pain, and somnolence, while important risks include suicidal thoughts in younger patients and increases in heart rate or blood pressure (DailyMed-Strattera).

Another non-stimulant medication widely used for the treatment of ADHD would be Qelbree, which is the brand name for an extended-release formulation of viloxazine hydrochloride. In 2021, Supernus Pharmaceuticals received FDA approval for ADHD (FDA Qelbree Review). It is thought to work by inhibiting norepinephrine reuptake, although its exact mechanism in ADHD is unclear (DailyMed-Qelbree). Clinical trials in children found that viloxazine significantly reduced ADHD symptoms compared with placebo (Nasser et al.). Common side effects include drowsiness, decreased appetite, fatigue, nausea, vomiting, insomnia, and irritability, and Qelbree also carries a warning for suicidal thoughts and behaviors (DailyMed-Qelbree).

Intuniv is another example of a non-stimulant method of mitigating the symptoms of ADHD. This drug is an extended-release formulation of guanfacine developed for ADHD by Shire Pharmaceuticals and FDA-approved for this use in 2009 (FDA Intuniv Clinical Review). Guanfacine is a selective central α2A-adrenergic receptor agonist, although its exact mechanism in treating ADHD is not known (DailyMed-Intuniv/Guanfacine). Clinical trials found that guanfacine extended-release was effective in reducing symptoms of ADHD. At the same time, adverse events were mild to moderate, did not interfere with improvements in attention, and rarely led to discontinuation (Sallee et al.). Common side effects include somnolence, fatigue, hypotension, nausea, dizziness, and abdominal pain, while bradycardia can occur and sudden discontinuation can cause rebound hypertension (DailyMed-Intuniv/Guanfacine).

Kapvay is the final drug on this list that was approved for ADHD treatment. This medication is an extended-release formulation of clonidine hydrochloride sponsored by Shionogi Pharma and FDA-approved in 2010 for the treatment of ADHD as monotherapy or as adjunctive therapy to stimulant medications (FDA Kapvay). Clonidine stimulates α2-adrenergic receptors in the brain, although its exact mechanism for treating ADHD is not known (DailyMed-Clonidine). Clinical trials in children and adolescents found that extended-release clonidine significantly improved ADHD symptoms compared with placebo (Jain et al.). Common side effects include somnolence, fatigue, irritability, nightmares, insomnia, constipation, and dry mouth, while hypotension, bradycardia, and rebound hypertension are important concerns

(DailyMed-Clonidine).

Although Qelbree, Strattera, Intuniv, and Kapvay are all non-stimulant treatments for ADHD, they affect noradrenergic signaling in different ways. Qelbree is thought to inhibit norepinephrine reuptake, while Strattera selectively inhibits the presynaptic norepinephrine transporter (DailyMed-Qelbree; DailyMed-Strattera). In contrast, Intuniv acts as a central α2A-adrenergic receptor agonist, while Kapvay stimulates α2-adrenergic receptors in the brain (DailyMed-Clonidine; DailyMed-Intuniv/Guanfacine).

2.3. Non-pharmacological Interventions

ADHD cannot be explained simply as having "too much" or "too little" of a single neurotransmitter (Pliszka et al.). Dopamine and norepinephrine are important for higher-level executive functions, such as attention, planning, and self-control, which are often affected in ADHD (Del Campo et al.). Because of these factors, some non-drug approaches may help support brain function or help people compensate for attention and executive-function difficulties.

2.3.1 Physical Activity and Exercise

Exercise is one non-drug approach that may be connected to ADHD biology. Stimulant medications (Wigal et al.) target the same dopamine and norepinephrine systems that exercise affects.

- Aerobic Exercise: Activities such as running, swimming, and cycling can improve executive functions, including inhibitory control, working memory, and cognitive flexibility in children with ADHD (Yang et al.).
- Skill-Based Exercise: Exercises that require quick reactions, decision-making, and coordination can improve executive functioning, particularly inhibitory control and cognitive flexibility, in people with ADHD (Qiu et al.).

In my experience, physical activity has almost always positively affected my focus and mood. Both aerobic exercise and skill-based exercise have helped me concentrate better, release extra energy, and feel more mentally balanced afterward. Personally, I tend to prefer skill-based activities because they involve movement, coordination, quick reactions, and decision-making, which keep my brain more engaged. Team-based activities can also be especially meaningful because collaboration and working toward a shared goal can build strong bonds and friendships that are difficult to find in other settings. However, this is also a personal preference. Different people with ADHD may respond better to different types of exercise, so the most important thing may be finding a form of physical activity that is enjoyable, social, and realistic to continue.

2.3.2 Dietary and Nutritional Approaches

Diet can affect the body and brain, including the nutrients used to make neurotransmitters. However, current evidence is not strong enough to recommend diet changes as

a general treatment for ADHD (Uldall Torp and Thomsen). Because nutrition can affect brain function, researchers have studied specific nutrients to see whether they may play a supporting role in ADHD management.

2.3.2.1 Protein and Tyrosine

Protein provides amino acids such as tyrosine, which the brain uses as a precursor to produce dopamine and norepinephrine (Fernstrom and Fernstrom). Still, this does not mean that simply eating more protein or tyrosine can treat ADHD. It is better to describe nutrition as one possible supporting factor, not a cure.

2.3.2.2 Omega-3 Fatty Acids

Omega-3 fatty acids can change properties of brain cell membranes, which can affect dopamine and serotonin neurotransmission. A meta-analysis found that omega-3 supplementation had a small but significant effect on ADHD symptoms, but the effect was modest compared with standard ADHD medications (Bloch and Qawasmi). These findings suggest that omega-3 fatty acids may provide a small benefit for some people with ADHD, but they should not be considered a replacement for established ADHD treatments.

2.3.2.3 Micronutrients

Iron and zinc support brain function and neurotransmitter-related processes, including dopamine biology (Granero et al.). However, evidence for using iron, zinc, or magnesium supplements as ADHD treatments is still limited. Lower nutrient levels cannot be said to cause ADHD, but correcting a real deficiency may be reasonable in some cases (Hariri and Azadbakht; Villagomez and Ramtekkar). Therefore, micronutrient supplements may be most useful when someone has a confirmed deficiency rather than as a general treatment for ADHD.
Medical professionals have not always systematically discussed diet and nutrition as part of ADHD care, based on my experience. For that reason, many families end up exploring questions about food, nutrients, sleep, energy, and daily routines mostly on their own. I think nutrition is an important area to explore more carefully because the body and brain are connected, and small changes may affect focus, mood, or energy for some people. At the same time, it is important not to overstate the evidence. Nutrition is not a universal cure for ADHD, and what helps one person may hurt another. Any dietary or supplement changes should be considered person by person and case by case.

2.3.3 Psychological and Brain-Training Approaches

Psychological treatments for ADHD usually do not “fix” brain chemistry directly. Instead, they teach skills and strategies that reduce the daily impact of ADHD symptoms. The Australian ADHD Clinical Practice Guideline describes cognitive-behavioral approaches as tools that can help reduce the day-to-day effects of ADHD and improve functioning and well-being (Australian ADHD Clinical Practice Guideline).

- Cognitive Behavioral Therapy (CBT): ADHD-focused CBT teaches skills such as organization and planning, reducing distractions, and changing unhelpful/negative thinking patterns. One clinical trial found that CBT improved ADHD symptoms in adults who still had symptoms even while taking medication (Safren et al.).
- Neurofeedback: Neurofeedback uses real-time information about a person's brain-wave activity to help them learn better control over certain patterns of brain activity (Omejc et al.). However, a meta-analysis found that neurofeedback did not significantly improve ADHD symptoms when outcomes were measured using probably blinded ratings, so well-controlled evidence did not support it as an effective ADHD treatment (Cortese, Ferrin, et al.).
- Mindfulness: Mindfulness-based therapies may help some people with ADHD become more aware of attention, impulses, and emotions. A meta-analysis found improvements in inattention and hyperactivity/impulsivity, although more high-quality studies are still needed (Cairncross and Miller).

Many therapists, schools, social workers, psychologists, and support groups use CBT-based strategies and mindfulness as important tools to help people with ADHD manage attention, impulsivity, organization, and emotions. However, these strategies are usually part of a larger support plan, not a cure on their own.

3.1 Reflections and Implications

Overall, there is no single ADHD medication or treatment plan that is best for every person. Short-term evidence supports methylphenidate as a preferred first medication for many children and adolescents and amphetamines for many adults, although long-term evidence is still more limited (Cortese, Adamo, et al.).

This matches my experience with ADHD treatment. Over the years, I have personally tried several medications, including Intuniv, Strattera, Adderall, and Ritalin, with Adderall being the last medication I used. Some medications helped me in certain ways, but I also experienced side effects and challenges. During my teenage years, growth spurts and body changes made treatment even more complicated. About seven months ago, I stopped taking medication completely and shifted my focus toward physical activity and exercise, dietary and nutritional balance, Cognitive Behavioral Therapy (CBT), and mindfulness. These approaches have positively improved my daily life, focus, and mood. My experience does not prove what is best for everyone, but it shows why ADHD treatment should be flexible and regularly re-evaluated. In the future, I will continue monitoring myself and seeking other supportive approaches if needed. Rather than trying to find one medication that perfectly “corrects” ADHD, the most practical goal may be to find an individualized combination of supports that improves daily functioning while causing the fewest unacceptable side effects.

3.2 Limitations & Future Directions

ADHD treatment research still has several important limitations. Many randomized medication trials are short-term. One major review found enough evidence to compare ADHD medications at about 12 weeks, but there was not enough evidence to make reliable comparisons at 26 or 52 weeks (Cortese, Adamo, et al.). Research quality is another concern. A review of 81 adult ADHD medication trials found that many studies had possible sources of bias, long-term side effects were not well reported, and many trials did not fully represent real-life ADHD patients with other psychiatric conditions (Elliott et al.). Future studies should follow patients for longer periods, include people who better represent real ADHD populations, and examine not only symptom improvement but also side effects, quality of life, and daily functioning.

Another important future direction is to study how different treatments can work together, instead of only comparing one treatment against another. The 14-month Multimodal Treatment Study of Children with ADHD compared medication management, behavioral treatment, combined treatment, and community care. Medication treatment alone was particularly effective for core ADHD symptoms, while combined treatment with behavioral interventions showed some additional benefits in broader areas of functioning (Hinshaw et al.; MTA Cooperative Group). Future research could examine when to add, stop, or switch between medication and non-pharmacological approaches such as behavioral treatment, exercise, mindfulness, dietary support, and school/work accommodations, depending on each person's response.

This individualized approach also connects to the idea of neurochemical imbalance discussed throughout this paper. ADHD cannot be explained as simply having too much or too little of one neurotransmitter. However, dopamine and norepinephrine signaling are still important parts of ADHD and executive functioning (Del Campo et al.; Pliszka et al.). Medications can directly affect these neurotransmitter systems, while behavioral, psychological, educational, and lifestyle approaches can help with other parts of ADHD functioning. Clinical guidelines also support combining medication with behavioral and educational interventions when appropriate, while regularly adjusting treatment based on benefits and side effects (NICE; Wolraich et al.).

Work Cited

Arnsten, Amy F. T., and Steven R. Pliszka. "Catecholamine Influences on Prefrontal Cortical Function: Relevance to Treatment of Attention Deficit/Hyperactivity Disorder and Related Disorders." Pharmacology Biochemistry and Behavior, Cognitive Enhancers for the Treatment of Neuropsychiatric Disorders: Clinical and Preclinical Investigations, vol. 99, no. 2, Aug. 2011, pp. 211–16. ScienceDirect, https://doi.org/10.1016/j.pbb.2011.01.020.

Asarnow, Robert F., et al. "Association of Attention-Deficit/Hyperactivity Disorder Diagnoses With Pediatric Traumatic Brain Injury: A Meta-Analysis." JAMA Pediatrics, vol. 175, no. 10, Oct. 2021, pp. 1009–16. Silverchair, https://doi.org/10.1001/jamapediatrics.2021.2033.

Australian ADHD Clinical Practice Guideline. 4.2 Cognitive-Behavioural Intervention Approaches -Australian ADHD Clinical Practice Guideline. 24 Sept. 2022, https://adhdguideline.aadpa.com.au/non-pharmacological/cognitive-behavioural-interventions/.

Berquin, P. C., et al. "Cerebellum in Attention-Deficit Hyperactivity Disorder." Neurology, vol. 50, no. 4, Apr. 1998, pp. 1087–93. neurology.org (Atypon), https://doi.org/10.1212/WNL.50.4.1087.

Berridge, Craig W., and David M. Devilbiss. "Psychostimulants as Cognitive Enhancers: The Prefrontal Cortex, Catecholamines, and Attention-Deficit/Hyperactivity Disorder." Biological Psychiatry, vol. 69, no. 12, June 2011, pp. e101–11. www.biologicalpsychiatryjournal.com, https://doi.org/10.1016/j.biopsych.2010.06.023.

Bloch, Michael H., and Ahmad Qawasmi. "Omega-3 Fatty Acid Supplementation for the Treatment of Children with Attention-Deficit/Hyperactivity Disorder Symptomatology: Systematic Review and Meta-Analysis." Journal of the American Academy of Child and Adolescent Psychiatry, vol. 50, no. 10, Oct. 2011, pp. 991–1000. PubMed, https://doi.org/10.1016/j.jaac.2011.06.008.

Cairncross, Molly, and Carlin J. Miller. "The Effectiveness of Mindfulness-Based Therapies for ADHD: A Meta-Analytic Review." Journal of Attention Disorders, vol. 24, no. 5, Mar. 2020, pp. 627–43. PubMed, https://doi.org/10.1177/1087054715625301.

Casella, Erasmo Barbante, and Beatriz Borba Casella. "Unraveling ADHD for the Pediatrician." Jornal de Pediatria, vol. 102, Mar. 2026, p. 101464. ScienceDirect, https://doi.org/10.1016/j.jped.2025.101464.

Castellanos, F. Xavier, et al. "Developmental Trajectories of Brain Volume Abnormalities in Children and Adolescents With Attention-Deficit/Hyperactivity Disorder." JAMA, vol. 288, no. 14, Oct. 2002, pp. 1740–48. Silverchair, https://doi.org/10.1001/jama.288.14.1740.

Castellanos, F. Xavier, and Erika Proal. "Large-Scale Brain Systems in ADHD: Beyond the Prefrontal–Striatal Model." Trends in Cognitive Sciences, vol. 16, no. 1, Jan. 2012, pp. 17–26. www.cell.com, https://doi.org/10.1016/j.tics.2011.11.007.

CDC. “Data on ADHD in Children.” Attention-Deficit / Hyperactivity Disorder (ADHD), 8 July 2026, https://www.cdc.gov/adhd/data/index.html.

CHADD. “Coexisting Conditions.” CHADD, 28 Aug. 2026, https://chadd.org/about-adhd/coexisting-conditions/.

Childress, Ann C., and Sally A. Berry. “Pharmacotherapy of Attention-Deficit Hyperactivity Disorder in Adolescents.” Drugs, vol. 72, no. 3, Feb. 2012, pp. 309–25. Springer Link, https://doi.org/10.2165/11599580-000000000-00000.

Cleveland Clinic. “Dopamine: What It Is, Function & Symptoms.” Cleveland Clinic, 23 Mar. 2022, https://my.clevelandclinic.org/health/articles/22581-dopamine.

———. “Norepinephrine: What It Is, Function, Deficiency & Side Effects.” Cleveland Clinic, 27 Mar. 2022, https://my.clevelandclinic.org/health/articles/22610-norepinephrine-noradrenaline.

———. “Serotonin: What Is It, Function & Levels.” Cleveland Clinic, 18 Mar. 2022, https://my.clevelandclinic.org/health/articles/22572-serotonin.

Collin, M. Samuel, et al. “Bioaccumulation of Lead (Pb) and Its Effects on Human: A Review.” Journal of Hazardous Materials Advances, vol. 7, Aug. 2022, p. 100094. ScienceDirect, https://doi.org/10.1016/j.hazadv.2022.100094.

Cortese, Samuele, Nicoletta Adamo, et al. “Comparative Efficacy and Tolerability of Medications for Attention-Deficit Hyperactivity Disorder in Children, Adolescents, and Adults: A Systematic Review and Network Meta-Analysis.” The Lancet Psychiatry, vol. 5, no. 9, Sept. 2018, pp. 727–38. www.thelancet.com, https://doi.org/10.1016/S2215-0366(18)30269-4.

Cortese, Samuele, Maite Ferrin, et al. “Neurofeedback for Attention-Deficit/Hyperactivity Disorder: Meta-Analysis of Clinical and Neuropsychological Outcomes From Randomized Controlled Trials.” Journal of the American Academy of Child and Adolescent Psychiatry, vol. 55, no. 6, June 2016, pp. 444–55. PubMed, https://doi.org/10.1016/j.jaac.2016.03.007.

DailyMed-Adderall. DailyMed - ADDERALL- Dextroamphetamine Saccharate, Amphetamine Aspartate, Dextroamphetamine Sulfate, and Amphetamine Sulfate Tablet. 29 May 2024, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=f22635fe-821d-4cde-aa12-419f8b53db81.

DailyMed-Clonidine. DailyMed - CLONIDINE HYDROCHLORIDE Tablet, Extended Release. 31 Oct. 2025, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=df80255e-f942-4d9f-8edd-cc95795772cb.

DailyMed-Intuniv/Guanfacine. DailyMed - INTUNIV- Guanfacine Tablet, Extended Release. 23 June 2025, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=b972af81-3a37-40be-9fe1-3ddf59852528.

DailyMed-Qelbree. DailyMed - QELBREE- Viloxazine Hydrochloride Capsule, Extended

Release. 5 Dec. 2025, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=aedf408d-0f84-418d-9416-7c39ddb0d29a.

DailyMed-Ritalin. DailyMed - RITALIN LA- Methylphenidate Hydrochloride Capsule, Extended Release. 23 Sept. 2025, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=effd952d-ac94-47bb-b107-589a4934dcca.

DailyMed-Strattera. DailyMed - STRATTERA- Atomoxetine Hydrochloride Capsule. 30 June 2026, https://dailymed.nlm.nih.gov/dailymed/lookup.cfm?setid=309de576-c318-404a-bc15-660c2b1876fb.

Del Campo, Natalia, et al. "The Roles of Dopamine and Noradrenaline in the Pathophysiology and Treatment of Attention-Deficit/Hyperactivity Disorder." Biological Psychiatry, vol. 69, no. 12, June 2011, pp. e145-157. PubMed, https://doi.org/10.1016/j.biopsych.2011.02.036.

Dopheide, Julie A. "ASHP Therapeutic Position Statement on the Appropriate Use of Medications in the Treatment of Attention-Deficit/Hyperactivity Disorder in Pediatric Patients." American Journal of Health-System Pharmacy, vol. 62, no. 14, July 2005, pp. 1502–09. Silverchair, https://doi.org/10.2146/ajhp040600.

Elliott, Jesse, et al. "Pharmacologic Treatment of Attention Deficit Hyperactivity Disorder in Adults: A Systematic Review and Network Meta-Analysis." PLOS ONE, vol. 15, no. 10, Oct. 2020, p. e0240584. PLoS Journals, https://doi.org/10.1371/journal.pone.0240584.

Faraone, Stephen V., and Henrik Larsson. "Genetics of Attention Deficit Hyperactivity Disorder." Molecular Psychiatry, vol. 24, no. 4, Apr. 2019, pp. 562–75. www.nature.com, https://doi.org/10.1038/s41380-018-0070-0.

FDA. "Treating and Dealing with ADHD." FDA, Aug. 2023. www.fda.gov, https://www.fda.gov/consumers/consumer-updates/treating-and-dealing-adhd.

FDA Intuniv Clinical Review. FDA Access Data for Intuniv/Guanfacine Extended Release Tablet. 3 Nov. 2014, https://www.accessdata.fda.gov/drugsatfda_docs/pediatric/22037%20guanfacine%20CDTL%20clinical%20both.pdf.

FDA Kapvay. FDA Supplement Approval: Kapvay (Clonidine Hydrochloride) Extended-Release Tablets, NDA 022331/S-001/S-002. 28 Sept. 2010, https://www.accessdata.fda.gov/drugsatfda_docs/appletter/2010/022331s001%2Cs002ltr.pdf.

FDA Pediatric Labeling–Strattera. New Pediatric Labeling Information Database - Strattera. 26 Nov. 2002, https://www.accessdata.fda.gov/scripts/sda/sdDetailNavigation.cfm?sd=labelingdatabase&id=14CE0322580E1B7DE053564DA8C071F2&rownum=756.

FDA Qelbree Review. Approval Package for Application Number 211964Orig1s000. Approval

Package for Application Number 211964Orig1s000, 2 Apr. 2021, https://www.accessdata.fda.gov/drugsatfda_docs/nda/2021/211964Orig1s000Approv.pdf.

Fernstrom, John D., and Madelyn H. Fernstrom. "Tyrosine, Phenylalanine, and Catecholamine Synthesis and Function in the Brain." The Journal of Nutrition, vol. 137, no. 6 Suppl 1, June 2007, pp. 1539S-1547S; discussion 1548S. PubMed, https://doi.org/10.1093/jn/137.6.1539S.

Granero, Roser, et al. "The Role of Iron and Zinc in the Treatment of ADHD among Children and Adolescents: A Systematic Review of Randomized Clinical Trials." Nutrients, vol. 13, no. 11, Nov. 2021, p. 4059. PubMed, https://doi.org/10.3390/nu13114059.

Greenhill, Laurence L., et al. "Practice Parameter for the Use of Stimulant Medications in the Treatment of Children, Adolescents, and Adults." Journal of the American Academy of Child & Adolescent Psychiatry, vol. 41, no. 2, Feb. 2002, pp. 26S-49S. www.jaacap.org, https://doi.org/10.1097/00004583-200202001-00003.

Hariri, Mitra, and Leila Azadbakht. "Magnesium, Iron, and Zinc Supplementation for the Treatment of Attention Deficit Hyperactivity Disorder: A Systematic Review on the Recent Literature." International Journal of Preventive Medicine, vol. 6, 2015, p. 83. PubMed, https://doi.org/10.4103/2008-7802.164313.

Harshitha, P., et al. "Neurodevelopmental Consequences Induced by Early-Life Lead Exposure Through Disruption of Neurotransmitter Pathways and Molecular Mechanisms: A Systematic Review." Developmental Neurobiology, vol. 86, no. 3, 2026, p. e70043. Wiley Online Library, https://doi.org/10.1002/dneu.70043.

Hinshaw, Stephen P., et al. "Attention-Deficit Hyperactivity Disorder, Multimodal Treatment, and Longitudinal Outcome: Evidence, Paradox, and Challenge." Wiley Interdisciplinary Reviews. Cognitive Science, vol. 6, no. 1, 2015, pp. 39–52. PubMed, https://doi.org/10.1002/wcs.1324.

Hunter, Catherine, et al. "A Closer Look at the Role of Nutrition in Children and Adults with ADHD and Neurodivergence." Frontiers in Nutrition, vol. 12, July 2025. Frontiers, https://doi.org/10.3389/fnut.2025.1586925.

Huss, Michael, et al. "Methylphenidate Dose Optimization for ADHD Treatment: Review of Safety, Efficacy, and Clinical Necessity." Neuropsychiatric Disease and Treatment, vol. 13, no. null, July 2017, pp. 1741–51. tandfonline.com (Atypon), https://doi.org/10.2147/NDT.S130444.

Jain, Rakesh, et al. "Clonidine Extended-Release Tablets for Pediatric Patients with Attention-Deficit/Hyperactivity Disorder." Journal of the American Academy of Child and Adolescent Psychiatry, vol. 50, no. 2, Feb. 2011, pp. 171–79. PubMed, https://doi.org/10.1016/j.jaac.2010.11.005.

Kesner, Raymond P., and John C. Churchwell. "An Analysis of Rat Prefrontal Cortex in Mediating Executive Function." Neurobiology of Learning and Memory, vol. 96, no. 3, Oct. 2011, pp. 417–31. ScienceDirect, https://doi.org/10.1016/j.nlm.2011.07.002.

Larsson, H., et al. "The Heritability of Clinically Diagnosed Attention Deficit Hyperactivity

Disorder across the Lifespan." Psychological Medicine, vol. 44, no. 10, July 2014, pp. 2223–29. Cambridge University Press, https://doi.org/10.1017/S0033291713002493.
Martínez, Lía, et al. "Executive Dysfunctions: The Role in Attention Deficit Hyperactivity and Post-Traumatic Stress Neuropsychiatric Disorders." Frontiers in Psychology, vol. 7, Aug. 2016, p. 1230. PubMed Central, https://doi.org/10.3389/fpsyg.2016.01230.
Michelson, David, et al. "Once-Daily Atomoxetine Treatment for Children and Adolescents with Attention Deficit Hyperactivity Disorder: A Randomized, Placebo-Controlled Study." The American Journal of Psychiatry, vol. 159, no. 11, Nov. 2002, pp. 1896–901. PubMed, https://doi.org/10.1176/appi.ajp.159.11.1896.
Mousa, Alyaa M. A., et al. "Effect of Developmental Lead Exposure on Neurogenesis and Cortical Neuronal Morphology in Wistar Rats." Toxicology and Industrial Health, vol. 34, no. 10, Oct. 2018, pp. 665–78. SAGE Journals, https://doi.org/10.1177/0748233718781283.
MTA Cooperative Group. "A 14-Month Randomized Clinical Trial of Treatment Strategies for Attention-Deficit/Hyperactivity Disorder. The MTA Cooperative Group. Multimodal Treatment Study of Children with ADHD." Archives of General Psychiatry, vol. 56, no. 12, Dec. 1999, pp. 1073–86. PubMed, https://doi.org/10.1001/archpsyc.56.12.1073.
Nasser, Azmi, et al. "A Phase III, Randomized, Placebo-Controlled Trial to Assess the Efficacy and Safety of Once-Daily SPN-812 (Viloxazine Extended-Release) in the Treatment of Attention-Deficit/Hyperactivity Disorder in School-Age Children." Clinical Therapeutics, vol. 42, no. 8, Aug. 2020, pp. 1452–66. www.clinicaltherapeutics.com, https://doi.org/10.1016/j.clinthera.2020.05.021.
NICE. Recommendations | Attention Deficit Hyperactivity Disorder: Diagnosis and Management | Guidance | NICE. NICE, 14 Mar. 2018, https://www.nice.org.uk/guidance/ng87/chapter/recommendations.
Nutt, David, et al. "Development of a Rational Scale to Assess the Harm of Drugs of Potential Misuse." The Lancet, vol. 369, no. 9566, Mar. 2007, pp. 1047–53. www.thelancet.com, https://doi.org/10.1016/S0140-6736(07)60464-4.
Oades, Robert D., et al. "The Influence of Serotonin- and Other Genes on Impulsive Behavioral Aggression and Cognitive Impulsivity in Children with Attention-Deficit/Hyperactivity Disorder (ADHD): Findings from a Family-Based Association Test (FBAT) Analysis." Behavioral and Brain Functions, vol. 4, no. 1, Oct. 2008, p. 48. Springer Link, https://doi.org/10.1186/1744-9081-4-48.
Omejc, Nina, et al. "Review of the Therapeutic Neurofeedback Method Using Electroencephalography: EEG Neurofeedback." Bosnian Journal of Basic Medical Sciences, vol. 19, no. 3, Aug. 2019, pp. 213–20. PubMed, https://doi.org/10.17305/bjbms.2018.3785.
Pliszka, S. R., et al. "Catecholamines in Attention-Deficit Hyperactivity Disorder: Current Perspectives." Journal of the American Academy of Child and Adolescent Psychiatry, vol. 35, no. 3, Mar. 1996, pp. 264–72. PubMed,

https://doi.org/10.1097/00004583-199603000-00006.
Potter, Alexandra S., et al. “Central Nicotinic Cholinergic Systems: A Role in the Cognitive Dysfunction in Attention-Deficit/Hyperactivity Disorder?” Behavioural Brain Research, vol. 175, no. 2, Dec. 2006, pp. 201–11. ScienceDirect, https://doi.org/10.1016/j.bbr.2006.09.015.
PubChem. Methylphenidate Hydrochloride. 5 Sept. 2026, https://pubchem.ncbi.nlm.nih.gov/compound/9280.
Qiu, Chunyue, et al. “Effects of Practicing Closed- vs. Open-Skill Exercises on Executive Functions in Individuals with Attention Deficit Hyperactivity Disorder (ADHD)—A Meta-Analysis and Systematic Review.” Behavioral Sciences, vol. 14, no. 6, June 2024, p. 499. PubMed Central, https://doi.org/10.3390/bs14060499.
Rong, Ying, et al. “Prevalence of Attention-Deficit/Hyperactivity Disorder in Individuals with Autism Spectrum Disorder: A Meta-Analysis.” Research in Autism Spectrum Disorders, vol. 83, May 2021, p. 101759. ScienceDirect, https://doi.org/10.1016/j.rasd.2021.101759.
Safren, Steven A., et al. “Cognitive Behavioral Therapy vs Relaxation with Educational Support for Medication-Treated Adults with ADHD and Persistent Symptoms: A Randomized Controlled Trial.” JAMA, vol. 304, no. 8, Aug. 2010, pp. 875–80. PubMed, https://doi.org/10.1001/jama.2010.1192.
Sallee, Floyd R., et al. “Guanfacine Extended Release in Children and Adolescents with Attention-Deficit/Hyperactivity Disorder: A Placebo-Controlled Trial.” Journal of the American Academy of Child and Adolescent Psychiatry, vol. 48, no. 2, Feb. 2009, pp. 155–65. PubMed, https://doi.org/10.1097/CHI.0b013e318191769e.
Seidman, Larry J., et al. “Dorsolateral Prefrontal and Anterior Cingulate Cortex Volumetric Abnormalities in Adults with Attention-Deficit/Hyperactivity Disorder Identified by Magnetic Resonance Imaging.” Biological Psychiatry, vol. 60, no. 10, Nov. 2006, pp. 1071–80. www.biologicalpsychiatryjournal.com, https://doi.org/10.1016/j.biopsych.2006.04.031.
Sharma, Alok, and Justin Couture. “A Review of the Pathophysiology, Etiology, and Treatment of Attention-Deficit Hyperactivity Disorder (ADHD).” Annals of Pharmacotherapy, vol. 48, no. 2, Feb. 2014, pp. 209–25. SAGE Journals, https://doi.org/10.1177/1060028013510699.
Shaw, P., et al. “Attention-Deficit/Hyperactivity Disorder Is Characterized by a Delay in Cortical Maturation.” Proceedings of the National Academy of Sciences, vol. 104, no. 49, Dec. 2007, pp. 19649–54. pnas.org (Atypon), https://doi.org/10.1073/pnas.0707741104.
Shaw, Philip. “ADHD: 10 Years Later.” Cerebrum: The Dana Forum on Brain Science, vol. 2013, Sept. 2013, p. 11.
Staley, Brooke S., et al. “Attention-Deficit/Hyperactivity Disorder Diagnosis, Treatment, and Telehealth Use in Adults — National Center for Health Statistics Rapid Surveys System, United States, October–November 2023.” MMWR. Morbidity and Mortality Weekly Report, vol. 73, 2024. www.cdc.gov, https://doi.org/10.15585/mmwr.mm7340a1.

Swanson, James M., and Nora D. Volkow. “Psychopharmacology: Concepts and Opinions about the Use of Stimulant Medications.” Journal of Child Psychology and Psychiatry, vol. 50, nos. 1–2, 2009, pp. 180–93. Wiley Online Library, https://doi.org/10.1111/j.1469-7610.2008.02062.x.

Tau, Gregory Z., and Bradley S. Peterson. “Normal Development of Brain Circuits.” Neuropsychopharmacology, vol. 35, no. 1, Jan. 2010, pp. 147–68. www.nature.com, https://doi.org/10.1038/npp.2009.115.

Tistarelli, Naomi, et al. “The Nature and Nurture of ADHD and Its Comorbidities: A Narrative Review on Twin Studies.” Neuroscience & Biobehavioral Reviews, vol. 109, Feb. 2020, pp. 63–77. ScienceDirect, https://doi.org/10.1016/j.neubiorev.2019.12.017.

Uchida, Mai, et al. “Assessing the Magnitude of Risk for ADHD in Offspring of Parents with ADHD: A Systematic Literature Review and Meta-Analysis.” Journal of Attention Disorders, vol. 25, no. 13, Nov. 2021, pp. 1943–48. SAGE Journals, https://doi.org/10.1177/1087054720950815.

Uldall Torp, Nanna Maria, and Per Hove Thomsen. “The Use of Diet Interventions to Treat Symptoms of ADHD in Children and Adolescents - a Systematic Review of Randomized Controlled Trials.” Nordic Journal of Psychiatry, vol. 74, no. 8, Nov. 2020, pp. 558–68. PubMed, https://doi.org/10.1080/08039488.2020.1769187.

Vanicek, Thomas, et al. “The Norepinephrine Transporter in Attention-Deficit/Hyperactivity Disorder Investigated With Positron Emission Tomography.” JAMA Psychiatry, vol. 71, no. 12, Dec. 2014, pp. 1340–49. Silverchair, https://doi.org/10.1001/jamapsychiatry.2014.1226.

Villagomez, Amelia, and Ujjwal Ramtekkar. “Iron, Magnesium, Vitamin D, and Zinc Deficiencies in Children Presenting with Symptoms of Attention-Deficit/Hyperactivity Disorder.” Children [Basel, Switzerland], vol. 1, no. 3, Sept. 2014, pp. 261–79. PubMed, https://doi.org/10.3390/children1030261.

Wernicke, J. F., and Christopher J. Kratochvil. “Safety Profile of Atomoxetine in the Treatment of Children and Adolescents with ADHD.” The Journal of Clinical Psychiatry, 63 Suppl 12, 2002, pp. 50–55.

Wigal, Sharon B., et al. “Exercise: Applications to Childhood ADHD.” Journal of Attention Disorders, vol. 17, no. 4, May 2013, pp. 279–90. SAGE Journals, https://doi.org/10.1177/1087054712454192.

Wolraich, Mark L., et al. “Clinical Practice Guideline for the Diagnosis, Evaluation, and Treatment of Attention-Deficit/Hyperactivity Disorder in Children and Adolescents.” Pediatrics, vol. 144, no. 4, Oct. 2019, p. e20192528. Silverchair, https://doi.org/10.1542/peds.2019-2528.

Yang, Gao, et al. “Effect of Aerobic Exercise on the Improvement of Executive Function in Children with Attention Deficit Hyperactivity Disorder: A Systematic Review and Meta-Analysis.” Frontiers in Psychology, vol. 15, June 2024. Frontiers, https://doi.org/10.3389/fpsyg.2024.1376354.

From Viral Trends to Business Growth: The Impact of TikTok and Instagram Marketing Strategies on The Entrepreneurial Success of Emerging Dance Businesses in the United States By Annie Li

Abstract

Throughout the recent years, TikTok and Instagram have become common platforms people frequent on. These social media platforms allow companies, particularly emerging businesses, to enhance, promote and engage. TikTok and Instagram utilize comprehensive algorithms that align to the taste of the user. With billions of monthly users these two platforms claim a wide local and global audience. Moreover, marketing has become a prevalent feature of both platforms as companies utilize the complex algorithms to secure consumers. This study investigates a research question: how do TikTok and Instagram marketing strategies influence the entrepreneurial success of emerging businesses in the United States dance industry? By exploring an under-studied niche industry, this paper bridges the gap between digital marketing theory and practical entrepreneurial success in the dance community. Research on the topic was conducted using existing literature that demonstrated the effects of social media, specifically TikTok and Instagram, on business performance and entrepreneurial success. Datasets, interviews and past studies were used to conduct findings. Results show that the reinforcement of proper TikTok and Instagram marketing tactics will likely enable emerging dance businesses to perform better in regards to enrollment and engagement rate. Future research on this topic should prioritize primary collection of evidence through surveys and interviews specific to this topic.

Keywords United States dance industry, dance entrepreneurship, social media marketing strategies, TikTok marketing, Instagram marketing and consumer engagement.

1 Introduction

1.1 Overview

With the recent advancements of technology, various social media platforms have become increasingly enticing, making them the perfect marketing tools. As globally renowned social media platforms, marketing done on TikTok and Instagram leaves a persisting impact on consumers: from children to adults alike. Now more than ever, it is imperative that businesses, especially entrepreneurial or start-ups, utilize this prevalence of digital marketing to their advantage. The dance industry especially, must use these marketing strategies to enhance longevity and exposure. Various types of commercial and cultural dance are housed within the dance industry. Some of which include hip-hop, african dance, classical ballet and modern dance (Kowal et al.). Moreover, the American dance industry has been noted to hold a very significant competition culture. It is common for studios to send their students to competitions such as Starpower, Showstopper and others, although some studios may choose not to participate in competitions (Weisbrod).

1.2 The United States Dance Industry

This study uses the United States dance industry as its focus geographic location. The purpose of choosing the United States dance industry as the geodemographic for this study is because the United States dance industry has specifically demonstrated the culture of the nation at large by consistently demonstrating the current trends within the country (Weisbrod). As noted before, the U.S. dance industry has a very core competition culture. It is also commonly seen within these dance competitions for choreographic themes and music choices to demonstrate and reciprocate popular culture (Weisbrod).

The U.S dance industry also includes a wide variety of different dance styles and cultures. Therefore, American dance competitions and the dance industry itself allow bridges to form between various cultures and people. Chinese and Peruvian immigrants in Minnesota and

New Jersey demonstrated how they maintained their cultural authenticity through dance. By continuing ethnic dances, the immigrants shared that they felt closer to the homes they had left behind and were able to share their cultural identity with their new society (Wilcox and Busse). Specifically, Chinese immigrants mentioned that while living in contemporary American society, they are able to hold on and share their ancestry through Chinese dance (Wilcox).

Unfortunately, the American dance industry is incredibly reliant on tuition and continues to receive cuts in public funding. Dance studio revenue is generated primarily from tuition, which can be problematic when companies do not acquire enough enrollments (Dalal). In IBIS World's recent data and analysis of the US dance industry, it shows revenue to have grown at a CAGR of 0.8% to $5.0 billion and was projected to expand to 2.3% in 2025 (Dalal). In terms of entrepreneurship within the U.S dance industry, it is common for dance entrepreneurs to be self-employed and working on a freelance basis (Bornman and James). Interestingly, studies have also found that even as a dance entrepreneur, studio owners will often still refer and display themselves as artists rather than a business owner (Bornman and James). This study will help identify what marketing tactics can be utilized to increase enrollment rates, benefiting the United States dance industry at large.

1.3 TikTok and Instagram Marketing

With the proliferation of digital marketing, the United States dance industry—like many other industries—has started to utilize TikTok and Instagram marketing tactics to improve business success and enrollment rates (Dalal). Existing research demonstrates that social media, specifically TikTok and Instagram, can be used to build a brand image, awareness, improve consumer engagement and business growth for entrepreneurs (Haenlein et al.).

In recent years, TikTok has established itself as a major entertainment platform and influential sales channel (Melgarejo-Espinoza et al.). TikTok has become a tool for businesses due to its adaptive, advanced algorithm that allows users to receive relevant content and products (Melgarejo-Espinoza et al.).

Instagram has also demonstrated a shift from being a simple social media platform to a notable marketing tool. Having enticing ads and colorful, and engaging visuals allows Instagram marketing to leave a lasting impression on users (Salunke and Jain).

1.4 Defining Entrepreneurial Success

For the purposes of this paper's research, entrepreneurial success refers to the increase in enrollment rate or level of consumer engagement following the usage of TikTok and Instagram marketing tactics. Positively influenced entrepreneurial success will indicate that there will be a high rate of student acquisition or consumer engagement. In contrast, negatively influenced entrepreneurial success will indicate there is a lower or equal rate of student enrollment or consumer engagement.

1.5 Research Question and Hypothesis

This paper explores a question: how do TikTok and Instagram marketing strategies influence the entrepreneurial success of emerging businesses in the United States dance industry? Findings from this study will aid emerging dance studio owners and dance entrepreneurs in understanding the effects of TikTok and Instagram marketing and influence. Additionally, the study will contribute to existing research on social media marketing's impact on entrepreneurship by addressing a niche industry that has received limited scholarly attention. It is hypothesized that TikTok and Instagram marketing strategies will increase the enrollment rate and level of consumer engagement in emerging dance businesses in the United States dance industry.

2 Literature Review

2.1 General Marketing Principles

For a business to last, they must have a superior marketing strategy. No matter how sensational or innovative a product or service is, it may still not attract any clients without sufficient marketing (Zaggelidou et al., 2013). Some refer to marketing as a business process that manages and creates opportunities for both sides to obtain their desires: consumers the goods and sellers the revenue. Others will refer to marketing as an exchange that satisfies both parties by the process of carrying out pricing, promotion and distribution of products (Stegemiller). Regardless of the exact definition, it is agreed upon that good marketing tactics can draw the clear line between failure and success (Zaggelidou et al., 2013). Below is a list of steps that the marketer should follow in developing a marketing plan as stated in a study conducted by Zaggelidou et al.

1. Define the product (dance)
2. Define the market (audience/students)
3. Define the goals to be attained (The goals must be measurable)
4. Develop the strategy (list the techniques involved)
5. Develop a plan of evaluation (every marketing plan must be able to be evaluated)

To be a good marketer, one must take many aspects into account. They must think themselves into the shoes of the consumer and look at all factors: cost, timing, accessibility, content and quality of program and service (Zaggelidou et al., 2013).

2.2 Marketing Strategies Within the Dance Industry

Like all other industries, the dance industry has had a long history of utilizing various marketing strategies to promote itself. Studies have shown that the most effective dance

marketing tactics involve being personal and developing close instructor-client relationships within the geographic location of the studio (Zaggelidou et al., 2013). Interest in the dance industry, whether it be for entertainment or personal health, often arises from communication and technological advances. The rise of these same developments has in turn enabled the entertainment industry to gain publicity (Zaggelidou et al., 2013). In correlation, digital marketing has repeatedly proven to pose an important tool for the dance industry (Gagliano). Dance can be classified as a form of art or entertainment; art and entertainment industries are incredibly reliant on mainstream televising and publicization through digital media. This makes digital media even more critical to success in the dance industry (Zaggelidou et al., 2013).

Digital media has become a popular marketing tactic in only recent years. Radio television, print advertising and physical distribution were what used to market the performing arts (Gagliano). These late strategies proved to be less efficient as they had the sole goal of promotion rather than using a more benefactory approach. According to Gagliano, current arts' marketing prioritizes "how" over "what" and which platforms are used to market the product is more important than the product itself. This new approach has proven to increase engagement in the dance industry as shown by the table below.

Table 1. Ballet Opera de Paris: Instagram (April 7, 2020)

	Instagram	
# of Followers	419,000 Followers	
Engagement on First Post (Oct. 2, 2014)	211 Likes	3 Comments
Engagement on Most Recent Post (April 7, 2020)	13,128 Likes	79 Comments
Years Between Posts	6 Years	
Increase of Engagement	12, 917 Likes	76 Comments

Source: (Galiano)

For the dance industry, audience development is vital to the longevity of a company. For private studios, students are a necessity and they often will join the dance competition culture. If they choose to omit competitions, studios will apply that as a marketing tool to attract enrollment from those who do not wish to participate in competitions (Zaggelidou et al., 2013).

2.3 TikTok Marketing Strategies

TikTok, a well-known social media platform, has an advanced algorithm that adapts to user taste. Other than simply being a source of entertainment, TikTok has consistently demonstrated its effectiveness as a marketing tool (Melgarejo-Espinoza et al.). In the recent decade, short video platforms have become increasingly popular. By 2020, the number of daily active users engaged on short video platforms has surged to 800 million which is two times higher than other video platforms such as YouTube (Mou). Additionally, there has been a clear

trend of short video platforms becoming incorporated within e-commerce (Cai). TikTok, being one of these short video platforms, has established an ecosystem that provides various marketing opportunities for emerging brands.

Should the marketing strategies be adequately utilized, business owners will largely benefit from an increased consumer engagement and profit rate (Mou). TikTok can be utilized as a marketing tool that allows businesses to interact with customers and generate personalized experiences (Melgarejo-Espinoza et al.). TikTok short videos hold a high brand communication value. Once a business is able to identify what users desire and care about, they can then push out a relatable video that generates high interest allowing the company to incite purchasing power. This strategy will then in turn often result in a strong influence for consumers to actively promote the product or service, and word-of-mouth marketing is generated (Mou).

2.4 Instagram Marketing Strategies

Similarly to TikTok, Instagram too has short video content and marketing opportunities. Unlike TikTok, Instagram places an emphasis on posts, stories and curating a marketable feed (Mou). Following its rapid growth, Instagram accommodates to a multitude of business opportunities. Advertising on Instagram is enticing to users due to the visual nature of the platform aiding the app in having a successful marketing route (Salunke and Jain). Instagram has also been recognized as a widely used social media platform as shown by the table below.

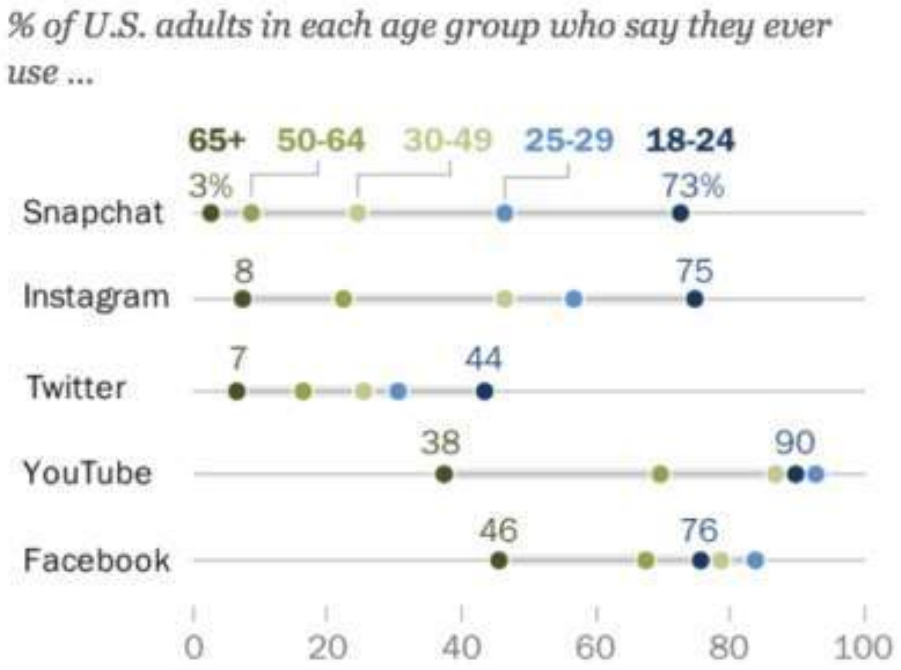

Source: (Mou)

Since its inception in 2010, Instagram has gathered a widely dispersed demographic of engaged consumers. Consumers became vulnerable to the advertising processes embedded within this social media platform. As user quantity continued to increase, the business industry

became incentivized to utilize Instagram to market their corporate organizations (Rahayuningrat et al.). Instagram marketing has proven to be so powerful that some businesses may start as a digitally native brand to become a renowned company. This includes Glossier, a beauty brand specializing in lip products, created through social media platforms including Instagram (Mou).

2.5 Social Media Marketing Impact on Dance and Other Businesses

Needless to say, social media marketing strategies have repeatedly proven effective. As audience numbers continue to decrease in dance theatres, strong social media marketing tactics become increasingly valuable. When a professional ballet company utilized social media to market themself, they found that the theatre profited from the increase in stakeholder engagement following digital marketing (Gagliano). TikTok and Instagram have become platforms for more than mere entertainment, the marketing qualities they possess show just how important they are in terms of promoting and advancing the dance industry (Lisowski).

In a study conducted in the United States on a small business, social media marketing was found to have increased the number of new customers, awareness, enhanced relationships with customers, aided in promoting local businesses and enabled the business to connect with consumers on a larger, global scale (Jones et al.). The usage of various social media platforms to market a brand has proven to increase brand equity, consumer engagement and loyalty, and business performance. Studies have also found that social media users who receive good customer service will tell three times as many people as non-users, positively impacting the brand image (Funk).

Entrepreneurs can benefit from proper social media marketing. By utilizing social media marketing to focus on relationship-based interactions with consumers, entrepreneurs are found to be much more lucrative (Bizhanova et al. 2019). In regards to student enrollment, studies have found that digital marketing strategies will strongly influence students in their enrollment decisions (Wijaya et al.). Generally, social media marketing has continuously demonstrated significant, beneficial effects on business performance and consumer engagement.

3 Methodology

3.1 Overview

This study uses qualitative secondary data obtained from academic dissertations to test and support the hypothesis that TikTok and Instagram marketing strategies will increase the entrepreneurial success of emerging dance businesses within the United States dance industry. Sources were identified by inputting keywords such as TikTok Marketing, Instagram marketing, dance entrepreneurship, entrepreneurship and United States dance industry into search engines and databases such as Google Scholar, JSTOR, Semantic Scholar and more. Due to time restrictions, no study was conducted with the participation of humans; therefore, institutional review board (IRB) approval was not required.

3.2 Choosing Platforms

TikTok and Instagram were the chosen social media platforms for this study. In

evaluating which platforms to conduct the study on, comprehensive research was done on varying social media platforms and their current reach, marketing and visual content. The platforms were chosen based on how large their day-to-day audience is, their recommendation systems, and impact on the audience. Upon thorough research of platforms such as Facebook, TikTok, Instagram, Youtube, X (formerly known as Twitter), etc, it was found that TikTok and Instagram marketing and reach had the most impact on consumers and audience. Hence, TikTok and Instagram were chosen as the social media platforms involved in this particular study.

3.3 Hypothesis Testing Process

To test the hypothesis, this study collected evidence of the likelihood of business and entrepreneurial success following Instagram and TikTok marketing tactics from a variety of sources. Sources used did not necessarily contain all the necessary components of this study, however. For example, one source may have demonstrated the impact of Instagram marketing on a recently established business, while not specifically discussing its impact on the dance industry or the geographical demographic. When researching for ideal evidence and papers, notable sites and relevant publication dates were taken into account. Sources were read in depth and annotated. Specific, preeminent points were highlighted within the text to make it more efficient when synthesizing and analyzing data in the later stages. Data was collected using a dataset from Kaggle; past papers found on Sage Journals, Journal of Dance Education, SCITEPRESS, Canadian Journal of Business and Information Studies, International Journal of Advanced Computer Science and Applications, and more; and interviews from studies done on a certain aspect of this study. Examples of the data collected include how effective TikTok and Instagram marketing strategies are, how social media marketing impacts the dance industry, how social media influences consumers, and how social media marketing affects the enrollment rate of students for general education. This broader data was then analyzed to form a coherent response towards the research question.

4 Results

4.1 TikTok and Instagram Marketing Influence on Consumers

Using a Kaggle dataset on social-media impact, this analysis focuses on the 28 participants who selected both TikTok and Instagram among the platforms they frequented. The full dataset included 300 U.S. participants. The source reports an age range of 22-40; however, the dataset documentation and sampling limitations do not establish that this range represents the broader population of prospective dance students or purchasers. This age-range concern is addressed in the limitations section.

Figure 1: Influence of TikTok and Instagram on Consumers

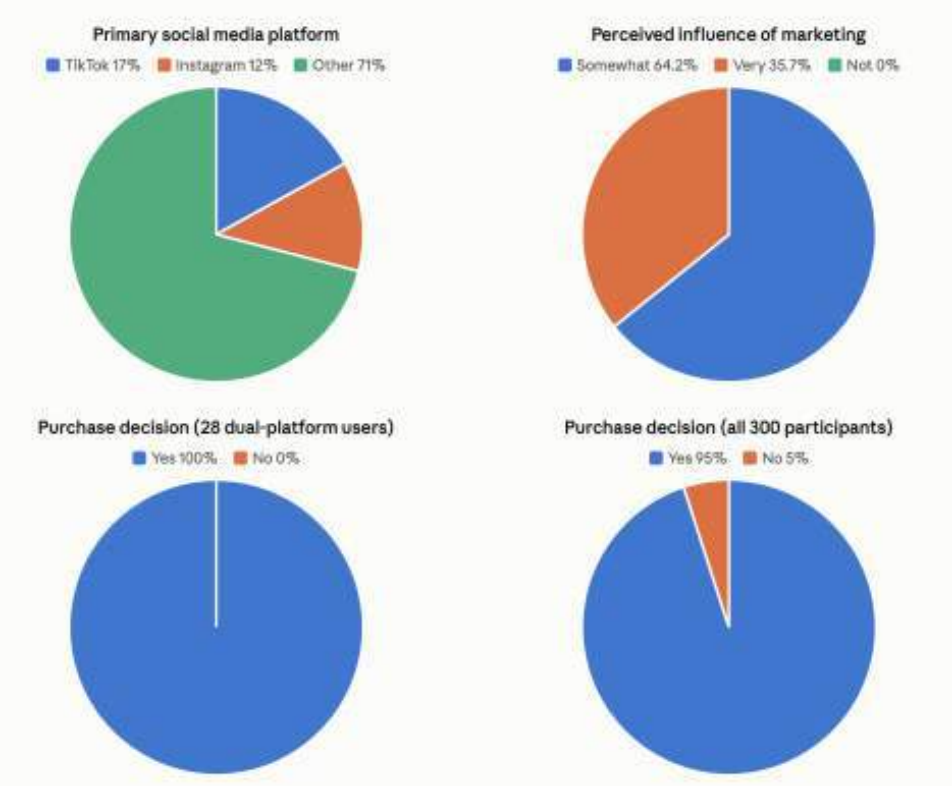

Source: ("Impact of Social Media (Dataset)")

Among the 28 participants in the TikTok-and-Instagram subgroup, 18 reported that marketing was somewhat influential; none selected "not influential." All 28 participants also reported making a purchase based on something seen on social media. In the full sample, 285 of 300 participants reported such a purchase. Thus, the subgroup result matches the strong purchasing pattern in the baseline sample rather than exceeding it. Since the subgroup is small and the dataset is observational, these counts indicate an association but should not be interpreted as causation that TikTok or Instagram caused the purchases.

4.2 Social Media Marketing Impact on the United States Dance Industry

The following data was collected in various papers on the subject of how social media marketing has transcended the United States dance industry. This table shows a qualitative review of the type of company, what marketing strategy they used and how that tactic changed or impacted the company as a whole.

Table 1: Social Media Marketing Impact on U.S Dance Industry

Type of Dance Business	Social Media Marketing Strategy Used	Observed Change and Findings
Professional Ballet Companies	Used social media to build their brand, reputation, and connect with audiences.	Social media increased audience engagement, attendance, and profits.

Hip-Hop	Used Instagram and Facebook to promote classes and share videos/posts with current and potential students.	Social media helped increase class attendance and attract more students. Teachers also used it to connect with other dancers and find teaching opportunities.
Commercial Dance Industry (freelance/small personal brands)	Dancers used TikTok and YouTube to promote themselves and their work. They also used Instagram DMs to communicate with customers.	Social media helped dancers build their careers and companies, gain visibility, network, and promote themselves. TikTok was especially useful for networking and self-promotion.
Professional Dance Industry (Nashville Ballet)	Used keywords in bios, posted at strategic times, tagged celebrities, and promoted upcoming performances and summer programs.	Social media increased engagement, especially during performances and summer programs. Tagging celebrities also helped attract more attention.

Sources in Order: (Gagliano), (Zimányi and Lanszki), (Lisowski), (Lambert)

The findings of Table 1 demonstrated in this table show that the application of social media marketing strategies has largely benefited the dance industry and community. Although one study claimed that social media marketing could get complex and stressful, it still highlighted that social media provides an excellent platform for marketing commercial dance. Following the usage of social media marketing, there has been an increase in profit, engagement, audience, networking and visibility. One study even asserted that social media can market teachers as well as the company and connect these teachers with global opportunities. These findings justify the idea that social media marketing is an essential tool that is exceptionally effective towards promoting and enhancing the dance industry and community.

4.3 Social Media Marketing Impact on Entrepreneurs

Using prior research on social-media marketing and entrepreneurship, Table 2 presents two illustrative cases. These examples show how entrepreneurs have used analytics and social platforms, but they do not validate the paper's specific hypothesis about emerging U.S. dance studios.

Table 2: Two Illustrative Cases of Social Media Use by Entrepreneurs

Form of Social Media Marketing Done	Impact on Entrepreneur
Using data and analytics to improve social media campaigns.	Helped entrepreneurs identify which strategies worked best and improve their performance.
Using Instagram and other platforms to market businesses and show real-life business issues.	Entrepreneurs viewed social media as essential for business growth. All 11 interviewees said it was important, and one said 90% of their business came from social media.

Sources in Order: (Geho and Dangelo), (Olanrewaju et al.)

Case 1 - Geho and Dangelo: Entrepreneurs used data and analytics to evaluate social-media campaigns, identify more effective strategies, and improve performance.

Case 2 - Olanrewaju et al.: Interviewees described social media as important to business

growth; all 11 interviewees considered it important, and one attributed 90% of the business to social media.

Together, these cases illustrate plausible mechanisms through which social media may support entrepreneurial decision-making and visibility. However, the cases involve different business and geographic contexts and therefore provide contextual support rather than direct validation of the dance-studio hypothesis.

4.7 Social Media Marketing's Impact on Enrollment Rate

Since no local study directly measured the effect of TikTok and Instagram campaigns on enrollment in emerging U.S. dance studios, this paper uses Indonesian higher-education research as an illustrative proxy. Figure 2 reports effectiveness scores assigned to social-media marketing strategies in that setting.

Figure 2: Reported Effectiveness of Social Media Marketing Strategies for Enrollment

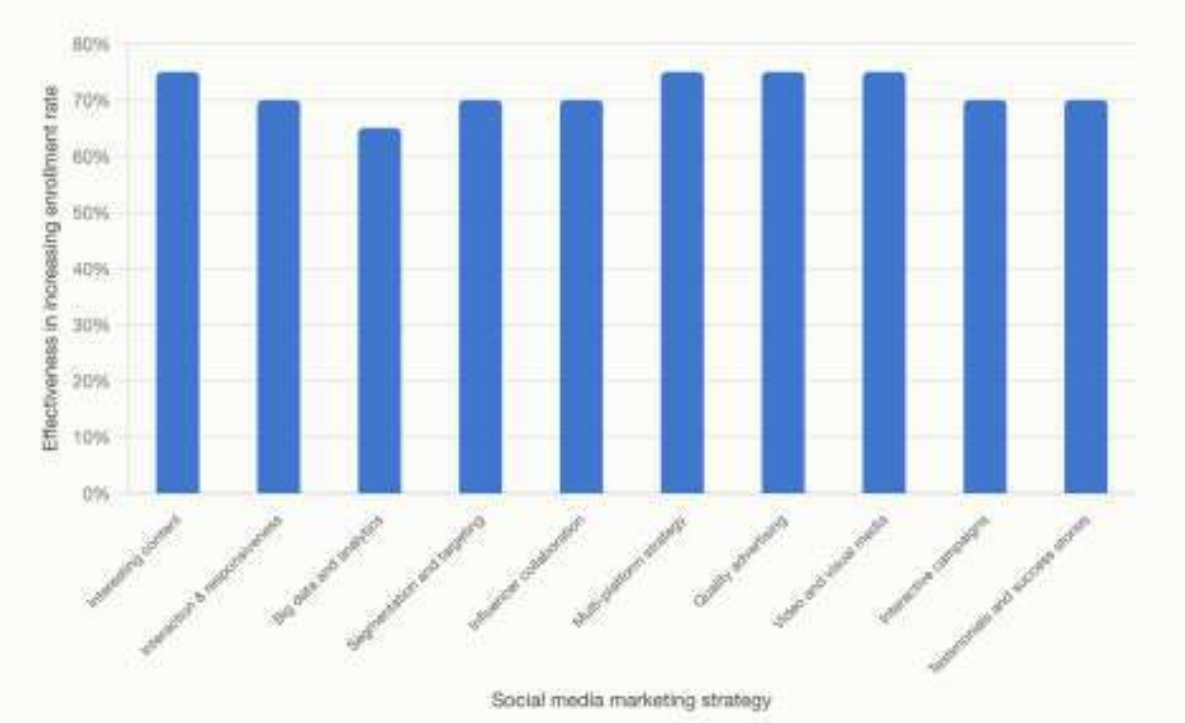

Source: (Siminto et al.)

The proxy differs from the target population in several important ways: it concerns higher-education institutions rather than private dance studios; it was conducted in Indonesia rather than the United States; and it addresses prospective college students rather than potential dance students and their parents or guardians. Though the figure does not directly estimate dance studio enrollment outcomes, it identifies potentially relevant tactics, such as visually engaging content and multi-platform communication. It demonstrates the association between effectiveness in increasing enrollment rate and social media marketing.

5 Discussion

The results of this study demonstrate how TikTok and Instagram marketing tactics will positively influence the entrepreneurial success of an emerging dance business in the United States. The initial hypothesis was partially validated, as the data revealed a positive association between TikTok and Instagram marketing efforts, elevated consumer engagement, and higher proxy efficiency in enrollment . Since higher education data was used as a proxy for enrollment in the dance industry, claims about enrollment must remain tentative.

5.1 Analysis

There were four main pieces of data collected to support this study. These include how TikTok and Instagram marketing influences consumers, social media marketing's impact on the U.S dance industry, social media marketing's impact on entrepreneurs and their business success, and how effectively social media marketing increases enrollment rates in institutions of higher education. While there is no data collected to support the specific idea that TikTok and Instagram marketing tactics will increase enrollment rates of an emerging dance business in the United States, it can be inferred based on the consistency of data throughout different sectors.

By combining the four data points a clear sequence emerges from initial consumer interest to sustained entrepreneurial success. Figure 1 provides the fundamental capability of TikTok and Instagram marketing as the data shows that the platforms directly drive consumer purchasing decisions. Table 1 grounds this effect within the dance industry, demonstrating that social media marketing can increase audience engagement, attendance, and visibility. Table 2 then goes on to present data that entrepreneurs find social media marketing on such platforms is motivational and this mental aspect may help the company perform better in the long run. Social media marketing includes uploading posts, videos, stories and more on platforms such as TikTok and Instagram promoting a service or product to a certain audience. Finally, Figure 2 demonstrated how an increase in enrollment rates corresponded to specific tactics, something company owners should look at to enhance their own social media marketing. Marketing tactics such as influencer collaboration, interesting and visual content, and multi-platform usage will largely benefit business success as they showed high effectiveness in increasing enrollment rates in higher education institutions.

These findings suggest that using these platforms and specific content catered towards prospective future students the audience will make the decision to enroll or engage with the company. Therefore, new dance companies should utilize TikTok and Instagram marketing to promote their business and gain attraction.

Overall, results align with findings from past literature. Gagliano asserted that modern arts marketing prioritizes the delivery channel, "how", over the content itself, "what"; which aligns with the data in Table 1 showing that dance businesses generated increased engagement and attendance by adopting specialized social media marketing tactics over traditional outreach. Furthermore, Zaggelidou et al. emphasized that dance marketing depends on building close, personal relationships; the data from Table 1 demonstrates through the personal interactions and community building using Instagram DMs and TikTok comments. The enrollment data in Figure 2 validates Wijaya et al. by proving that targeted digital strategies directly drive enrollment

decisions, offering a reliable method to stabilize tuition revenues amid the funding cuts highlighted by Dalal.

As dance studios in the United States receive cuts in funding, it is inevitable for business owners to be educated on how to properly and effectively market themselves on social media. The findings of this study are significant because they will aid dance entrepreneurs in optimizing their resources to elevate their business.

5.3 Limitations

This study had one major limitation: there was no primary data collected that directly identified dance entrepreneurs in the United States and had them test Instagram and TikTok marketing tactics. Therefore, data was not catered towards answering the specific question this study posed. That being said, the trends in data still indicate that the influence of TikTok and Instagram marketing is positive and universally beneficial.

Other limitations included limited polling geodemographic, socioeconomic background, demographic control, and outdated context. Research used in this study targeted large ballet companies such as the Nashville Ballet, dance businesses in more suburban places, or did not specify the geodemographic. Additionally, there was no specific data collected on the socioeconomic background of those involved in the studies and the demographic of those who were polled and researched upon was not specified within the papers that were utilized. Lastly, some research and findings were from more than a decade prior and certain circumstances may have changed in the recent years.

6 Conclusion

As social media continues to exert a profound influence on modern culture, platforms such as TikTok and Instagram serve as vital tools for optimizing marketing, sales, and overall business success. This study demonstrates the strong correlations between TikTok and Instagram marketing and its influence on consumer choices and success within the United States dance industry. While there was no primary data collected to directly answer the research question, the findings of this study indicate that Instagram and TikTok marketing will allow an emerging business in the United States dance industry to achieve higher enrollment rates and consumer engagement.

Using data collected from this study it can be concluded that emerging dance businesses in the United States should utilize specific Instagram and TikTok marketing strategies to benefit their company. New studios aiming for a higher enrollment rate should use interesting and visual content to engage prospective students. Content that is interesting has been shown to be more effective than data and statistics when it comes to increasing enrollment rates. That being said, content should be a mixture of data, statistics, visual and interesting content. While it has been found that interesting content leads to an increase in enrollment, statistics and data will allow consumers to understand the credibility of the company. Additionally, it has been found that creating content targeting a specific audience will increase chances of reaching those who are more likely to become direct customers. Collaborating with influencers and ambassadors greatly accelerates the rate of enrollment. Therefore, studios should invest in quality ambassadors and

influencers within the field to create engaging, relatable posts that will increase consumer engagement. Ultimately, a mix of content will enable consumers to understand and engage with the company in more ways, increasing potential student enrollments.

Future continuations on this study should work to incorporate direct, primary data conducted to clearly target the research hypothesis. Specifically, future studies should track quantitative changes in student enrollment and revenue across newly founded private dance studios and compare companies who implement structured TikTok and Instagram campaigns and those who do not. Data collection should occur at regular intervals, such as every two weeks for a year or monthly for three years. Additionally, future studies should conduct qualitative interviews of the dance entrepreneurs in order to clarify how TikTok and Instagram marketing alters an entrepreneur's business strategy, offering deeper insight into how social media adoption directly enhances and impacts the studio.

Works Cited

Bizhanova, Kenzhegul, et al. “Impact of Digital Marketing Development on Entrepreneurship.” E3S Web of Conferences, edited by D. Rudoy and V. Murgul, vol. 135, 2019, p. 04023. DOI.org (Crossref), https://doi.org/10.1051/e3sconf/201913504023.

Bornman, Dawid Alwyn Jacobus, and Augustina Geertruida James. “Analyzing Entrepreneurial Identity and Capabilities as Determinants for Business Sustainability in the Dance Industry.” Journal of Dance Education, vol. 0, no. 0, Apr. 2026, pp. 1–13. Taylor and Francis+NEJM, https://doi.org/10.1080/15290824.2025.2515442.

Cai, Yang. “The Integration Communication Revolution of Short Video Platforms: Innovative Approaches and Business Practices of TikTok Marketing Strategies:” Proceedings of the 2nd International Conference on E-Commerce and Modern Logistics [Singapore, Singapore], 2025, pp. 151–57. DOI.org (Crossref), https://doi.org/10.5220/0013840400004719.

Dalal, Michael. “Dance Studios in the US Industry Analysis, 2025.” Ibis World, 2025, https://www.ibisworld.com/united-states/industry/dance-studios/5350/.

Funk, Tom. Advanced Social Media Marketing: How to Lead, Launch, and Manage a Successful Social Media Program. Apress, 2014.

Gagliano, Alexandra Lord. “Digital Media Marketing for Dance: Enhancing Audience Engagement with Ballet Companies and Concert Dance Performance.” UF Journal of Undergraduate Research, vol. 22, Nov. 2020. journals.flvc.org, https://doi.org/10.32473/ufjur.v22i0.121994.

Geho, Patrick R., and Jennifer Dangelo. “THE EVOLUTION OF SOCIAL MEDIA AS A MARKETING TOOL FOR ENTREPRENEURS.” ProQuest, 2015, https://www.proquest.com/openview/92a27d7460667a0d79b2084328a15deb/1?pq-origsi t e=gscholar&cbl=28211.

Haenlein, Michael, et al. “Navigating the New Era of Influencer Marketing: How to Be Successful on Instagram, TikTok, & Co.” California Management Review, vol. 63, no. 1, Nov. 2020, pp. 5–25. SAGE Journals, https://doi.org/10.1177/0008125620958166. “Impact of Social Media (Dataset).” 2023, https://www.kaggle.com/datasets/lastman0800/impact-of-social-media-dataset. Jones, Nory, et al. “Impact of Social Media on Small Businesses.” Journal of Small Business and Enterprise Development, vol. 22, no. 4, Nov. 2015, pp. 611–32. Silverchair, https://doi.org/10.1108/JSBED-09-2013-0133.

Kowal, Rebekah J., et al. The Oxford Handbook of Dance and Politics. Oxford University Press, 2017.

Lambert, Lindsay. Applying Social Media Marketing Content Strategies from the Entertainment Field to Professional Dance. 2017.

Lisowski, Luke. Dancing in the Digital Age: How Social Media Shapes the Commercial Dance Industry. 2025.

Melgarejo-Espinoza, Raquel, et al. “Impact of the TikTok Algorithm on the Effectiveness of Marketing Strategies: A Study of Consumer Behavior and Content Preferences.” International Journal of Advanced Computer Science and Applications, vol. 16, Feb. 2025. ResearchGate, https://doi.org/10.14569/IJACSA.2025.0160215.

Mou, Jessie Boxin. Study on Social Media Marketing Campaign Strategy -- TikTok and Instagram. 2020, https://dspace.mit.edu/entities/publication/f6b62f8e-a6f9-49db-8a9d-ffee1aea5da5 Massachusetts Institute of Technology.

Olanrewaju, Abdus-Samad Temitope, et al. “The Influence of Social Media on Entrepreneur Motivation and Marketing Strategies in a Developing Country.” Challenges and Opportunities in the Digital Era, edited by Salah A. Al-Sharhan et al., Springer International Publishing, 2018, pp. 355–64. Springer Link, https://doi.org/10.1007/978-3-030-02131-3_32.

Rahayuningrat, Prahita Sri, et al. “The Effectiveness of Instagram Advertising for Digital Marketing Strategy: Systematic Literature Review.” International Journal of Tourism and Hospitality in Asia Pacific, vol. 7, no. 1, Feb. 2024, pp. 1–13. DOI.org (Crossref), https://doi.org/10.32535/ijthap.v7i1.2888.

Salunke, Parth, and Varsha Jain. “Instagram Marketing (2015–2021): A Review of Past Trends, Implications, and Future Research.” Exploring the Latest Trends in Management Literature, edited by Sudhir Rana et al., vol. 1, Emerald Publishing Limited, 2022, p. 0. Silverchair, https://doi.org/10.1108/S2754-586520220000001007.

Siminto, et al. “Social Media Marketing Strategy in Increasing Student Enrollment and Retention.” Al-Fikrah Jurnal Manajemen Pendidikan, vol. 12, June 2024, pp. 202–13. ResearchGate, https://doi.org/10.31958/jaf.v12i1.5799.

Stegemiller, Terresa Michelle. Dance Marketing Strategies: Application of Theory to Practice - ProQuest. 1995, https://www.proquest.com/openview/f2249594981d7f02873090dbaca2338e/1?pq-origsite =gscholar&cbl=18750&diss=y.

Wijaya, Hadi, et al. “ANALYSIS OF DIGITAL MARKETING STRATEGIES ON INTEREST AND ENROLLMENT DECISIONS OF PROSPECTIVE NEW STUDENTS IN PRIVATE HIGHER EDUCATION INSTITUTIONS IN INDONESIA (A CASE STUDY OF JAKARTA GLOBAL UNIVERSITY).” Klabat Journal of Management, vol. 4, no. 2, Sept. 2023, pp. 147–62. ejournal.unklab.ac.id, https://doi.org/10.60090/kjm.v4i2.1007.147-162.

Wilcox, Hui. “Movement in Spaces of Liminality: Chinese Dance and Immigrant Identities.” Ethnic and Racial Studies, vol. 34, no. 2, Feb. 2011, pp. 314–32. Taylor and Francis+NEJM, https://doi.org/10.1080/01419870.2010.528439.

Wilcox, Hui Niu, and Erika Busse. “‘Authentic’ Dance and Racialized Ethnic Identities in Multicultural America: The Chinese in Minnesota and Peruvians in New Jersey.” Sociology of Race and Ethnicity, vol. 3, no. 3, July 2017, pp. 355–69. SAGE

Journals, https://doi.org/10.1177/2332649216674224.
Zaggelidou, Eleni, et al. “The Effect of Marketing on Dance Activity.” Journal of Physical Education and Sport, vol 13, no. 2, 2013, pp. 220-230. https://doi.org/10.7752/jpes.2013.02037
Zimányi Gabriella, and Lanszki Anita. “The Influence of Social Media on Hip-Hop Dancers and Their Classes.” Tánc és Nevelés, vol. 1, nos. 1–2, Aug. 2020, pp. 80–96. ojs.mtak.hu, https://doi.org/10.46819/TN.1.1.80-96.

The Effects of Dance Movement Therapy on Executive Function in individuals with Autism Spectrum Disorder By Christopher Tsai

Abstract

Autism Spectrum Disorder (ASD) is associated with differences in executive function (EF), a set of cognitive processes involved in planning, organizing, inhibiting responses, maintaining attention, and adapting to changing demands. Dance Movement Therapy (DMT) is a movement-based therapy that consists of physical movement, rhythm, emotional expression, and social interaction, serving as a potential intervention for children with ASD. This literature review examines the existing evidence regarding the potential effects of DMT and related rhythmic movement interventions on EF in children with ASD. The review draws on empirical studies, systematic reviews, and meta-analyses. Existing research suggests that movement and rhythm-based interventions may support aspects of EF, particularly inhibitory control, cognitive flexibility, and motor planning. However, much of the available evidence examines movement broadly, with limitations in sample size and control conditions. Further research using larger samples, active control groups, standardization, and child-specific ASD populations is needed to determine the effectiveness of DMT as an EF intervention.

Keywords Dance Movement Therapy, Autism Spectrum Disorder, executive function, intervention, children, fMRI, behavioral neuroscience

Introduction

The Centers for Disease Control and Prevention (CDC) provided an estimate that 1 out of 31 children in the United States are diagnosed with Autism Spectrum Disorder (ASD) (Centers for Disease Control and Prevention). ASD is a neurodevelopmental condition that is defined by deficits in two main categories: social communication/interaction and restricted/repetitive patterns of behaviour (American Psychiatric Association 50). Since ASD is referred to as a “spectrum,” it presents itself slightly differently in every individual, with varying challenges and needs. The current Diagnostic and Statistical Manual-5 categorizes ASD into three levels of severity. Level 1 is "Requiring support," where individuals show challenges in social communication, inflexibility, and behaviors of sameness across more than one context. Next, Level 2 is "Requiring substantial support": marked by deficits in both verbal and nonverbal social communication skills, specifically with restricted and repetitive behaviors that are frequently observed. Lastly, Level 3 is "Requiring very substantial support": in which there are severe deficits in verbal and nonverbal social communication that cause differences in daily functioning. Individuals in Level 3 have limited initiation of social interactions and require significant coping and behavioral support when faced with change (American Psychiatric Association 52).

ASD is often associated with a range of characteristics that could affect everyday functioning. Some challenges that individuals with ASD face are with social-emotional

regulation; motor skills, which may result in restricted body language and simplified motor movements; and difficulties in Executive Function (EF) (Lord et al. 2). Executive functioning skills are a set of cognitive skills that include the combination of planning, organizing, initiation, sequencing, and acting (Diamond 136). One important aspect of EF is inhibitory control which includes self-control and controlling attention. These skills allow individuals to inhibit impulses and filter out distracting external stimuli. In particular, inhibitory control of attention focuses on selective attention and the suppression of interfering external stimuli (Diamond 137). EF in specific is of interest to study due to its proposed role in contributing to certain challenges in children with ASD. This domain refers to a broad set of purposeful higher-order neuropsychological abilities which include but are not limited to ideation, planning, organization, goal-directed actions, abstract reasoning, decision-making, and social regulation (Demetriou et al. 1198). Planning, organization, decision-making, and goal-directed actions refer to higher-order EF skills that require intentional concentration and attention (Diamond 137). Similarly, abstract reasoning relies on cognitive flexibility, which allows individuals to mentally manipulate ideas, think creatively, and view things from multiple different perspectives (Diamond 149). EF is important to address for neurodivergent children because without concentrative and attentive thinking, children with ASD may be more likely to act on their impulses rather than regulate their behavior (Diamond 141).

While a variety of treatment options exist to address EF deficiencies, one form of intervention that could be used in children with ASD is Dance Movement Therapy (DMT). Meta-analysis study and data show that physical/movement interventions in ASD can improve inhibitory control (Liang 84). The study identified 14 studies with a mix of performance tasks and rating scales, where the Behavior Rating Inventory of Executive Function (BRIEF) concluded that inhibitory control and cognitive flexibility reached average effects while working memory had minimal statistical changes (Liang 78). The American Dance Therapy Association (ADTA) defines DMT as "the psychotherapeutic use of movement to promote emotional, social, cognitive, and physical integration of the individual, for the purpose of improving health and well-being" (Millman et al. 24). Currently, research regarding DMT remains sparse, however, there is existing research that explains how exercise therapy can improve the executive functioning skills of children with ASD (Wang et al. 285). The results of the study suggested that greater gains were observed with school age children with ASD as opposed to preschool age children, and that improvement in executive functioning and inhibition were noted, with minimal gains in working memory (Liang, 79). Existing research on DMT suggests that techniques such as reflection, synchronous movements, mirroring, rhythm, body awareness, dance, and relaxation can support changes in young children (Tortora 6). Studying this topic is valuable because evaluating the existing evidence of DMT as an intervention for improving executive function in neurodivergent children strengthens understanding of existing evidence base and identifies what needs to continue being explored in this topic (Tortora 7).

This paper draws on empirical evidence and systematic reviews to further explore the effects of DMT on EF in children with ASD. Specifically, this paper will synthesize findings

from existing research on how DMT can potentially offer unique benefits in enhancing executive function in children with ASD, emphasizing its non-traditional approach that integrates movement, emotional expression, and cognitive engagement to strengthen development in EF (Tortora 10).

This research examines EF in children with ASD and discusses the differences in EF in ASD compared to the neurotypical pediatric populations. Then, the discussion will shift into the foundations of DMT and how movement may positively influence EF and other neurodevelopmental outcomes. Next, this study will examine empirical research to evaluate if DMT can improve EF in children with Autism Spectrum Disorder (Tortora 7). Discussion of the limitations of existing literature will help propose potential directions for future research and evaluate if DMT can be a positive intervention for improving executive function in children with ASD.

Autism Spectrum Disorder and Executive Function

Executive Function is a major area of interest when studying children with ASD, as it is a proposed cause of challenges in a variety of cognitive, inhibition, and social domains.

In children with ASD, EF encompasses the processes responsible for: regulation, planning, adapting to changes of routines, or selective attention (Demetriou et al. 1199). Autistic individuals exhibit varying levels of strengths and needs; hence the term spectrum within the diagnostic label. This means that some children with ASD will face more symptoms of EF challenges while others may experience reduced difficulties in EF skills. Demetriou et al., conducted a meta-analysis that used existing databases such as Embase, Medline, and PsycINFO, that included studies from the introduction of autism into DSM-III in 1980 all the way through June 2016, to compare EF performance in ASD with neurotypical controls A total of 235 studies comprising 6,816 participants with ASD and 7,265 control participants were within the analysis. Each study's EF difference was converted into a standardized effect size (Hedges' g), which yielded a moderate overall EF impairment in ASD (Demetriou et al. 1201). The evidence indicates that individuals with ASD experience moderate challenges across all major EF domains when compared with neurotypical individuals. While there is no single EF domain that shows significantly greater impairment, these findings suggest that overall differences in executive function in ASD is widespread and relatively uniform across executive processes (Demetriou et al. 1201).

Additionally, Silk et al. examined brain networks involved in visuospatial processing and executive function between neurotypical adolescents and adolescents with ASD. The study consisted of a total of 16 male participants, split into 7 with ASD and 9 typically developing comparison participants. The participants' mean age was 14.7 years (ranging from 11–18). The ASD group met DSM-IV criteria for autistic disorder or Asperger's disorder, while the control group of neurotypical children was free of known medical, neurological, or psychiatric disorders. The study employed a mental rotation task that used “three-dimensional Shepard-Metzler-type block shapes” as stimuli, with a target stimulus and four test stimuli presented concurrently for

10 seconds (Silk et al. 1441). The participants were instructed to select one of four buttons on a fiber-optic button panel to identify which of the test shapes corresponded to the rotated target shape. A baseline condition required the participants to indicate which test stimulus was the closest visual match to the target stimulus, when no rotation was involved. In this task, researchers presented a total of 18 rotation stimuli and 18 baseline stimuli, in alternating blocks of three rotation trials followed by three baseline trials. This covered a total scanning duration of 396 seconds. A follow up functional Magnetic Resonance Imaging (fMRI) scan, described above, conducted on the group of individuals with ASD showed metabolic reduction in the prefrontal cortex. Particularly in divisions of the anterior cingulate gyrus, a region that is critically involved in EF skills and supports the individual's behavior, controls attention, regulates motivation, and manages distractions (Silk et al. 1442). In summary, there is evidence that challenges in EF skills are evident within the ASD group. Through identification of specific brain regions associated with executive function, this raises the possibility that movement-based interventions such as DMT could help address these differences, though this remains to be directly tested.

EF shows an evident difference in functioning for neurotypical children compared to ASD children. Gentil-Gutiérrez et al, compared EF skills between school-aged children and adolescents with ASD (requiring level 1 and level 2 support levels) with their neurotypical peers (4). The BRIEF-2 results found that there was significantly greater executive dysfunction across all domains in the ASD group. This means that for skills such as regulating behaviour, adapting to changing situations, planning, organizing, and controlling attention, children with ASD are more likely to experience challenges compared to their neurotypical peers (Gentil-Gutiérrez et al. 3). However, by identifying and monitoring the strengths and weaknesses of each child's individual EF skills, teachers and other caregivers can receive valuable guidance to expand their range of potential, corrective intervention options to optimize school performance. (Gentil-Gutiérrez et al. 6)

Dance Movement Therapy and Executive Function

DMT is a creative arts therapy that uses movement to integrate the emotional, cognitive, physical, and social aspects of an individual (Martin 545). Initially pioneered by Marian Chace, this therapeutic approach first emerged in the 1940s and was based on the principle that bodily and psychological changes can reciprocally influence one another (Millman et al. 25). Liu et al. developed an 8-week DanceSport program to examine cognitive flexibility, inhibitory control, and working memory in typically developing preschoolers, and whether effects differ by age. This double-blind RCT enrolled children from a Beijing kindergarten (201 analyzed) randomized to an experimental group (two 45-minute DanceSport sessions per week for 8 weeks) or a control group. The experimental group demonstrated significant gains in inhibitory control, cognitive flexibility, and working memory (Liu et al.).

While DMT encompasses a range of approaches, there are three main models of DMT that have been identified (Karkou et al. 3). The first model centers primarily on dance and

movement engagement, using structured or free movement, with or without the therapist moving alongside the client (Karkou et al. 3). The second model prioritizes non-verbal interaction between client and therapist. A leading example was developed by the American DMT pioneer Marian Chace, who introduced mirroring as a key technique (Sandel et al.). This strategy includes the therapist who reflects the client's movements to them, thereby engaging in empathy and non-verbal communication (Karkou et al. 5). The third model concentrates on improvised movement directed toward internal experience and creative expression in the presence of the therapist. A classic example is Authentic Movement, developed by the American choreographer Mary Whitehouse (Whitehouse 3). In this form of DMT, the therapist remains still and observes, serving as an empathic witness, and music is typically omitted so that it does not shape the client's movement (Whitehouse 3-7). Of these three, the second model appears to hold the most promise for improving executive function (EF) in children with autism spectrum disorder (ASD), because of its central reliance on mirroring. The Mirroring Intervention Protocol, developed in 2012, aims to promote empathy, social competence, body awareness, and awareness of the boundary between self and others (Koch and Kercher 107).

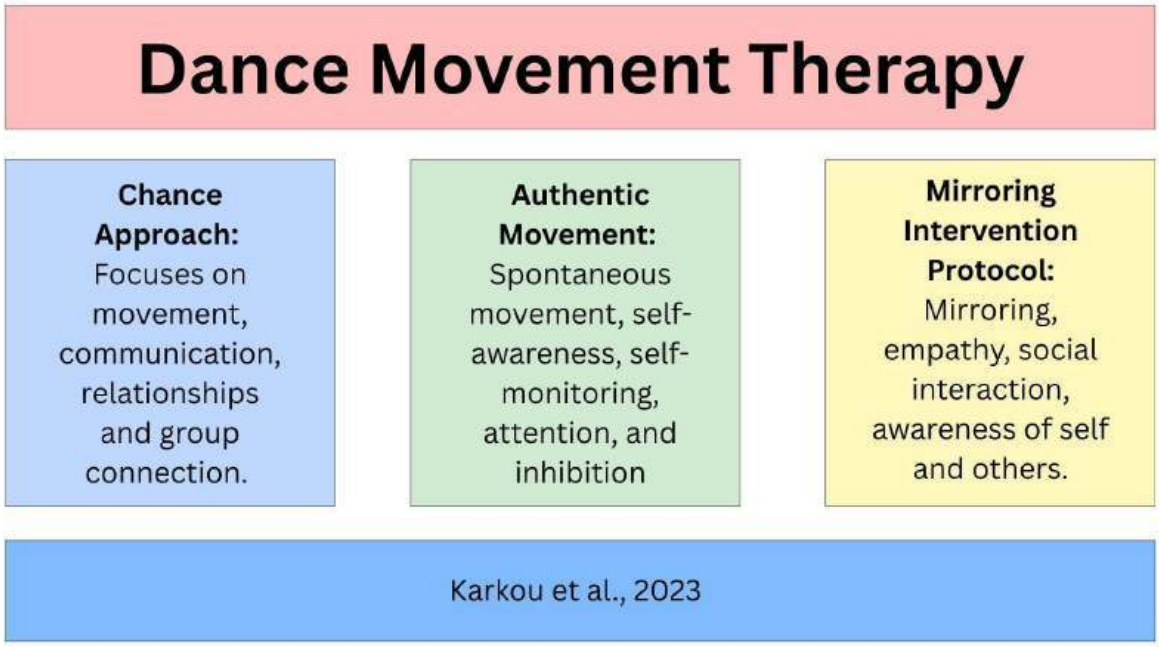

Figure 1: Core models and focus areas within Dance Movement Therapy interventions

Additionally, the three models of DMT described above share a common feature: recruitment of motor planning, rhythmic movement synchrony, mirroring, and body awareness, which require the premotor and frontostriatal circuits that Silk et al. found to be underactive in ASD (1442). The networks implicated in executive functions such as working memory, attention,

and cognitive control therefore suggest plausible targets addressed during dance movement–based intervention. In addition, Gentil-Gutiérrez et al. reported that children with ASD demonstrated challenges compared to their neurotypical peers across all nine BRIEF-2 clinical scales, with specific deficits in inhibition, shifting, emotional control, working memory, planning, and monitoring (4). This further suggests that DMT and related rhythmic/movement interventions can address these clusters of functions and serve as a potentially effective intervention modality.

Existing Studies of DMT in Children with ASD

In Vazou et al.'s study, there was a total participant count of 39 children aged 6–11 years who participated in either a rhythmic program or a physical education program. Of those 39 children, 22 were grouped in the rhythmic program and 17 in the physical education program. The groups shared many similarities in categories such as age, IQ, fitness, height, weight, and BMI. Both programs lasted a total of 7 weeks, with 2 sessions per week and 30 minutes per session (Vazou et al. 3).

Children in the Rhythmic Program learned to move alongside music and rhythm. Some of the activities participants engaged in included: clapping, jumping, hopping, walking to music, drumming on stability balls, using drumsticks, bouncing balls to the beat, creating their own movement patterns, partner work, group performances, rhythmic games, and Philippine-inspired "Tinikling" (Vazou et al. 10). Comparatively, children in the Physical Education Program participated in more standard PE activities such as: throwing, catching, kicking, dribbling, running, obstacle courses, tumbling, locomotor skills, and team sports (Vazou et al. 10). The main difference between the Physical Education Program and Rhythmic Program is the PE program lacked emphasis on incorporating rhythm or music. Children in both groups of the Rhythmic Program and Physical Education Program improved significantly. Some of the improvements that occurred for EF skills included: faster reaction times, better accuracy, stronger inhibition, better working memory, and greater cognitive flexibility. Interestingly, the biggest improvements were seen in the most difficult of Flanker tasks, which suggests that children became better at complex executive processing.

Researchers measured Executive Functions during these programs using the Flanker Test (Vazou et al. 3). This test was broken down into three levels: Standard Flanker which measures inhibition (the ability to ignore distractions and control impulses); Reverse Flanker which measures inhibition and working memory (the ability to hold and use information in the mind); and lastly Mixed Flanker which measures inhibition, working memory, and cognitive flexibility (the ability to switch between tasks and ideas).

In conclusion, results of this study showed that there were no major performance changes between the standard Physical Education Program vs Rhythmic Program. However, since rhythmic activities simultaneously require listening, timing, movement coordination, attention, memory, inhibition, creativity, and cooperation, this helps create an example of embodied cognition, where thinking and movement begin to develop together (Vazou et al. 11). This is

important to recognize because through combining cognitive and physical demands, brain networks that involve executive function are strengthened to improvements that transfer to other everyday tasks. While results showed that there was no superiority of one program over the other, embodied cognition provides promising potential for rhythm-based movement. Overall, since there was no non-movement control, the study cannot confirm that instructed movement was responsible for the observable gains.

While this study does not specifically explore how rhythmic movement affects children on the Autism Spectrum, this research provides foundational evidence that rhythmic movement may positively influence aspects of child development regarding EF skills. Overall, it offers clarity on the potential of DMT as a positive intervention, since significant benefits in neurotypical children were displayed through movement. Future studies should replicate this study and adjust by adding control conditions to better isolate whether the observed effects are specific to children with ASD and attributable to DMT itself. Most importantly, future studies should specifically include children with ASD to obtain more diverse, accurate, and reliable results regarding the effectiveness of DMT for improving EF skills in neurodivergent children.

Furthermore, in another study, Hildebrandt et al. employed a “randomized controlled trial that investigated the effects of a 10-week manualized dance and movement therapy intervention on negative symptoms in participants with autism spectrum disorder" (15). For some context, this study consisted of a total of 78 individuals with ASD who agreed to participate in the study. The breakdown was 12 female, 63 male, and three that didn’t provide data on their gender, where the mean age was 22.5 years old (ranging from 14–53). One major criteria of this study was that participants had to have had no previous experience with DMT or similar therapeutic approaches (Hildebrandt et al. 6).

Some symptoms of the ASD participants included social communication differences, restricted or repetitive behaviors, a strong need for sameness, highly specific interests, and sensory differences. The experiment was intended to be a 2:1 treatment-to-control ratio, so approximately 52 participants would be placed in the DMT group while approximately 26 would be placed in the control group (Hildebrandt et al. 2). Participants in the DMT group received ten weekly sessions of manualized DMT while participants in the control group continued with their individual routine. The DMT group participated in sessions that took 60 min and consisted of three mirroring exercises and one verbal processing element. The control group’s individual routine consisted of typical day-to-day activities that didn’t comprise other forms of psychotherapy (Hildebrandt et al 2).

The results of this study show that from the DMT group’s experiences, DMT leads to an encouraging trend towards stronger symptom reduction in the treatment group for overall negative symptoms in almost all subtypes. The mean symptom reduction in the treatment group was 15.27%, while symptom severity actually worsened by 6.99% in the control group. The observed symptom reduction of 15.27% is “classified as a small effect and corresponds with a CGI-I (Clinical Global Impression-Improvement) score of “minimally improved, which is

regarded as clinically substantial" (Hildebrandt et al.10). This shows further evidence that body-oriented therapy can support participants and patients with ASD.

To conclude, the intervention also helps to enhance self-perception by adding improvements in all aspects of the social functioning skills. In the future, research should therefore focus on control interventions that differ from the movement therapy only in categories of mirroring, self-perception, and emotion (Hildebrandt et al. 3). Since the findings showed potentially influential factors such as the importance of dyadic mirroring with a non-autistic partner, future research can continue to learn from this and build upon it. Overall, this study explored DMT intervention in individuals with ASD and how it addressed challenging symptoms. While it didn't focus specifically on EF or study children/individuals under 18, it offered a strong body of information to formulate an argument for this research topic. Future research could benefit from focusing on including children and examining EF, but seeing symptom reduction being clinically significant shows promising potential for this area of research.

Limitations of Current Research

Throughout the review, the major limitation of DMT associated with the outcomes of EF is the limited research within this subject area. To start, a limitation across the existing research connects with limited sample sizes. For example, Silk et al. included only seven male adolescents with ASD (1442), while Vazou et al. included only 39 children (3). These small sample sizes are a significant limitation because it can be difficult for researchers to generalize their findings to the broader population, specifically the ASD population.

Another major limitation is that existing studies have difficulty isolating the specific effects that movement-based interventions and DMT have on EF skills. In Vazou et al., both the rhythmic activity and physical education groups demonstrated improvements (11) which made it challenging to determine whether the improvements were specifically attributable to rhythm and music. Similarly, Hildebrandt et al. compared DMT with participants' usual motor routines rather than an active intervention which made it difficult to separate the effects of DMT from other factors such as attention gathered from the therapist, social engagement, and participation in a structured task (10). In the future, research should prioritize using clear control groups which would help better determine whether DMT itself produces improvements in EF skills.

Additionally, future researchers can also investigate how different genres and styles of dance can influence emotional control and mood in children with ASD, and the area of EF has not been thoroughly evaluated with a controlled group, with explicit gaps in the outcomes of DMT and the impact of EF.

Conclusion

This literature review examined whether Dance Movement Therapy (DMT) may be an effective intervention for improving executive function (EF) in children with Autism Spectrum Disorder (ASD). The reviewed literature demonstrates that children with ASD frequently

experience differences in EF, such as challenges in inhibition, cognitive flexibility, working memory, planning, attention, and regulation (Demetriou et al. 1201). The existing research also provides a promising foundation for considering movement-based interventions as a potential approach to EF (Karkou et al. 3; Tortora 6). Vazou et al. found that children participating in both rhythmic and physical education programs demonstrated improvements in several EF-related outcomes, including reaction time and inhibition (3). Although the study did not find the rhythmic program to be superior to the physical education program, its findings suggest that structured movement can engage multiple cognitive processes associated with EF (Vazou et al. 5; Silk et al. 1441). In addition, there is neurological evidence that also provides a basis for further investigation into movement-based interventions. Silk et al. identified differences in activation within the prefrontal, premotor, and anterior cingulate in adolescents with ASD during an executive-function-related task (1442). Because these neural networks are involved in processes such as cognitive control, motor planning, attention, and working memory, their involvement provides a possible framework for understanding why movement-based interventions could influence EF-related skills (1442).

Overall, the evidence provides a promising theoretical foundation for the use of DMT as an intervention for EF-related skills; however, the current research does not provide sufficient evidence to conclude that DMT independently improves EF in children with ASD. Several limitations within the current literature further reinforce the need for additional research. Small sample sizes, limited control conditions, and differences in intervention approaches and outcome measures make it difficult to determine the specific effects of DMT on EF (Silk et al. 1442; Vazou et al. 3; Hildebrandt et al. 6). Future research should directly investigate DMT and EF in children with ASD and examine which components of DMT may contribute most strongly to EF-related outcomes. Specific elements such as mirroring, rhythm, choreography, motor planning, movement sequencing, and memorization may provide particularly useful areas of investigation (Karkou et al. 3). Specifically, Hildebrandt et al.'s mirroring protocol provides a useful direction towards improving EF-related skills, such as attention, inhibitory control, and cognitive flexibility, by encouraging nonverbal communication and attention to someone else's movements. Although more rigorous research is necessary, DMT provides a potentially valuable approach to supporting cognitive and developmental needs in children with ASD (Hildebrandt et al., 6; Karkou et al.5; Millman et al. 24). Overall, while DMT and related movement interventions may provide emerging outcomes in the area of executive function-relevant skills in children with ASD, the current evidence is limited to draw firm conclusions. Therefore, given that no existing studies examine DMT and its effect on executive function in autistic children, further research is required to better understand this research topic and address the gap in the literature.

Works Cited

Aithal, S., and V. Karkou, editors. Arts Therapies Research and Practice with Persons on the Autism Spectrum: Colourful Hatchlings. Routledge, 2023. https://doi.org/10.4324/9781003201656.

American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders. 5th ed., American Psychiatric Association, 2013. https://doi.org/10.1176/appi.books.9780890425596.

Centers for Disease Control and Prevention. "Autism Spectrum Disorder (ASD)." 21 May 2026, https://www.cdc.gov/autism/index.html.

Demetriou, E. A., et al. "Autism Spectrum Disorders: A Meta-Analysis of Executive Function." Molecular Psychiatry, vol. 23, no. 5, 2018, pp. 1198–1204. https://doi.org/10.1038/mp.2017.75.

Diamond, Adele. "Executive Functions." Annual Review of Psychology, vol. 64, 2013, pp. 135–168. https://doi.org/10.1146/annurev-psych-113011-143750.

Gentil-Gutiérrez, A., et al. "Executive Functions in Children and Adolescents with Autism Spectrum Disorder, Grade 1 and 2, vs. Neurotypical Development: A School View." International Journal of Environmental Research and Public Health, vol. 19, no. 13, 2022, article 7987. https://doi.org/10.3390/ijerph19137987.

Hildebrandt, M. K., S. C. Koch, and T. Fuchs. "'We Dance and Find Each Other': Effects of Dance/Movement Therapy on Negative Symptoms in Autism Spectrum Disorder." Behavioral Sciences, vol. 6, no. 4, 2016, article 24. https://doi.org/10.3390/bs6040024.

Karkou, V., et al. "Dance Movement Therapy for Dementia." Cochrane Database of Systematic Reviews, no. 8, 2023. https://doi.org/10.1002/14651858.CD011022.pub3.

Koch, S. C., and K. Kercher. "The Impact of Mirroring in Dance Movement Therapy on Empathy and Relating in Autism Spectrum Disorder: A Secondary Reflexive Research." Arts Therapies Research and Practice with Persons on the Autism Spectrum: Colourful Hatchlings, edited by S. Aithal and V. Karkou, Routledge, 2023, pp. 101–116. https://doi.org/10.4324/9781003201656-10.

Landers, M. J. F., et al. "A Systematic Review of the Use of Subcortical Intraoperative Electrical Stimulation Mapping for Monitoring of Executive Deficits and Neglect: What Is the Evidence So Far?" Acta Neurochirurgica, vol. 164, no. 1, 2022, pp. 177–191. https://doi.org/10.1007/s00701-021-05012-w.

Liang, X., et al. "The Effects of Exercise Interventions on Executive Functions in Children and Adolescents with Autism Spectrum Disorder: A Systematic Review and Meta-Analysis." Sports Medicine, vol. 52, no. 1, 2022, pp. 75–91. https://doi.org/10.1007/s40279-021-01545-3.

Liu, X., et al. "The Effects of Early Childhood DanceSport Intervention on Executive Functions in Preschoolers." *Scientific Reports*, vol. 15, no. 1, 2025, doi:10.1038/s41598-025-10751-2. https://doi.org/10.1038/s41598-025-10751-2.

Lord, C., et al. "Autism Spectrum Disorder." The Lancet, vol. 392, no. 10146, 2018, pp.

508–520. https://doi.org/10.1016/S0140-6736(18)31129-2.
Martin, M. “Moving on the Spectrum: Dance/Movement Therapy as a Potential Early Intervention Tool for Children with Autism Spectrum Disorders.” The Arts in Psychotherapy, vol. 41, no. 5, 2014, pp. 545–553. https://doi.org/10.1016/j.aip.2014.10.003.
Millman, L. S. M., et al. “Towards a Neurocognitive Approach to Dance Movement Therapy for Mental Health: A Systematic Review.” Clinical Psychology & Psychotherapy, vol. 28, no. 1, 2021, pp. 24–38. https://doi.org/10.1002/cpp.2490.
Sandel, S. L., S. Chaiklin, and A. Lohn, editors. Foundations of Dance/Movement Therapy: The Life and Work of Marian Chace. Marian Chace Memorial Fund of the American Dance Therapy Association, 1993. Silk, T. J., et al. “Visuospatial Processing and the Function of Prefrontal-Parietal Networks in Autism Spectrum Disorders: A Functional MRI Study.” American Journal of Psychiatry, vol. 163, no. 8, 2006, pp. 1440–1443. https://doi.org/10.1176/ajp.2006.163.8.1440.
Tortora, Suzi. The Dancing Dialogue: Using the Communicative Power of Movement with Young Children. Redleaf Press, 2005. Vazou, S., et al. “Rhythmic Physical Activity Intervention: Exploring Feasibility and Effectiveness in Improving Motor and Executive Function Skills in Children.” Frontiers in Psychology, vol. 11, 2020, article 556249. https://doi.org/10.3389/fpsyg.2020.556249.
Wang, H., G. Cheng, and M. M. Li. “The Effectiveness and Sustained Effects of Exercise Therapy to Improve Executive Function in Children and Adolescents with Autism: A Systematic Review and Meta-Analysis.” European Journal of Pediatrics, vol. 184, no. 5, 2025, article 286. https://doi.org/10.1007/s00431-025-06120-w.
Whitehouse, Mary Starks. "C. G. Jung and Dance Therapy: Two Major Principles." Eight Theoretical Approaches in Dance-Movement Therapy, edited by Penny Lewis Bernstein, Kendall/Hunt, 1979, pp. 51–70. https://doi.org/10.1007/BF02579607

The Case for Edible Coatings in Bread Waste Prevention By Ryan Wu

Abstract

Bread is a critical source of essential nutrients for people around the world, with some populations relying on bread for almost all of their daily calories. Despite the importance of bread, millions of tons are still wasted due to two factors: staling, caused by starch retrogradation; and microbial spoilage, mainly caused by molds. Starch retrogradation is the process by which starches within the bread recrystallize, leading to moisture migration outside of the bread. Microbial spoilage occurs due to fungal spores in the air that come in contact with the bread after baking. First, this paper explores established chemical methods incorporated within the bread dough, such as α- and β-amylases and sodium hydroxide, in order to prevent starch retrogradation and microbial spoilage. Secondly, this study explores and advocates for novel edible coatings as a cheap, effective, and noninvasive solution to the problem of bread waste through preventing retrogradation and/or microbial spoilage. Natural edible coatings, mallow flower extract and grapefruit seed extract, were preferred due to public aversion against unnatural ones.

Introduction

Wheat plays a critical role in ensuring food security, and with the global production volume of wheat being almost 785 million metric tons, wheat bread is a staple food for many in the world. Wheat bread contains fibers, carbohydrates, lipids, proteins, vitamins, and minerals, which supplies a significant amount of energy. Bakery products provide 75% of total caloric intake as well as 90% of needed protein in some parts of the world (“Bread Staling Measurement Techniques”). Worldwide, bread consumption per capita was 23.5 kg in 2022 (Ollinger et al.).

The most significant problem concerning bread waste is its short shelf life, which is a result of staling and microbiological contamination (Deseta et al.). Approximately 7-30% of bread worldwide is wasted due to inferior texture as a result of staling (Wang et al.). The staling process is caused by starch retrogradation, the process by which starch, a type of carbohydrate, begins to change structures after cooling (Niu et al.). Two main starch molecules, amylose and amylopectin, undergo a process called starch gelatinization in which they absorb water and swell, leading to the formation of an orderly crystalline structure when these molecules cool down (Ning et al.). Amylose recrystallizes faster than amylopectin, with the former rearranging itself into double helix structures. The recrystallization of amylopectin, which may take days or weeks, is the main driver behind the drying and hardening of bread after a long period of time (Ning et al.). Moisture migration promotes amylose and amylopectin recrystallization as well as retrogradation, where starch-based foods become dry in low humidity environments but softer in high humidity ones (Ning et al.).

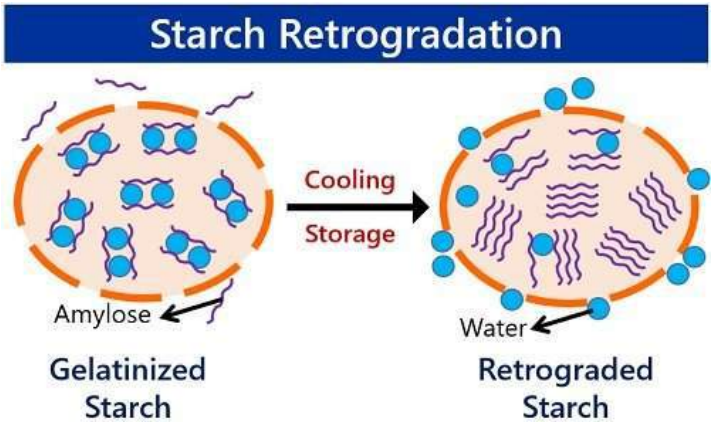

Fig 1: Starch retrogradation mechanism (Supriya N, 2022)

Bread is also vulnerable to microbial attack, mainly due to fungi or molds that originate from raw materials and bakery air (Garcia et al.). The baking process transforms the starch of the dough into a digestible gelatinized form in bread, promoting mold growth (Hafyan et al.). The most common method to prevent mold in bread is the use of preservatives, but this defense can be insufficient against the resistance of some fungal strains (Garcia et al.).

In order to prevent food waste and starch retrogradation, economical anti-retrogradation and anti-microbial methods must be found that work under standard temperature and pressure conditions. Previous research has shown that enzymatic, fermentation, physical, and natural product addition methods can replace the chemical additives typically used in the food industry, which would ensure better safety and nutritional quality (Niu et al.). For example, (Tao et al.) found that decreasing starch granule surface proteins on bread reduced bread quality as well as promoted retrogradation. (Liu et al.) also revealed that black rice anthocyanins can prevent retrogradation. However, the economic feasibility of these methods and impact of their production processes must be better studied (Niu et al.).

This paper will explore edible coatings, safe to eat layers on the surface of foods, as a cost-effective way to prevent starch retrogradation and microbiological contamination. Edible coatings such as ethanol, propionic acid, sorbic acid, and acetic acid, and their salts are frequently used as antimicrobial agents on bread (Ning et al.). Because bread is a staple food and is the main source of nutrients for people across the world, finding a low-cost way to prevent bread waste has the possibility of decreasing starvation for a large amount of people.

Enzyme Advances in Starch Retrogradation

Enzymes are commonly used in starch-based products to delay starch retrogradation because of their ability to directly affect the formation of the starch network, recrystallize the amylopectin, and move water during processing and storage (Ning et al.). Examples of these enzymes are α-amylase and β-amylase, both used at an industrial scale. α-amylase substantially improves the softness of breadcrumbs after three days of storage (Niu et al.). Malto-oligosaccharides, especially maltotetraose, block retrogradation due to their ability to hold

water molecules and inhibit starch-starch interactions (Wang et al.). Enzymes such as α-amylases and maltogenic amylase, commonly used in various starch-based foods for anti-retrogradation purposes, yield little of these Malto-oligosaccharides (Wang et al.). A novel malteratose-forming amylase has just recently been found, which reduced hardness by 43.7% (Wang et al.).

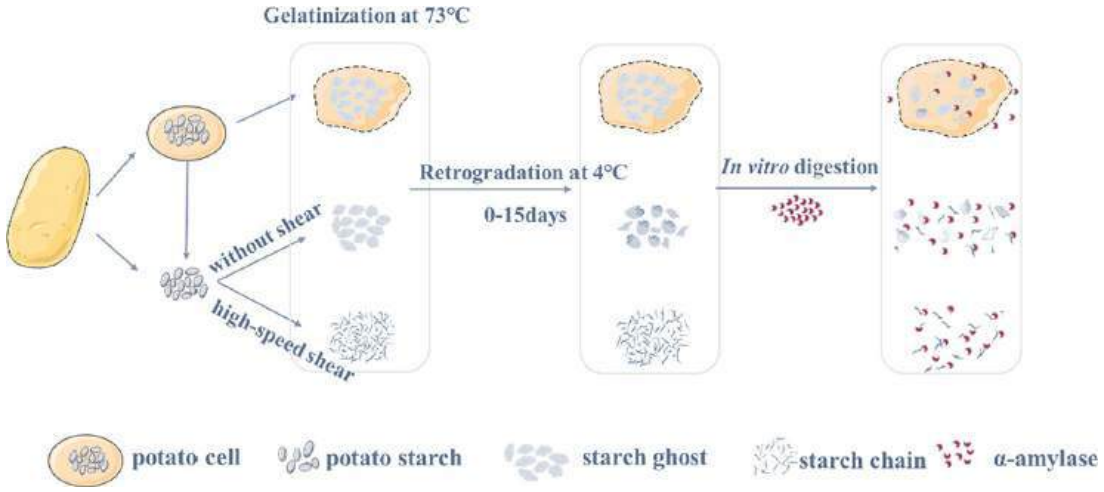

Fig 2: A-amylase effect on starch retrogradation (Shu et al.).

In rice cakes, starch that was pretreated with β-amylase was found to lead to rice cakes with a lower hardness than for an untreated group (Niu et al.). This effect can be attributed to the increased proportion of short branch chains, which inhibited amylopectin rearrangement (Niu et al.). One problem with using β-amylase is that its economic feasibility is low. Presently, industrial β-amylase is usually extracted from plants such as sweet potatoes and soybeans. This extraction process is complex, production costs are high, and the large demand for grains could contribute to food supply shortages (Niu et al.).

Methods

Rice cake studies were considered in this review because of their similar composition to bread. Both bread and rice cakes are highly starchy foods that suffer from starch retrogradation (Lee et al.). Because starch retrogradation is determined by the behavior of amylose and amylopectin, which are the main carbohydrates that make up both foods, anti-retrogradation methods on rice cakes are believed to have a similar effect on bread and vice-versa. For example, this paper treats mung bean starch's anti-retrogradation effect on rice cakes as one that most likely applies to bread as well. Additionally, the contaminating spores that cause mold on bread are present in areas where rice cakes are stored and use the same carbohydrates for energy (Ning et al.). Thus, anti-microbial effects on rice cakes or bread were considered applicable to the other.

Other Protein Advances in Starch Retrogradation

Experimental methods to prevent retrogradation in bread involve starch granule surface proteins (SGSPs) and starch granule-associated proteins (SGAPs), which are proteins that are adsorbed to the surface of and embedded inside starch molecules, respectively (Tao et al.). In one

experiment, the absence of SGSPs and SGAPs lowered the bread's air holding capacity, increasing the cell density. This increased density promoted water migration and evaporation by reducing the cell's water holding capacity (Tao et al.). This finding elucidates the role of SGSPs and SGAPs in slowing starch retrogradation, yet it requires future research into the addition of SGSPs and SGAPs to bread (Tao et al.).

In general, proteins can inhibit or promote starch retrogradation. One effect of retrogradation is the release of water from the bread's starch matrix, causing the starch to recrystallize. Proteins that are more hydrophilic inhibit starch retrogradation by sequestering this water and reducing water mobility, thereby preventing recrystallization (Scott and Awika). Conversely, proteins that are more hydrophobic repel water from starch granules, promoting retrogradation (Scott and Awika).

Edible Coatings

Edible coatings are a variable and versatile tool to inhibit starch retrogradation. By preventing the movement of oxygen, water, and solutes without changing the bread's original composition, edible coatings preserve the quality of the bread from environmental damage while still being functional and economical (Ning et al.). Edible coatings are applied through dipping, spraying, and brushing onto food surfaces (Deseta et al.). They are already commonly used in fresh fruits and vegetables such as apples, oranges, and onions to maximize shelf life and quality (Perez-Vazquez et al.). In general, edible coatings should be hydrophobic and only slightly permeable to prevent the movement of oxygen and water molecules that is associated with starch retrogradation. Common materials used in edible coating include polysaccharides, proteins, resins, lipids and combinations of the above (Chen et al.).

Since the bread crust is sterile immediately after the baking process, exposure to contaminating spores on the surface must be kept at a minimum after the bread is made. Compared to other bread preservation methods such as enzymes and refrigeration, edible coatings have much greater potential to incorporate antimicrobial substances due to their ability to directly bind with them (Ning et al.). These antimicrobial compounds disrupt cell wall and membrane functions, slowing the growth of bacteria or fungi. These coatings' effectiveness has led to potential microbial resistance. Additionally, in recent years, microbial resistance and broad perceptions of these preservatives have harmed consumer acceptability, who demand natural and effective solutions (Deseta et al.). Edible coatings formed by natural substances such as polysaccharides, proteins, and lipids do not have any side effects on humans after consumption (Lee et al.). Therefore, this section will discuss five different naturally-sourced edible coatings that have anti-retrogradation or antimicrobial properties.

Chen et al. discovered that a 20% candelilla wax coating (solution of 20% candelilla wax in sunflower oil), 20% beeswax coating (solution of 20% beeswax in sunflower oil) and HPMC oleogel coating (coating containing hydroxypropyl methyl cellulose oleogel) were the most effective out of all the wax treatments they tried due to their hydrophobic nature and low permeability (Chen et al.). In this study, bread weight loss reduced from 30% to 13% and crumb

firmness (the maximum force during bread compression) reduced from over 500 N to 34 N after two weeks with the candelilla wax treatment.

Additionally, (Chen et al.) compared moisture loss and crust texture associated with enzymatic and chemical crust treatments, α-amylase and NaOH respectively, against those of edible coatings. As discussed in the enzymes section of this paper, α-amylase is an effective and commonly used protein that reduces retrogradation (Wang et al.). NaOH is a chemical that is frequently used in the pretzel industry in order to promote brown coloration through the Maillard reaction, which may also promote moisture retention (Chen et al.). This study found that α-amylase can randomly hydrolyze starches, creating a disrupted structure that does not retain water as well as the candelilla wax coating. Similarly, NaOH was found to have hydrolyzed proteins in the crust, which formed inter- and intramolecular bonds through the Maillard reaction while baking. These bonds made the bread denser and firmer, which was found to have reduced water migration and thus displayed a more significant reduction in retrogradation and water migration away from the bread. However, this effect comes at the expense of a less appealing bread appearance, which can also cause food waste. Due to protein hydrolysis and the deformation of the gluten network, NaOH caused the bread crust to become dotted with pits. This deformation is in combination with the hard texture and browning (like a pretzel) due to the Maillard reaction. α-amylase also causes a much darker crust because of the glucose release from enzymatic hydrolysis of the starches, creating more reactants for the Maillard reaction to further progress. As such, (Chen et al.) concluded that the smooth crust resulting from the candelilla wax coating, in conjunction with its anti-retrogradation properties, make it an effective treatment to prevent bread waste.

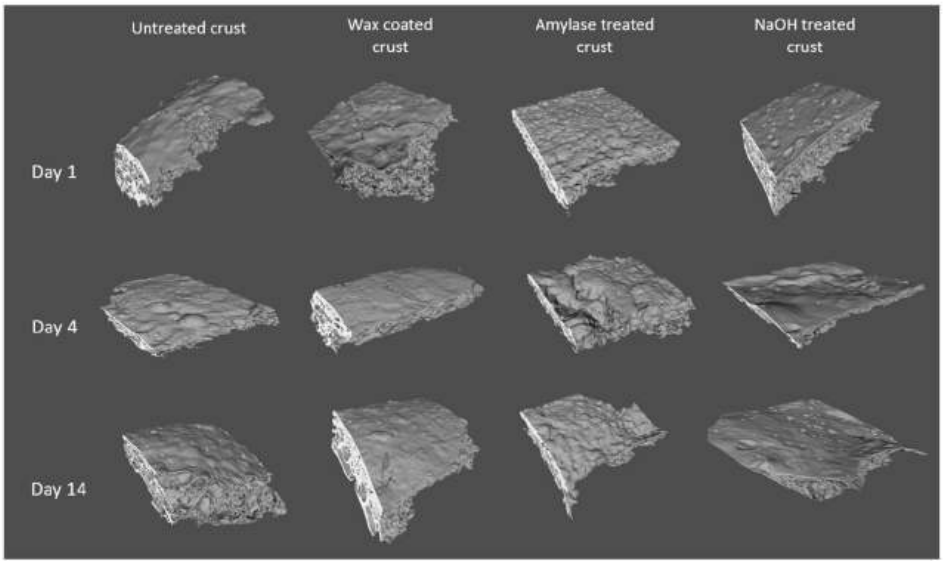

Fig 3: Crust surface texture following different coating treatments (wax, amylase, and NaOH) (Chen et al.)

Another edible coating method that uses egg white protein nanoparticles was explored by (Deseta et al.), which fought the problem of microbiological contamination as well as starch retrogradation. Both of these difficulties constitute the main cause of bread waste in the world.

Egg white protein nanoparticles (EWPn) are a good choice from which to create an edible coating due to their relatively low cost, high abundance, and effective film-forming properties (Deseta et al.). The researchers analyzed the encapsulation of bioactive compounds (BC) carvacrol, thymol, and trans-cinnamaldehyde (the main constituent of cinnamon essential oil) by EWPn nanocomplexes. Bioactive compounds have been proven as natural preservatives that can inhibit fungal growth. The study found that EWPn-BC nanoparticles were more effective in preventing fungal growth, especially the A. niger species, than pure BC ethanol solutions. In bread coated with EWPn-carvacrol and EWPn-thymol, antifungal properties extended the shelf life of preservative free-bread to seven days, doubling the normal shelf life of preservative free bread. Unfortunately, when the researchers explored the anti-starch retrogradation effects of EWPn-BC nanocomplexes, they found that these coatings influenced neither starch retrogradation nor texture.

In another study by (Alexieva et al.), the antimicrobial and anti-retrogradation properties of Malva sylvestris L. (mallow) flowers' extract on the edible coatings composed of three types of polysaccharides, pectin (P), xanthan (x), and carboxymethylcellulose (C). The addition of mallow extract in the xanthan coating strengthened its fungicidal ability, and its presence in xanthan and carboxymethylcellulose strengthened their yeasticidal properties. All three coatings exhibited the greatest moisture retention when mallow extract was incorporated, with samples losing 2.40% moisture in the crumb without the extract, compared to 0.67% with the extract on the third day. In the xanthan and pectin coatings, mallow extract significantly retarded the staling of the crumb, with the greatest effect observed in the xanthan coating (Alexieva et al.).

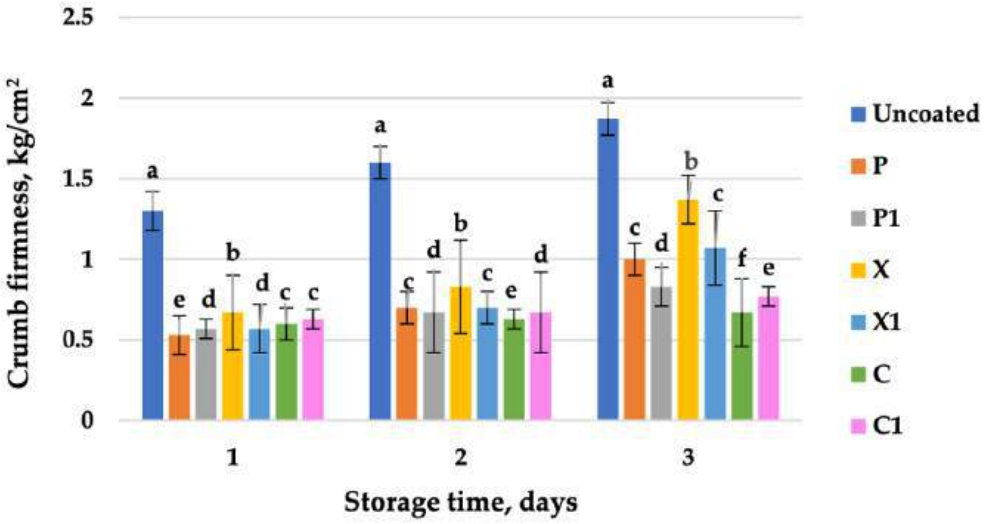

Fig 4: Changes in crumb firmness of uncoated and coated bread during 20°C storage for three days. Coatings include pectin (P), xanthan (X), and carboxymethylcellulose (C) with/without mallow flower extract (1) (Alexieva et al.).

(Lee et al.) experimented with mung bean starch and guar gum edible coatings to slow retrogradation and microbial growth. Starch-based coatings, such as mung-bean derived ones, are

commonly used due to their transparency, film-forming ability, cost effectiveness, and bland taste. Compared to other starch-based coatings, mung bean starch shows further research potential due to its ability to create a more dense and cohesive film. Mung bean starch and other starch-based coatings are brittle and easily allow water to leave the product, so other gums and lipids are required to prevent retrogradation. Guar gum is a polysaccharide that can be incorporated into the mung bean starch edible coating, increasing the viscosity of oil droplets within the bread and thus trapping them in its aqueous matrix. These oil droplets are hydrophobic, which can improve the ability of the moisture barrier of the aqueous matrix to retain water (Lee et al.).

Rice cake hardness decreased by 29% and crystallization rate decreased by 24% with the presence of mung bean starch coatings. Adding 0.8% grapefruit seed extract also strengthened the antimicrobial activity within the rice cakes, as shown with the decrease of B. cereus and P. citrinum populations by 0.68 log CFU/g and 1.18 log CFU/g, respectively. Although 2% grapefruit seed extract exhibited higher antimicrobial properties, the researchers discovered that 0.8% concentration was optimal because it had no physical impacts on the rice cake (Lee et al.).

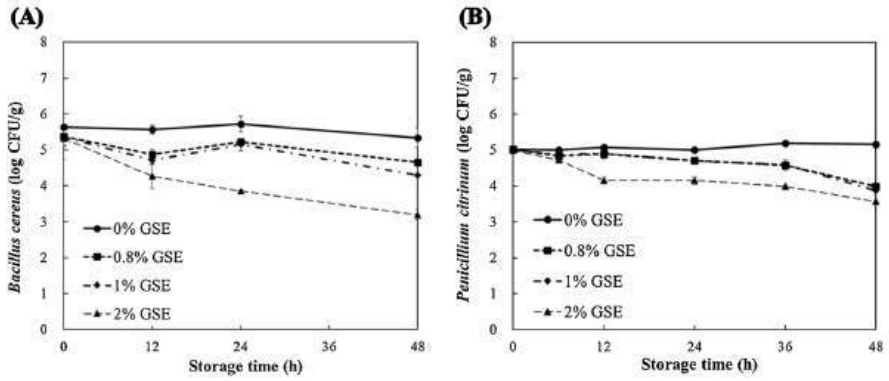

Fig 5: Growth curves of Bacillus cereus (A) and Penicillium citrinum (B) in rice cakes with emulsion coatings containing 0% (●), 0.8% (■), 1% (◆), and 2% (▲) grapefruit seed extract (GSE) (Lee et al.).

It is difficult to compare the antimicrobial properties of the mallow flower extract and grapefruit seed extract, as (Lee et al.) attempted to simulate real conditions after baking. Thus, the grapefruit seed extract experiment relied on submerging the bread in bacteria and fungi suspensions and then applying the coating. (Alexieva et al.), on the other hand, sterilized the bread and then immediately applied the mallow flower extract coatings before putting the bread under controlled, humid conditions for three days. This study did not submerge the bread in bacteria, and on the first day they did not detect any of their tested bacteria. Over the next three days, they observed that no colony-forming units formed. In this case, it seemed as though mallow flower extract prevented the formation of bacteria colonies, while it seemed as though in (Lee et al.) the grapefruit seed extract actively decreased the number of bacteria. Because these methods were different, a conclusion cannot be reached as to which antibacterial is more effective. There may also be a difference in effect due to differences in rice cake and bread

composition, so additional testing of the grapefruit seed extract's antimicrobial properties on bread are needed in the future. Regardless of challenges in comparing these methods, both methods show much potential as cheap, natural, and nontoxic antimicrobial agents in bread.

Conclusion

Since bread is a staple food consumed in every part of the world by people of varying socioeconomic statuses, it is imperative that methods to improve bread shelf life remain as cheap, practical, and effective as possible while being to work at room temperature. Substantial food waste has been observed due to the inability of vendors to sell bread before it goes bad (Niu et al.). The two main causes of decreased shelf life are starch retrogradation and microbial spoilage (Deseta et al.). This paper advocates for edible coatings as the optimal method for preventing these two events because of its noninvasiveness and cost-effectiveness (Ning et al.).

This paper first explored chemical advances in starch retrogradation to compare their implementation and effectiveness with edible coatings. Chemicals such as β-amylase were found to have high production costs and complex extraction processes, lowering their economic feasibility (Niu et al.).

Next, this paper examined different types of edible coating and their anti-retrogradation and antimicrobial properties. (Chen et al.) discovered that a 20% candelilla wax coating could significantly reduce bread weight loss and crumb firmness compared to normal chemical methods, but the study did not explore its antimicrobial effectiveness. (Deseta et al.) found that egg white protein nanoparticles could encapsulate bioactive compounds and create an edible coating that is much more effective against microbes than bioactive compounds in ethanol solutions. Thus, this paper suggests that mallow flower extract in xanthan edible coating is an effective method to reduce food waste as much as possible via anti-retrogradation as well as antimicrobial methods. (Lee et al.) proved that mung bean starch and guar gum edible coatings could also exhibit antimicrobial properties in addition to their anti-retrogradation ability if used in conjunction with 0.8% grapefruit seed extract. This study opens a possible route of investigation regarding the mallow flower and xanthan edible coating, as grapefruit seed extract may strengthen the antimicrobial property of the original formulation.

Bread waste is also highly nuanced, with type of bread and socioeconomic factors contributing as well. In Australia, bread rolls, sliced bread and flat bread are the three most wasted bread products in households. Since the studies mentioned in this paper each only used one type of bread, it will be interesting to explore how breads with different surface area and shape interact differently with each edible coating. Additionally, higher grocery shopping frequency, "just-in-case" buying, and to-go meal ordering are associated with higher bread waste (Ananda et al.). The extension of shelf life provided by edible coatings may affect each of these factors differently, "just-in-case" buying may be more affected by longer shelf life than the amount of children families have. Edible coatings and bread preservatives in general cannot solve the bread waste problem in its entirety, but it is still a significant step in the right direction.

Works Cited

Mohammadi-Moghaddam, Toktam, et al. “Bread Staling Measurement Techniques: A Review.” Journal of Food Chemistry & Nanotechnology, vol. 10, May 2024. DOI.org (Crossref), https://doi.org/10.17756/jfcn.2024-178.

Alexieva, Iordanka, et al. “Development and Application of Edible Coatings with Malva Sylvestris L. Extract to Extend Shelf-Life of Small Loaf.” Foods, vol. 11, no. 23, Nov. 2022, p. 3831. DOI.org (Crossref), https://doi.org/10.3390/foods11233831.

Ananda, Jayanath, et al. “Breaking Bread: Assessment of Household Bread Waste Incidence and Behavioural Drivers.” Journal of Cleaner Production, vol. 471, Sept. 2024, p. 143377. ScienceDirect, https://doi.org/10.1016/j.jclepro.2024.143377.

Chen, Yi, et al. “Crust Treatments to Reduce Bread Staling.” Current Research in Food Science, vol. 4, 2021, pp. 182–90. DOI.org (Crossref), https://doi.org/10.1016/j.crfs.2021.03.004.

Deseta, María Laura, et al. “Nanocomplexes Based on Egg White Protein Nanoparticles and Bioactive Compounds as Antifungal Edible Coatings to Extend Bread Shelf Life.” Food Research International, vol. 148, Oct. 2021, p. 110597. DOI.org (Crossref), https://doi.org/10.1016/j.foodres.2021.110597.

Garcia, Marcelo Valle, et al. “The Fungal Problem in Bread Production: Insights of Causes, Consequences, and Control Methods.” Current Opinion in Food Science, Trends in sustainable food processing • Food Mycology, vol. 29, Oct. 2019, pp. 1–6. ScienceDirect, https://doi.org/10.1016/j.cofs.2019.06.010.

Hafyan, Rendra Hakim, et al. “Bread Waste Valorization: A Review of Sustainability Aspects and Challenges.” Frontiers in Sustainable Food Systems, vol. 8, Feb. 2024. Frontiers, https://doi.org/10.3389/fsufs.2024.1334801.

Lee, Eun-sil, et al. “Effects of Mung Bean Starch/Guar Gum-Based Edible Emulsion Coatings on the Staling and Safety of Rice Cakes.” Carbohydrate Polymers, vol. 247, Nov. 2020, p. 116696. DOI.org (Crossref), https://doi.org/10.1016/j.carbpol.2020.116696.

Liu, Shuyun, et al. “Enhancing Bread Shelf-Life with Black Rice Anthocyanins: Optimization, Baking Properties, and Starch Retrogradation.” Journal of Cereal Science, vol. 127, Feb. 2026, p. 104365. DOI.org (Crossref), https://doi.org/10.1016/j.jcs.2026.104365.

N, Supriya. “Retrogradation of Starch - Meaning, Process, Examples & Factors Affecting.” Biology Reader, 17 Feb. 2022, https://biologyreader.com/retrogradation-of-starch.html.

Ning, Jiyang, et al. “Starch Retrogradation in Starch-Based Foods: Mechanisms, Influencing Factors, and Mitigation Strategies.” International Journal of Biological Macromolecules, vol. 300, Apr. 2025, p. 140354. DOI.org (Crossref), https://doi.org/10.1016/j.ijbiomac.2025.140354.

Niu, Wanyi, et al. “Recent Advances on Anti-Retrogradation Technologies of Starch-Based Foods.” Trends in Food Science & Technology, vol. 168, Feb. 2026, p. 105480. DOI.org (Crossref), https://doi.org/10.1016/j.tifs.2025.105480.

Ollinger, Nicole, et al. “Mycotoxin Contamination in Moldy Slices of Bread Is Mostly Limited to the Immediate Vicinity of the Visible Infestation.” Food Chemistry: X, vol. 23, June 2024, p. 101563. PubMed Central, https://doi.org/10.1016/j.fochx.2024.101563.

Perez-Vazquez, Ana, et al. “Edible Coatings as a Natural Packaging System to Improve Fruit and Vegetable Shelf Life and Quality.” Foods, vol. 12, no. 19, Sept. 2023, p. 3570. PubMed Central, https://doi.org/10.3390/foods12193570.

Scott, Gabrielle, and Joseph M. Awika. “Effect of Protein–Starch Interactions on Starch Retrogradation and Implications for Food Product Quality.” Comprehensive Reviews in Food Science and Food Safety, vol. 22, no. 3, May 2023, pp. 2081–111. DOI.org (Crossref), https://doi.org/10.1111/1541-4337.13141.

Shu, Ling, et al. “Starch Retrogradation in Potato Cells: Structure and in Vitro Digestion Paradigm.” Carbohydrate Polymers, vol. 286, June 2022, p. 119261. ScienceDirect, https://doi.org/10.1016/j.carbpol.2022.119261.

Tao, Han, et al. “Insight into the Promoted Recrystallization and Water Distribution of Bread by Removing Starch Granule - Surface and - Associated Proteins during Storage.” Food Chemistry, vol. 446, July 2024, p. 138829. DOI.org (Crossref), https://doi.org/10.1016/j.foodchem.2024.138829.

Wang, Yu-chuan, et al. “Enzymatic Modification of Wheat Starch by a Novel Maltotetraose-Forming Amylase from Atopomonas Hussainii to Retard Retrogradation and Improve Bread Quality.” Carbohydrate Polymers, vol. 348, Jan. 2025, p. 122909. DOI.org (Crossref), https://doi.org/10.1016/j.carbpol.2024.122909.

Modeling 1D Plasma Transport & Magnetic Shear Dynamics Using Burgers' Equation as an Analytical Proxy By Ethan Meline

Abstract

We develop a deliberately simplified one-dimensional toy model in radius and time to distinguish two mechanisms that can produce steep radial gradients: nonlinear convective steepening, and locally reduced transport associated with a reversed-shear magnetic geometry. The model is built in stages: linear advection, advection-diffusion, the inviscid and viscous Burgers' equations, and finally linear advection with a spatially varying diffusivity tied to a prescribed reversed-shear safety-factor profile. Each stage is solved with an explicit finite-difference scheme and compared against analytical or theoretical diagnostics. The inviscid Burgers' stage produces convective wave breaking at the theoretically predicted breaking time. When uniform diffusion is restored, shock formation is suppressed, illustrating the competition between nonlinear convection and dissipation. In the shear-dependent stage, diffusivity is prescribed to be lowest near the minimum of the safety factor, so the transported profile slows and steepens relative to a uniform-diffusivity control. This is a qualitative illustration of internal-transport-barrier phenomenology, not a predictive tokamak transport model. The construction separates amplitude-driven convective steepening from geometry-driven suppression in magnetic-confinement fusion plasmas.

Introduction

Controlled thermonuclear fusion offers a path toward a dense, low-greenhouse-gas-emission energy source by reproducing the reactions that power stars. Since the 1950s, the leading magnetic confinement concept has been the tokamak: a toroidal device that uses strong magnetic fields to hold a hot plasma away from material walls long enough for fusion reactions to occur. Magnetic confinement is required because a plasma at fusion-relevant temperature cannot be contained by a solid vessel; charged particles follow magnetic field lines, so a closed toroidal field geometry is used to isolate the plasma from the wall. The conditions required for ignition are extreme. The Lawson criterion demands combinations of sufficient density, temperature, and energy confinement time; in practice this corresponds to core temperatures of order 15 keV, or 8 roughly 10 K (Lawson). At such temperatures, confinement depends almost entirely on magnetic geometry.

A tokamak plasma is not uniform. The core must be substantially hotter and denser than the edge, producing macroscopic radial gradients in temperature and density. These gradients are thermodynamically necessary but also destabilising. They drive microinstabilites–among them ion temperature gradient (ITG) modes, electron temperature gradient (ETG) modes, trapped electron modes (TEM), and drift wave instabilities–that generate turbulent fluctuations and enhance radial energy and particle transport (Garbet et al.). The resulting anomalous transport far exceeds the energy and particle fluxes predicted by classical collisional theory and is widely regarded as the principal obstacle to achieving economical fusion performance

(Bickerton).

In the past decades, through experiments and theory, scientists have identified control parameters that help control this turbulent transport. One of the most important is **magnetic shear.** Within tokamak research, the safety factor q(r) measures how many times a field line warps around the toroidal length per each poloidal length. Effectively, it measures the ratio of toroidal to poloidal flux gradients. **Magnetic shear** describes how q varies with radius,

$$s(r) = \frac{r}{q}\frac{dq}{dr}$$

When shear is strong, turbulent mode growth rates can be reduced, (Drake et al.) and combined with E × B flow shear, magnetic shear can quench fluctuation and allow **internal transport barriers (ITBs)** to form. An ITB is a radial region of locally suppressed heat and particle transport in which temperature or density profiles become locally steep (Burrell). Reversed-shear configurations, in which q(r) has a minimum q_{min} in the plasma interior, are of particular interest because transport can fall sharply around the q_{min} surface, improving core confinement without relying on edge barriers alone (Levinton et al.; Strait et al.; Fujita et al.).

Understanding and predicting anomalous transport has motivated a hierarchy of theoretical and numerical tools. At the most fundamental level, gyrokinetic codes resolve the microscopic dynamics of turbulence. Global gyrokinetic codes such as GENE and trapped-gyro-Landau-fluid (TGLF) transport models, have made substantial progress in connecting linear and nonlinear microinstability physics to turbulent flux levels (Görler et al.; Kinsey et al.). These approaches are indispensable for questions such as which modes dominate, at what wavenumber, and at what saturation amplitude.

Yet many practical and conceptual questions concern not individual modes but **macroscopic radial transport**: how net heat and particle flux respond when equilibrium profiles, magnetic shear, or heating change. A common strategy that is in integrated modeling and quasilinear theory is to average over turbulent scales and represent the effects of microinstabilities through the use of **transport coefficients** (diffusivity, pinch terms, and convective distribution) in reduced transport equations. Full gyrokinetic simulations return fields and fluxes in high dimensional numerical data, while reduced transport models return **equations** whose terms can be inspected, manipulated, and compared across different cases. This distinction matters when the goal is to study a physical behavior rather than reproduce a complete discharge.

Recent work illustrates this macroscopic layer of the hierarchy. TORAX implements coupled core transport equations in a fast, differentiable framework suitable for scenario exploration (Citrin et al.). Data-driven neural network models such as QLKNN predict heat and particle flux from local equilibrium quantities (van de Plassche et al.). These tools sit between first-principles turbulent simulations and experiment: they assume and learn from local plasma conditions, then evolve profiles over time. What they typically do not emphasize, because it is buried inside effective transport coefficients or neural networks mapping, is the **nonlinear convective structure** of the transport problem itself: how steepening, shock-like profile formation, and competition between advection and diffusivity phenomena arise even in the

simplest radial models.

That gap motivates the present work. **The research question is: in a deliberately minimal one-dimensional radial model, can convective steepening and shear-induced transport suppression be isolated as two distinct routes to steep profiles, and can each be checked against a simple analytical or control-case diagnostic?** We use one spatial coordinate for the minor-radius direction and one time coordinate. The physics is staged from the simplest case upward: linear advection, for which the method of characteristics yields and exact solution; advection-diffusion, representing turbulent spreading; nonlinear advection through Burgers' equation, isolating convective steepening and wave breaking; and finally a prescribed reversed-shear q(r) profile with a shear-dependent diffusivity intended to *mimic*, rather than derive, ITB formation.

Burgers' equation is used here as a tractable nonlinear PDE that contains both convective steepening and dissipation, not as a claim that tokamak turbulence is Burgers flow (Burgers). Its role is to provide a controlled setting in which steeping and dissipation can be analysed before a shear-dependent transport suppression is added.

Compared to global gyrokinetics models, our approach sacrifices spatial dimensionality, electromagnetic detail, and self-consistent turbulence. Compared to integrated transport codes, it also simplifies the equilibrium, sources, and coupling between different fields. In return, it offers three clear advantages. The model isolates convective steepening and shear-induced transport barriers without the added complexity of other simulations. Its equations also make the connection between the mathematical terms and physical transport processes explicit, making the underlying dynamics easier to interpret. Most importantly, the model helps illustrate why the steep profiles needed for fusion are continually smoothed by transport, and how magnetic shear can locally counteract this smoothing. The following section describes the model framework.

Methods

To examine nonlinear convective steepening in turbulent plasma transport, we developed a one spatial dimension and time toy model using an explicit finite-difference transport scheme. The model progresses from the simplest linear advection equation, to advection-diffusion to represent turbulent spreading, then to the nonlinear inviscid Burgers' equation, and finally couples radial transport to a magnetic shear through a shear-dependent diffusivity. The final model is intended to capture the basic physics of internal transport barriers (ITBs). Throughout the model, *r* represents the minor-radius direction (from the magnetic axis towards the plasma edge), and *t* represents time. All variables are treated in dimensionless form.

Linear Advection and the Methods of Characteristics

We first consider pure convective transport with constant advection speed a > 0:

$$\frac{\partial T}{\partial t} + a\frac{\partial T}{\partial r} = 0 \qquad (1)$$

where T(r, t) is a transported scalar profile (representing a quantity such as plasma temperature

or density), r is the radial coordinate, and a is a constant representing a steady outward radial convection speed. Eq. (1) states that the local time rate of change of T at a fixed radial position is driven entirely by the spatial advection of the profile at speed a.

For this linear equation, the Method of Characteristics provides an exact solution (Evans). The Method of Characteristics works by finding so-called characteristic curves, which are specific paths in the (r,t) plane along which the PDE reduces to an ordinary differential equation. For Eq. (1), characteristics are straight lines r (t) = at + r_o or equivalently , r - at = constant, and the solution is simply the initial profile:

$$T(r,t) = \varphi(r - at) \qquad (2)$$

where φ (r) = T(r, 0) is the initial condition. Physically, Eq. (2) means that a localised perturbation retains its shape and translates in r at speed a. There is no spreading, no flattening, and no creation of steep gradients beyond those present initially.
In our model, we take a Gaussian pulse centered at r = r_0

$$T(r,0) = \varphi(r) = exp[-(r - r_0)^2] \qquad (3)$$

Which represents a localised blob of heat or density deposited in the inner plasma. The exact solution (2) with (3) serves as a reference for the numerical advection scheme we're then developing.

Adding Diffusion: the Linear Advection-Diffusion Equation

Real plasma transport is not purely convective. Turbulent processes act to smooth gradients, so we next add a diffusion term with a constant coefficient χ:

$$\frac{\partial T}{\partial t} + a\frac{\partial T}{\partial r} = \chi\frac{\partial^2 T}{\partial r^2} \qquad (4)$$

The left-hand side describes directed transport, while the right-hand side represents Fickian diffusion. Increasing χ causes the profile to broaden and flatten more quickly. We solve Eq.(4) numerically for simplicity. This model highlights the central competition between the two processes: convection transports the profile, while diffusion spreads it toward equilibrium. Setting a = 0 isolates pure diffusion; setting χ = 0 isolates pure advection.

Nonlinear Advection: the Inviscid Burgers' Equation and Wave Breaking

To study convective steepening in isolation, we remove diffusion (χ = 0) and replace the constant advection speed term with a speed that depends on the local value of T itself. The transport speed becomes v_r = T , giving the inviscid Burgers' equation:

$$\frac{\partial T}{\partial t} + T\frac{\partial T}{\partial r} = 0 \qquad (5)$$

Eq. (5) is still a first-order equation, but it is nonlinear, meaning that different parts of the profile propagate at different speeds. Using the same Gaussian initial condition (3), the leading edge of the pulse (where ∂T/∂r < 0 is most negative) outruns the trailing edge, and characteristic curves

that were initially parallel begin to converge.

The Method of Characteristics still applies, but now each characteristic carries the constant value of T it started with and propagates at a speed equal to that value. Where neighbouring characteristics intersect, the inviscid solution would develop a vertical slope i.e. a shock in (r,t). The earliest time at which this occurs can be estimated from the initial profile. If $T_0(r) = \partial T(r,0)/\partial r$, then

$$t_b = -\frac{1}{min_r T_0(r)} \qquad (6)$$

For the Gaussian (3) with $r_0 = 2$ the steepest negative slope lies on the right flank of the pulse, so steepening first develops there. For $t < t_b$, the solution remains smooth; for $t > t_b$, the inviscid equation becomes ill-posed and must be regularised by dissipation or by a shock-capturing numerical scheme. We use Eq. (6) as a theoretical diagnostic and compare it against the time evolution of $\partial T/\partial r$ from the numerical solution.

Viscous Regularisation: the Viscous Burgers' Equation

In a real plasma, steep gradients are limited by dissipative effects. We therefore reintroduce diffusion now coupled to the nonlinear advection term:

$$\frac{\partial T}{\partial t} + T\frac{\partial T}{\partial r} = \chi\frac{\partial^2 T}{\partial r^2} \quad (7)$$

Eq. (7) is the viscous Burgers' equation. As $\chi \to 0$ it recovers Eq. (5); for $\chi > 0$ the diffusivity term prevents the formation of a true discontinuity and replaces the inviscid shock with a smooth but steep internal layer. This is the form solved in the default nonlinear configuration of the baseline model. It allows us to follow the pulse through the steepening phase and into a regularised shock-like structure, while retaining enough dissipation for the explicit finite-difference scheme to remain stable.

The linear case (1) with $\chi = 0$ and the nonlinear case (5) with $\chi = 0$ bracket two ideal limits; Eq. (7) and Eq. (4) are the physically regularised versions used for numerical integration.

Magnetic Shear and Spatially Varying Diffusivity

Having established a baseline transport solver, we extend the model to include magnetic equilibrium effects relevant to ITB formation. The full tokamak transport problem is three-dimensional and flux-surface averaged; here we retain only the radial dependency and the magnetic structure through the safety factor q(r), defined as the number of toroidal transit of a field line per poloidal transit. We adopt a reversed-shear parabolic profile centered at r_0

$$q(r) = q_{min} + \alpha(r - r_0)^2, \ \alpha = \frac{q_{edge} - q_{min}}{(r_{max} - r_0)^2} \quad (8)$$

where q_{min} is the minimum safety factor at the intended barrier location and q_{edge} is the edge

value at r = r_{max} . This produces reversed magnetic shear (q decreasing with r) for r< r_0 and normal positive shear for r > $r._0$
From q (r) we compute magnetic shear:

$$\hat{s} = \frac{r}{q(r)} \frac{dq}{dr} \quad (9)$$

which is zero r = r_0 , negative for r < r_0 , and positive for r > r_0. Turbulent transport is then represented by a phenomenological shear-dependent diffusivity. Equation (10) is not derived from linear gyrokinetic growth rates. It is prescribed closure that places a Gaussian dip in χ at the shear-reversal surface ŝ = 0 (that is, at q_{min}), following the experimental observation that reversed-shear ITBs organize around q_{min} . Where |ŝ | is large, x recovers the background value x_{bg}; at ŝ = 0 , x falls to a floor x_{min} :

$$\chi(r) = \chi_{bg} - (\chi_{bg} - \chi_{min}) \, exp[-\frac{\hat{s}(r)^2}{2\sigma_s^2}] \quad (10)$$

with σ_2 controlling the width of the barrier in ŝ . The transport equation becomes:

$$\frac{\partial T}{\partial t} + a\frac{\partial T}{\partial r} = \frac{\partial}{\partial r}[\chi(r) \frac{\partial T}{\partial r}] \quad (11)$$

where advection is held linear (a is a constant). The Burgers' nonlinearity is not combined with shear in the current extension, so any steepening in this stage is caused by the imposed χ(r), not by amplitude-dependent wave breaking.

Initial and Boundary Conditions

All simulations use the Gaussian pulse (3) as an initial condition unless otherwise noted. Homogeneous Dirichlet boundary conditions $T(r_{min}, t) = T(r_{max,} t) = 0$ are imposed, representing the rapid loss of the transported quantity at the axis and plasma boundary. Howev simulations are run for times short enough that the choice of boundary conditions does not affe the inside dynamics.

Numerical Implementation.

All variables are treated as dimensionless. The radial coordinate r is scaled so that the computational domain $r \in [0, r_{max}]$ with $r_{max} = 10$ corresponds to the interval from the magnetic axis to the plasma edge. A tokamak minor radius a is used. Time is scaled so that an advective speed of order unity crosses a distance of order unity in a time of order unity. The transported scalar T is normalized to its initial peak, so $T(r, 0)$ has a maximum value of 1. So locations such as $r = 2.71$ are therefore positions in this dimensionless minor-radius coordinate, no SI metrics, and they depend on the chosen initial pulse location $r_0 = 2.$

The equations are solved in Python/NumPy with explicit finite differences on $N_r = 300$ uniform nodes ($\Delta r \approx 0.0334$). Time stepping is forward Euler. Advection is first-order upwind:

$$(v\partial_r T)_i \approx v_i \frac{T_i - T_{i-1}}{\Delta r},$$

with $v_i = a$ (linear) or $v_i = T_i$ (Burgers). Uniform diffusion uses a second-order central stencil. $\chi(r)$ uses $\Gamma_{i+1/2} = \chi_{i+1/2}(T_{i+1} - T_i)/\Delta r$ with $\chi_{i+1/2} = (\chi_i + \chi_{i+1})/2$. Homogeneous Dirichlet conditions $T = 0$ are imposed at both ends. The time step is

$$\Delta t = 0.4 min(\frac{\Delta r}{max|v|}, \frac{(\Delta r)^2}{2\chi_{mac}}).$$

Default values are: $a = 1$, $T(r,0) = exp[-(r-2)^2]$, $\chi = 0$ *or* 1; shear stage $q_{min} = 2$, $q_{edge} = 6$, barrier at $r_0 = 4$, $\chi_{bg} = 1$, $\chi_{min} = 0.05$, $\sigma_s = 0.5$.

Results & Discussion

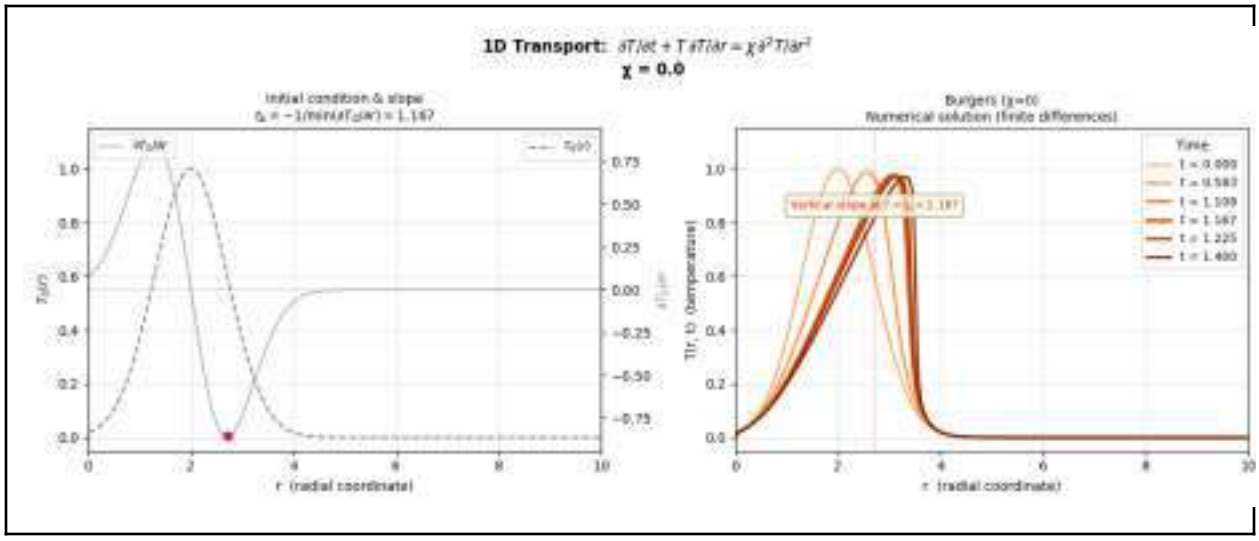

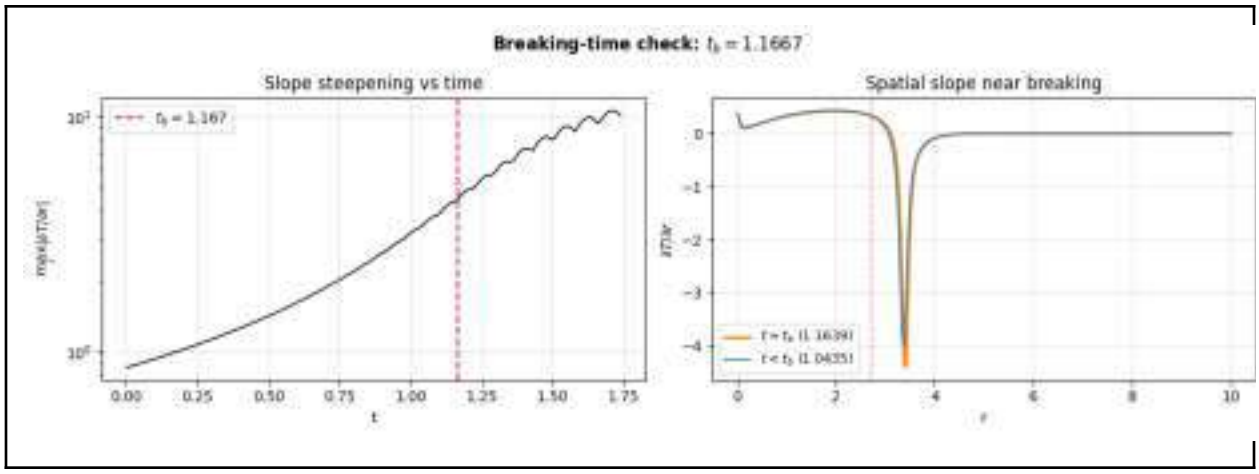

The numerical simulation of the inviscid Burgers' Equation ($\chi = 0.0$) confirms the theoretical prediction of nonlinear convective steepening leading to wave breaking. For the initial Gaussian profile $T_0(R) = \exp[-(r-2)^2]$, shown as a dashed curve in *Figure 1a,* the spatial slope $\partial T_0/\partial r$ reaches its steepest negative value at $r = 2.71$, where $\partial T_0/\partial r \approx -0.857$ (marked by the red dot). Eq. (6) predicts an inviscid breaking time of $t_b = 1.1667$ theoretical Method-of-Characteristics breaking time for the discrete initial slope.

Because the radial transport speed $v_r = T$ scales directly with local profile amplitude, higher amplitude regions move faster than the lower amplitude trailing regions. As time advances (*Figure 1b*), the outer flank of the pulse progressively steepens while the trailing edge flattens. Equation (6) provides a useful prediction for the onset of steepening, although it does not precisely predict the numerical shock time. At $t \approx t_b = 1.1667$, $\max_r(\partial T/\partial r) = 4.41$ slope continues to grow and peaks at 10.63 at $t = 1.7125$, showing a lag of 46.8% in t_b. That delay is expected as first-order upwind adds diffusion, so the discrete front cannot break as early as the inviscid characteristic solution. Equation is therefore a useful diagnostic for where and when steepening should begin, not a claim that the scheme reproduces the inviscid breaking time to high accuracy.

Figure 2 tracks $\max_r(\partial T/\partial r)$ versus time and shows that steepening is underway at t_b but that the peak slope occurs later, consistent with the 47% lag reported above.

Viscous Regularization and Competition with Diffusion ($\chi = 1.0$)

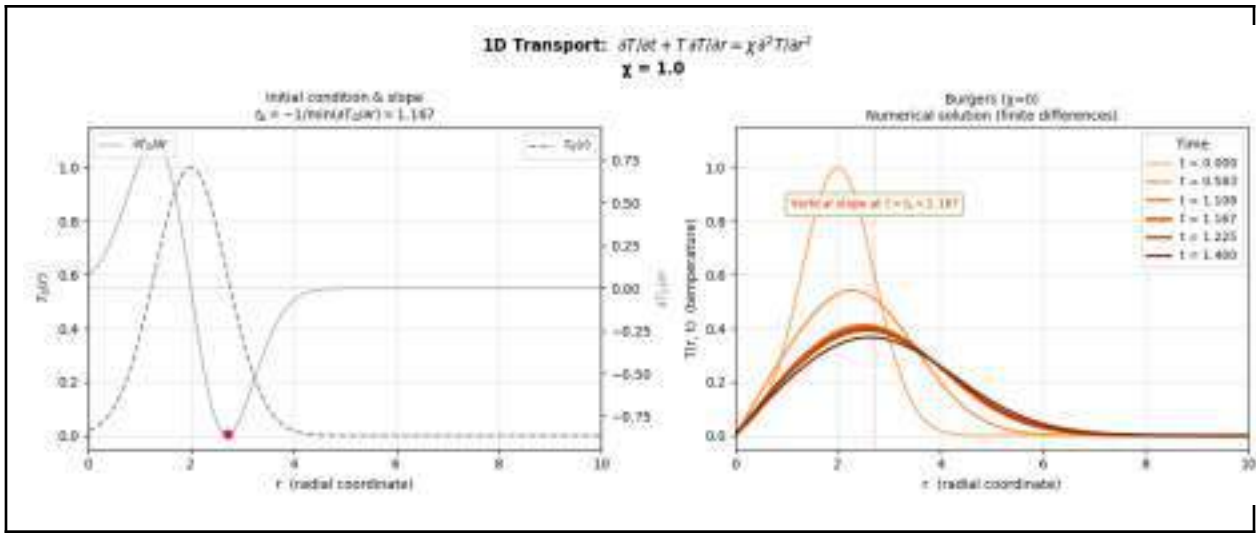

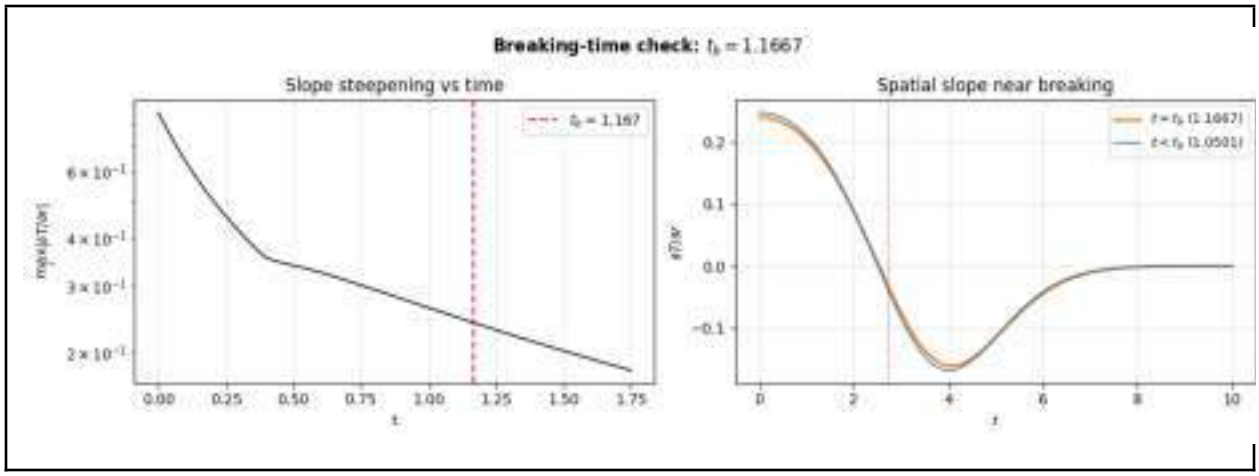

Reintroducing a uniform background diffusivity ($\chi = 1$, Eq. 7) qualitatively changes this outcome. The initial maximum slope is the same Gaussian value $max_r\,(\partial T/\partial r) = 0.857$. After a brief early transient (peak 0. 72 at t= 0. 058), the gradient then falls, reaching ≈ 0. 17 at t = 1. 75. There is no late-time steepening near t_b. If the inviscid peak-slope diagnostic is applied anyway, the time of maximum slope is t = 0. 058, a 95% discrepancy from $t_b = 1.1667$. That failure is correctly justified by the presence of the diffusion term $\chi\partial^2 T/\partial r^2$ in Eq.(7), which erodes gradients faster than the nonlinearity can sharpen them. The spatial slope at $t \approx t_b$ is nearly indistinguishable from its shape at t ≈ 1. 05, confirming that no steep layer forms. The pulse amplitude decays from 1. 0 to ≈ 0. 4 while spreading radially, rather than compressing into a front. Comparing $\chi = 0.0$ and $\chi = 1.0$ side by side, we isolate the diffusion-advection competition described in the Methods. We have identical initial data, and identical nonlinear terms, but the presence or absence of the right side dissipation term in Eq.(7) is the sole determinant of whether the solution shocks or relaxes.

Magnetic Shear and Internal Transport Barrier

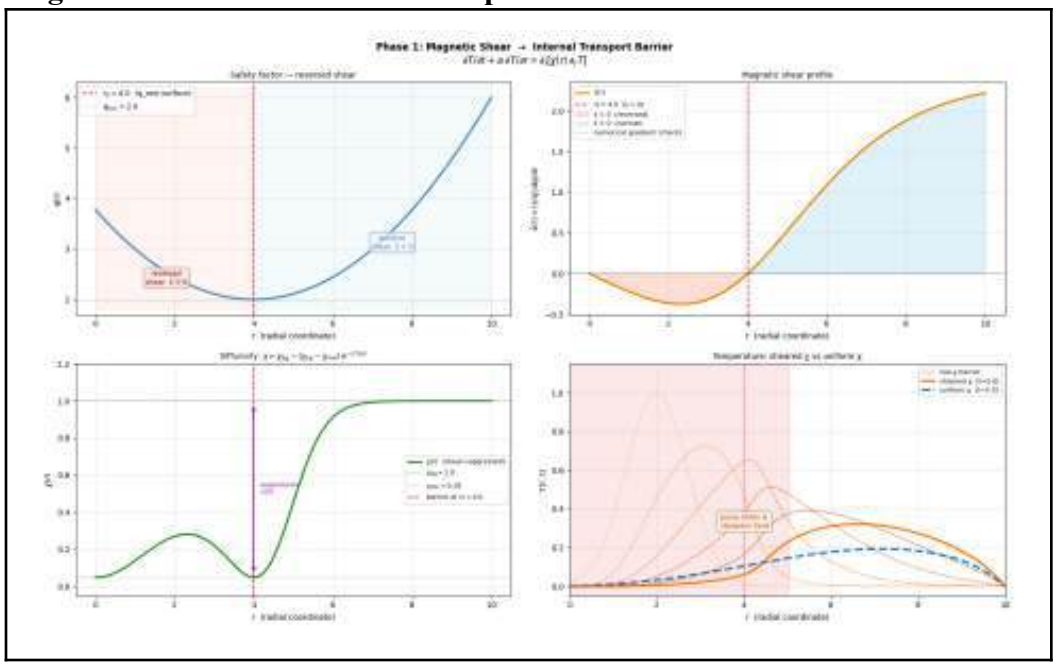

The reversed-shear safety factor profile (Eq.8) produces q(r) decreasing from q(0) ≈ 3. 8 to $q_{min} = 2$ at $r_0 = 4$ then rising to $q(r_{rmax}) = 6$. The magnetic shear profile computed from Eq.(9) crosses 0 at exactly $r_0 = 4$ and is negative in the reversed shear region (down to ≈ − 0. 4) , and positive (rising to ≈ 2. 2) beyond it. The resulting shear suppressed diffusivity χ(r) from Eq.(10) forms at dip centered at $r_0 = 4$ falling from the background value $\chi_{bg} = 1.0$ close to the neoclassical floor $\chi_{min} = 0.05$, with a local suppression of a factor of 20. This happens precisely at ŝ = 0, reproducing the phenomenon that transport barriers form when magnetic shear reverses and locally reduces turbulent diffusivity.

The consequence for radial transport is visible in the T (r,t) comparison between sheared and uniform diffusivity of Figure 4c. Both simultaneously start from the same pulse launched at $r_0 = 2$ and advect outwards. Beyond $r_0 = 4$ they both diverge. The control profile continues to advect and diffuse, arriving at a low amplitude profile (peaks ≈ 0. 2 at t = 5). The sheared profile visibly slows and steepens as it encounters the low χ regions; it retains a narrow, higher amplitude peak (≈ 0. 33 at t = 5) at ≈ 65% higher than the control profile. The amplitude and delay difference, obtained from otherwise similar initial and boundary conditions, isolates the effect of shear-induced transport suppression from ordinary convective-diffusive spreading, as intended by the control comparison described in the Methods.

The four-stage hierarchy isolates two distinct routes to steep radial gradients, letting each be attributed to its correct physical origin rather than lumped together as generic “barrier” behavior.

In the inviscid Burgers’ case (χ = 0), steepening is purely convective: since $v_r = T$, higher amplitude fluid outruns the trailing edge (figure 1b). The theoretical $t_b = 1.1667$ from Eq.(6), marks the start of rapid steepening in Figure 2, but the numerical peak slope occurs later (t = 1.

7125), consistent with the dissipation of a first-order upwind scheme. This is the same convective mechanism present in avalanche-like bursts in tokamaks (Diamond and Hahm), understood here as a qualitative analogue rather than a quantitatively timed shock.

Reintroducing diffusion ($\chi = 1$, Figure 3) removes this behavior: the profile decays (T_{peak} :1. 0 → 0. 40 at t = 1. 167 → 0. 37 at t = 1. 4) and broadens rather than sharpening, showing diffusion dominates the nonlinearity at this χ. So t_b is best read as an upper bound on unregularized steepening, not a general predictor.

The shear coupled run (Figure 4) shows a second, distinct mechanism: the pulse stalls and steepens not because of its own amplitude but because $\chi(r)$ collapses independently at the q_{min} surface ($r_0 = 4$) . The resulting barrier (with $T_{peak} \approx$ 0. 33vs. 0.20 for the uniform x control at T = 5) is a linear diffusivity effect, not self steepening.

Conclusion

The shear diffusivity closure (Eq.10) is grounded in established ITB phenomenology. Reversed shear was first linked to strongly improved core confinement (Levinton et al.)(Strait et al.)(Fujita et al.); in each case the improved confinement region traced q_min rather than a fixed radius consistent with the geometric organization built into Eq. (10). Negative shear was identified as a local stabilizing mechanism (Kessel et al.), later confirmed as a major ITB ingredient alongside E × B shear (Connor et al.)(Garbet et al).

The convective steepening stage connects to a separate strand of the literature on non-diffusive, avalanche like transport (Diamond et al.). proposed that transport near marginal stability proceeds via intermittent, ballistic events rather than local diffusion. This picture is supported by radial-propagation modelling showing that nonlinearities can steepen and accelerate transport fronts in a similar manner than the Burgers' dynamics we reproduced here. The convection-diffusion competition in Eq.(7) mirrors the avalanche versus diffusion competition
that determines whether transport in a near marginal plasma looks locally diffusive or non-Fickian (Sánchez and Newman).

The numerical approach chosen in Eq.(11) follows standard practice for hyperbolic/advection-diffusion systems with sharp diffusion (Leveque). The model is strictly 1D, omitting flux-surface geometry and poloidal asymmetry known to matter for real ITBs. The Burgers' nonlinearity V_r= T is a convenient proxy for convective transport, not a quantity derived from gyrokinetics. The shear closure (Eq.10) is static and lacks feedback of real ITBs. The explicit scheme used for the inviscid Burgers' case is only marginally stable near t_b.

Despite these simplifications, the staged construction provides a controlled way to attribute steep-gradient formation to its correct physical cause before extending towards more realistic transport solvers, making the model useful as a diagnostic and pedagogical tool rather than a quantitative predictor for any specific device. Future works could extend this model to include a more realistic turbulence-driven approach, providing a pathway from the present conceptual model toward quantitatively predicted simulations.

Works Cited

Lawson, J. D. “Some Criteria for a Power Producing Thermonuclear Reactor.” Proceedings of the Physical Society. Section B, vol. 70, no. 1, 1957, pp. 6–10.

Bickerton, R. J. “Magnetic Turbulence and the Transport of Energy and Particles in Tokamaks.” Plasma Physics and Controlled Fusion, vol. 39, 1997, p. 339.

Drake, J. F., et al. “Local Negative Shear and the Formation of Transport Barriers.” Physical Review Letters, vol. 77, no. 3, 1996, p. 494.

Burrell, K. H. “Effects of E×B Velocity Shear and Magnetic Shear on Turbulence and Transport in Magnetic Confinement Devices.” Physics of Plasmas, vol. 4, no. 5, 1997, pp. 1499–1518.

Görler, T., et al. “The Global Version of the Gyrokinetic Turbulence Code GENE.” Journal of Computational Physics, vol. 230, no. 18, 2011, pp. 7053–71.

Kinsey, J. E., et al. “The First Transport Code Simulations Using the Trapped Gyro-Landau-Fluid Model.” Physics of Plasmas, vol. 15, no. 5, 2008, article 055908.

Citrin, Jonathan, et al. “TORAX: A Fast and Differentiable Tokamak Transport Simulator in JAX.” arXiv, 2024.

Van de Plassche, K. L., et al. “Fast Modeling of Turbulent Transport in Fusion Plasmas Using Neural Networks.” Physics of Plasmas, vol. 27, no. 2, 2020.

Burgers, J. M. The Nonlinear Diffusion Equation. Springer Science & Business Media, 2013.
Evans, L. C. Partial Differential Equations. American Mathematical Society, 1998.

Diamond, P. H., and T. S. Hahm. “On the Dynamics of Turbulent Transport Near Marginal Stability.” Physics of Plasmas, vol. 2, no. 10, 1995, pp. 3640–49.

Levinton, et al. “Improved Confinement with Reversed Magnetic Shear in TFTR.” Physical Review Letters, vol. 75, no. 24, 1995, pp. 4417–20.

Strait, et al. “Enhanced Confinement and Stability in DIII-D Discharges with Reversed Magnetic Shear.” Physical Review Letters, vol. 75, no. 24, 1995, pp. 4421–24.

Fujita, et al. “Isolation of the Improved Core Confinement from High Recycling and Radiative Boundary in Reversed Magnetic Shear Plasmas of JT-60U.” Physical Review Letters, vol. 78, no. 7, 1997, pp. 1267–70.

Kessel, et al. “Improved Plasma Performance in Tokamaks with Negative Magnetic Shear.” Physical Review Letters, vol. 72, no. 8, 1994, pp. 1212–15.

Connor, J. W., and R. J. Hastie. “Microinstability in Tokamaks with Low Magnetic Shear.” Plasma Physics and Controlled Fusion, vol. 46, no. 10, 2004, pp. 1501–35.

Garbet, et al. “Physics of Transport Barriers in Tokamaks.” Plasma Physics and Controlled Fusion, vol. 46, no. 12B, 2004, pp. B557–74.

Diamond, P. H., and T. S. Hahm. “On the Dynamics of Turbulent Transport Near Marginal Stability.” Physics of Plasmas, vol. 2, no. 10, 1995, pp. 3640–49.

Sánchez, R., and D. E. Newman. “Self-Organized Criticality and the Dynamics of Near-Marginal Plasma Transport.” Plasma Physics and Controlled Fusion, vol. 57, no. 12, 2015, article 123002.

Leveque, R. J. Finite Volume Methods for Hyperbolic Problems. Cambridge UP, 2011.

Systematic Algorithmic Optimization: A case study on Cookie Clicker
By Namish Balakerthy Punyakoti

Abstract

This paper presents a systematic framework that makes great use of dynamic programming concepts to solve a complex, deterministic optimization problem found in the incremental game *Cookie Clicker* as efficiently as possible. The game is modeled as a state space search problem, where an agent (the player) must determine the optimal sequence of actions (building and upgrading purchase inputs), to reach the user-selected number of Cookies Baked (end goal) in as short a time as possible. The complexity resides in the sheer quantity of potential input sequences, as each purchase changes the game state, leading to a different future set of affordable actions. To address this, a graph-based model of the game's decision space is constructed and a systematic algorithm incorporating Pareto efficiency is employed for search space pruning to minimize the amount of combinations needing manual evaluation, and even those which aren't pruned are accounted for as dynamic programming principles are utilized to make their evaluation computationally efficient. The performance of this algorithm is empirically evaluated against several greedy heuristics within a custom Python-based simulation. The findings demonstrate that the principled algorithm significantly outperforms heuristic baselines. Finally, the computational complexity of the problem is discussed, drawing parallels to classic NP-hard problems like the Traveling Salesman Problem to contextualize the algorithmic trade-offs.

1. Introduction

Incremental games present an ideal computational challenge since like their nickname of 'idle-games' suggests, their core game loop amounts to carefully making discrete decisions to optimize a continuously accumulating currency, thus serving as an ideal testbed for exploring sophisticated optimization algorithms within such a domain since, the game's deterministic nature where the only variable affecting the rate of progress is the sequence of player inputs allows for a direct and quantifiable measurement of algorithmic efficiency, that being the number of frames required to reach a pre-defined goal.

This investigation models a restricted version of *Cookie Clicker* as a formal optimization problem, with the value being optimized for being the time, T, in terms of frames elapsed, to reach a user-defined 'Cookies Baked' goal, G.

This paper details the design and implementation of an algorithm that takes advantage of the nuances between the typical travelling salesman problem and the game's version of it to identify and then prune branches that can be proven to be definitively suboptimal. This proof is founded on concepts from graph theory and economics (specifically, the Pareto frontier), and its utilization makes the search for the definitively optimal path computationally feasible unlike the traditional NP-Hard travelling salesman problem.

2. Problem Definition and Background

2.1 Gamestate Representation

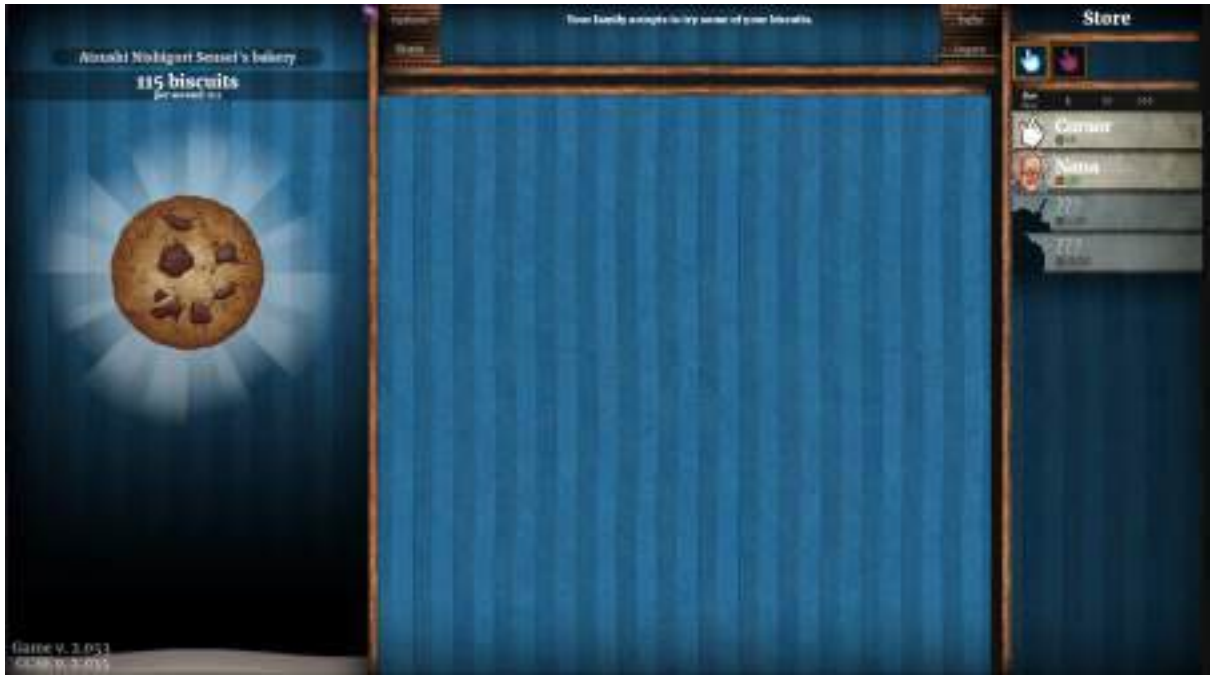

Figure 1: *A screenshot illustrating the game's interface. The Big Cookie is positioned on the left, the building buy options on the right with the amount currently owned indicated and the upgrade buy options above them.*

This paper exclusively interacts with the building and upgrade purchase systems in addition to 'clicking on the big cookie' as all other mechanics contain elements of randomness that compromise the perfect information our algorithm is dependent on.

A 'Frame' is the base unit for time within videogames and Cookie Clicker updates 30 times per second, meaning each frame in Cookie Clicker is 1/30th of a second or roughly 33 milliseconds. Optimal gameplay requires consistent frame precise inputs so rather than manually input actions, the optimal sequence of inputs is pre-calculated and then executed at the precise time at which they've been scheduled for within a python-based model of the game, which accounts for the feature restrictions mentioned before.

2.2 Clicking on The Big Cookie

In Cookie Clicker you can click on something known as 'The Big Cookie' to increment both the CookiesBaked and CookiesInBank statistics by a default of one. The difference between them is that CookiesInBank tracks only the cookies currently owned, whereas CookiesBaked tracks cookies gained over the entire lifetime of a playthrough thus making it far more convenient to use for the endpoint.

The game's code allows for multiple inputs to be made on the same frame meaning there is no opportunity cost associated with clicking on The Big Cookie, and since it always brings

gamestates closer to the goal, optimality demands maximization of this mechanism. Therefore clicking on The Big Cookie in all algorithms is modelled as a deterministic event that occurs at the earliest time possible, which is the game's hardcoded limit of every 40 milliseconds.

2.3 Building Purchases

When purchased, buildings automatically produce a set amount of Cookies per frame. There are multiple building types and in each gamestate there can be any quantity of a combination of building types.

This is where the rapid combinatorial explosion comes from, luckily however, each of the building types have their own unique core statistics, that being their BaseCost and BaseCookiesPerSecond, and these stats are quite high in the later building types meaning that many building purchases will not have to be considered as potential actions since their cost will be greater than that of the goal itself and as such can't be afforded before the simulation reaches the goal and terminates.

Below is a list of the first 5 building types

In game sprite	Name of building type	Base Cost (Cookies)	BaseCPS (Cookies Per Second)
	Cursor	15	0.1
	Grandma	100	1
	Farm	1,100	8
	Mine	12,000	47
	Factory	130,000	260

Table 1: *A list of the game's first 5 building types out of its total 20.*

2.4 Upgrade Purchases

Unlike buildings which are complex objects, upgrades are one-time purchases with a fixed price that simply act as a scalars in the formulae for their corresponding building type's Cookie production.

Upgrades do however need to be unlocked before they can be bought, which entails owning at least as many of their corresponding building type as their unlock condition states. The approach in this paper takes advantage of this by initially omitting upgrades from consideration, instead reintroducing them by checking which combinations determined to be pareto optimal contain enough of the corresponding building type for that particular upgrade to be unlockable.

Additionally, purchasing upgrades for the Cursor building type will also double the amount of Cookies you get per click on the Big Cookie.

In game sprite	Upgrade Name	Base Cost (Cookies)	Unlock Condition
	Reinforced Index Finger	100	Own 1 Cursor
	Carpal Tunnel Prevention Cream	500	Own 1 Cursor
	Forwards from grandma	1,000	Own 1 grandma
	Steel-plated rolling pins	5,000	Own 5 grandmas
	Fertilizer	55,000	Own 5 farms
	Sugar gas	120,000	Own 1 mine

Table 2: *A table with some of the upgrades in the game.*

$$CPS_i = BaseCPS_i \cdot N_i \cdot 2^{U_i} \quad , N \in \mathbb{Z} \geq 0 , \; U \in \mathbb{Z} \geq 0$$

Where

CPS_i = The cookies from building type ‘i’ per second

N_i = The number of building type ‘i’ owned

$BaseCPS_i$ = The hardcoded production rate values for building type ‘i’

U_i = The number of owned upgrades for building type ‘i’

3. Algorithm Design and Methodology

3.1 Problem formulation

Cookies are continuously being produced by buildings and through clicking, however the rate at which they are produced is constant, only ever changing when a purchase, or in other words an input from the player, is made.

Therefore any path to the goal can be expressed in terms of the sequence of purchase inputs made in that path, since deterministic elements can be retroactively simulated, but the purchase type and timing which are the only non-deterministic elements in any playthrough can't be.

$$S = (E_1, E_2, \ldots, E_k)$$

Where :

E_k is a purchase input that is the 'kth' purchase made, which could be either:

- The 'nth' purchase of a building type 'i'
- The 'nth' upgrade purchase for the building type 'i'

For any goal 'G' there exists the set $\delta(G)$ which contains all the sequences of purchases to the goal, and simply evaluating each of these sequences will eventually yield the optimal sequence.

$$\delta(G) = \{S = (E_1, E_2, \ldots E_n) \mid n \in \mathbb{Z} \geq 0\}$$

However, all combinations of purchases for a given goal, $\mathcal{C}(G)$, must be found for all the permutations to be calculated.

$$\mathcal{C}(G) = \{C = (E_1, E_2, \ldots E_n) \mid n \in \mathbb{Z} \geq 0,\ \Sigma P(C) < Goal\}$$

Where:

- $\mathcal{C}(G)$ = All combinations of purchase inputs valid for a goal 'G'
- $\Sigma P(C)$ = The sum total price of all purchases made within a combination of purchases 'C'

3.2 Scope constraints

This is the critical nuance between the classic travelling salesman problem and the version of it presented in Cookie Clicker. In the travelling salesman problem you must visit all the cities or nodes, but in Cookie Clicker we are limited to a fraction of the possible purchase combinations in the game due to the scope of the goal itself, as any combination of purchases which cost exceeds that of the goal would be unaffordable in a playthrough as the goal would've been reached and the simulation terminated before then.

This makes the problem more computationally feasible as unless the goal is set to

incredibly high values, past the millions, many building types will not even be affordable within the scope of the goal.

To find all of the valid combinations of purchases we must ensure that

$$\Sigma P(U_C) + \Sigma P(B_C) < Goal \quad , Goal \in \mathbb{R} \geq 0$$

Where:

- $\Sigma P(U_C)$ = All the upgrade purchases made in the combination of purchases 'C'
- $\Sigma P(B_C)$ = All the building purchases made in the combination of purchases 'C'

Performing these validation calculations is simple for upgrades due to their fixed prices however, buildings scale in price with each building of that type already owned, and as such more complex formulae such as those shown below are employed to calculate their sum price and by extension the maximum number of them that can be afforded within the Goal.

Building Price Formula:

$$P_i = BasePrice_i \cdot 1.15^{n_i - 1} \quad , n_i \in \mathbb{Z} \geq 0$$

Where

P_i = Price of buying a building of building type 'i'

$BasePrice_i$ = The hardcoded base price value of the building type 'i'

n_i = The number of buildings of building type 'i' currently owned

Sum total price of n buildings formula:

$$S_{ni} = \frac{BasePrice_i \cdot (1.15^n - 1)}{0.15} \quad , n \in \mathbb{Z} \geq 0$$

Which can be summated to find the total cost of all buildings in a combination.

$$\Sigma P(C_b) = \sum_{i=1}^{m} \frac{BasePrice_i \cdot (1.15^n - 1)}{0.15}$$

Where

$\Sigma P(C_b)$ = The sum total cost of all building purchases made within a sequence 'S'

m = The total amount of building types within the scope of the goal (If a building's base cost is higher than the goal then it is out of the scope)
n = The number of buildings of type 'i' bought in the combination of purchases 'C'

Max Amount of a Building Type Within Goal Formula:

By rearranging

$$S_{ni} = \frac{BasePrice_i \cdot (1.15^n - 1)}{0.15} \quad , n \in \mathbb{Z} \geq 0$$

To isolate n, the following is produced

$$n_i^* = \lfloor Log_{1.15}[(Gap \cdot \frac{0.15}{BasePrice_i}) + 1] \rfloor \quad , Gap \in \mathbb{R} \geq 0$$

$$Gap = Goal - \Sigma P(C_b)$$

Where

n_i^* = The maximum number of building type 'i' that is affordable within the gap
Gap = The amount of cookies that can be spent before the sum total price of the purchases exceeds the goal
$\Sigma P(C_b)$ = The sum total price of all building purchases made in a set of purchases 'C'

Using these formulae the combinatorially explosive nature of this problem reveals itself, as even with relatively small goals like 2000 Cookies Baked, a maximum of 21 purchases of the cheapest building type are affordable within the scope of the goal, and since three different building types have BaseCosts lower than that of the goal many more combinations can be created by omitting some of these purchases of the cheapest building type and redistributing those cookies to make purchases of the other building types, additionally various permutations that can be formed from these purchases.

Fortunately by identifying the combinations that lie on the Pareto frontier all others can be definitively and cheaply proven to be inefficient, allowing for them to be pruned and for the costly process of evaluation to only be applied to the handful of combinations on the Pareto frontier.

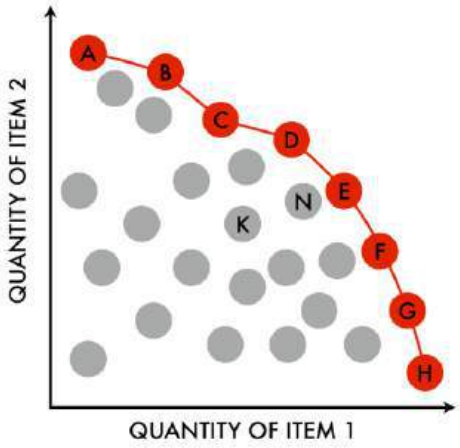

Figure 2: *A generic 2d, Pareto Frontier diagram*

The Pareto frontier describes a set of points on a graph, where an attempt to increment in one dimension requires the sacrificial decrement in another, and this almost perfectly models the decision making process in Cookie Clicker as some purchase combinations (red dots on fig 2) are trade offs, since the same cookies spent on one purchase could've been spent on a different, perhaps more optimal purchase, and there isn't any obvious way to prove that without manually evaluation, whereas other purchase combinations (grey dots on fig 2) opt to omit a purchase only to never reinvest the freed cookies in another purchase, leaving them unspent and thus definitively less efficient than the combinations on the Pareto frontier.

All points on the Pareto frontier can be found by first identifying one corner, by creating a purchase combination with maximum weightage given to one building type, and then systematically omitting purchases and redistributing the freed cookies from them to create alternate purchase combinations with increments in the other building types.

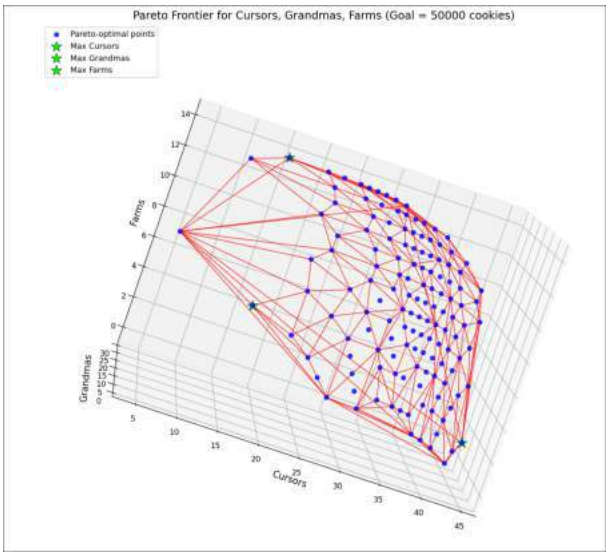

Figure 3: *A visualization of the Pareto frontier of Cookie Clicker that I created in matplotlib. It*

demonstrates that by starting at one of the corners, depicted in the figure via stars, and systematically omitting and redistributing the freed cookies the entirety of the Pareto frontier for this goal is created.

To avoid infinite loops and redundant calculations, when a purchase from a building type is omitted the freed cookies should never be redistributed to the same building type. Redistributions should also occur maximally, meaning if the omission of a purchase frees up enough cookies to make multiple purchases of the next building type then the maximum amount should be bought, and then the maximum amount of the next building type be bought with the leftover freed cookies and so on so forth. This is to ensure the combinations produced are truly Pareto optimal.

Upgrades in Cookie Clicker require more than just currency to be bought, an arbitrary amount of buildings of their corresponding building type must first be owned for the upgrade to be unlocked and purchaseable, and as such the Pareto optimal purchase combinations are finalized by checking to see which of them can unlock and afford to purchase upgrades within the scope.

Generating all permutations of a purchase combination isn't possible through simple combinatorics as upgrades have the unique feature of needing to be unlocked before they can be bought, meaning that every purchase sequence including an upgrade must have the upgrade purchase made after all of the building purchases for that upgrade to be unlocked have been made.

Manual evaluation of sequences of inputs has been referenced many times in this paper and below is the formula used to perform it. By applying it to all of the purchase sequences generated by the methodology detailed thus far, the optimum sequence of inputs which is guaranteed to be a part of this list will be revealed.

Time To Reach Goal Formula:

$$T_{Goal}(S) = \sum_{k=1}^{n} T(E_k \mid E_0, E_1, \ldots, E_{k-1}) + T_{Wait} \quad n \in \mathbb{Z} \geq 0,\ T \in \mathbb{Z},\ Goal \in \mathbb{R} \geq 0$$

Where:

$T_{Goal}(S)$ = Frames taken to reach the goal using the sequence of purchase inputs 'S'

n = The total number of purchases made before entering the wait phase

k = The last purchase to be made in the sequence

$T(E_k \mid E_0, E_1, \ldots, E_{k-1})$ = The amount of frames it would take for purchase 'k' to be afforded given the history of all purchases made thus far

T_{Wait} = The amount of frames in the waiting phase

This is where dynamic programming comes in, as the amount of time it takes for purchase 'k' to be afforded is dependent upon the amount of time it takes for purchase 'k-1' to be afforded and so on so forth, thus memoization was employed to cache segments that were

commonly shared across many sequences so that the answer may be reused rather than re-calculated at every instance, to make evaluation less costly and more computationally efficient and the entire problem far more computationally feasible for larger scoped goals.

Additional optimizations are present in the handling of ,T_{wait} , the period near the end of a simulation where all purchases have been made and the simulation simply needs to advance until the goal is reached. This idle period is guaranteed to exist for all sequences because the pruning removes all purchase combinations that result in purchases made on the exact frame of the simulation's termination since the purchase wouldn't provide any value at that point.

Purchases are always bought at the earliest instance, there is no benefit from delaying a purchase, opting to skip it in favour of another perhaps, but never any benefit from simply delaying it.

Combining this knowledge with information about the purchases in a sequence allows for, T_{wait}, to always be found cheaply, since the sum total cost and production rate increase from the purchases in a sequence can be calculated.

4. Implementation

4.1 Simulation Environment

The model game was made in Python 3.11. The game state is modeled through several interacting classes and functions that collectively track the player's cookie bank, total cookies baked, building counts, upgrade ownership, and production rate.

The core data structure is the BuildingType class, which encapsulates the static properties of each building type: its name, base price, and base cookies per second (CPS). The class provides methods for calculating bulk purchase costs using the geometric pricing formula, determining the cost of buying a specific quantity of buildings given the current ownership count, and finding the maximum number of buildings affordable within a given budget gap.

The Upgrade class similarly encapsulates upgrade properties, including its cost, the building type it affects, the unlock condition (minimum building count required), and whether it doubles click power.

The BuildingCombination class represents a single combination of building purchases. It stores the count of each building type owned, tracks the total cookies spent on that combination, and provides methods for adding purchases, omitting purchases (for the Pareto frontier traversal), and duplicating itself for branching operations.

The simulation advances by evaluating purchase sequences through the evaluate_sequence function. This function simulates the passage of time between purchases, accumulating cookies at the current production rate until sufficient funds are available for the next scheduled purchase. Each purchase is executed instantly, the production rate is updated, and the simulation continues. After all purchases are made, the simulation enters a final waiting phase where cookies accumulate until the goal is reached. The function returns the total number of frames required, calculated by multiplying total elapsed seconds by the frame rate and

applying a ceiling operation.

4.2 Data Structures

Two primary data structures support the Pareto frontier generation algorithm: 11

- Set for Frontier Storage: The frontier variable is a Python set that stores all discovered maximal building combinations. Because BuildingCombination implements hashing and equality based on its building counts tuple, membership checks are O(1), preventing redundant exploration of already-visited combinations.
- List as Processing Queue: The to_process variable is a Python list used as a stack (via .pop()) to maintain combinations that have been discovered but not yet explored. This depth-first traversal approach is memory-efficient and ensures all combinations are processed.

The upgrade subset generation uses itertools.combinations to enumerate all possible non-empty subsets of unlockable upgrades whose total cost fits within the remaining budget gap of a given building combination.

4.3 Baseline Algorithms

Two baseline strategies were implemented for comparison against the optimal algorithm:

- Cheapest-Building Greedy: This strategy meant always purchasing the cheapest available building type. It represents a naive approach where the player simply buys whatever is most immediately affordable without considering the benefits of long term strategy.
- Best-Ratio Greedy: This strategy calculates an efficiency ratio for each building type, defined as the increase in CPS divided by the purchase cost, and always buys the building with the highest ratio. This represents a more sophisticated heuristic that attempts to maximize immediate return on investment.

All greedy strategies generate their sequence in a single pass, making decisions based only on the current game state without considering future consequences. This is the fundamental difference between the heuristic approaches and the exhaustive optimal search.

5. Results

5.1 Experimental Setup

All algorithms were evaluated using a user-defined goal value, which must be greater than 15 cookies (the base cost of the cheapest building). The optimal algorithm performs an exhaustive search over all Pareto-optimal building combinations, all valid upgrade subsets for each combination, and all valid permutations of each resulting purchase set. The greedy algorithms generate their sequences in a single pass without branching.

5.2 Performance Comparison

The output of the simulation provides a direct comparison of all algorithms. A few examples are shown below.

Frames Taken vs Compute Time Algorithmic Performance Comparison Across Multiple Goals								
	Frames Taken To Reach End Goal				Compute Time (x10^{-5} seconds)			
Cookies Baked Goal	Optimal	Click Only	Cheapest Building	Best Ratio	Optimal	Click Only	Cheapest Building	Best Ratio
250	218	300	295	292	479.4	1.2	6.7	4.8
1000	549	1200	1136	1064	391610. 1	0.3	5.2	4.4
1250	620	1500	1396	1301	329869 7.5	0.4	7.3	6.3
1400	662	1680	1548	1438	941669. 74	0.5	8.2	7.6
1450	676	1740	1598	1581	132071. 027	0.5	7.5	2.9

Table 3: *A table comparing the performance of the algorithms across various Cookies Baked Goals*

```
Goal: 250.0 cookies
Algorithm                Frames      Purchases   Comp. Time (s)
-------------------------------------------------------------------
Optimal (exhaustive)     218         7           0.004794
Click Only               300         0           0.000012
Cheapest Building        295         8           0.000067
Best Ratio               292         4           0.000048
```

Fig 4: *A screenshot depicting the results the different algorithms produced for a goal of 250 Cookies*

```
Goal: 1000.0 cookies
Algorithm                Frames      Purchases   Comp. Time (s)
-------------------------------------------------------------------
Optimal (exhaustive)     549         13          3.916101
Click Only               1200        0           0.000003
Cheapest Building        1136        17          0.000052
Best Ratio               1064        11          0.000044
```

Fig 5: *A screenshot depicting the results the different algorithms produced for a goal of 1000 Cookies*

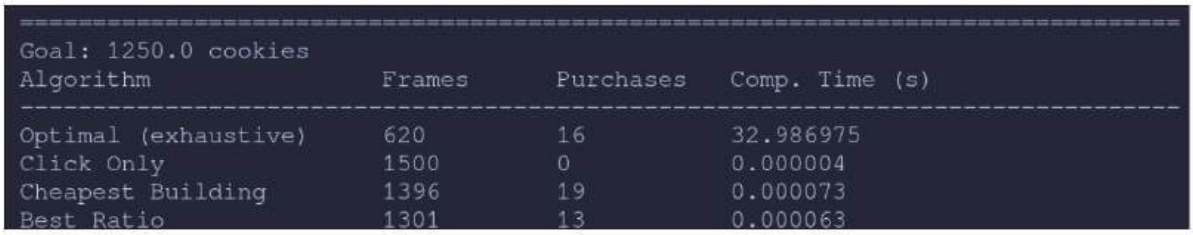

Goal: 1250.0 cookies

Algorithm	Frames	Purchases	Comp. Time (s)
Optimal (exhaustive)	620	16	32.986975
Click Only	1500	0	0.000004
Cheapest Building	1396	19	0.000073
Best Ratio	1301	13	0.000063

Fig 6: *A screenshot depicting the results the different algorithms produced for a goal of 1250 Cookies*

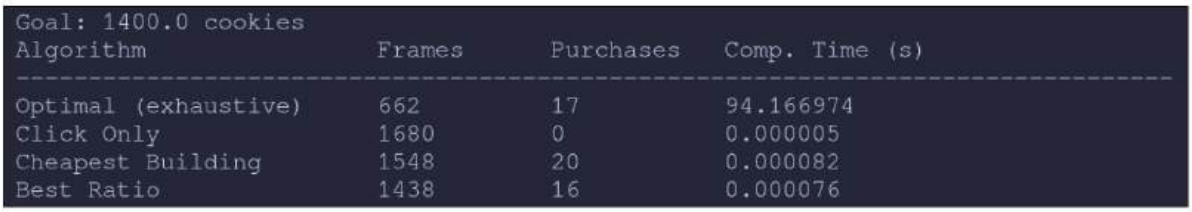

Goal: 1400.0 cookies

Algorithm	Frames	Purchases	Comp. Time (s)
Optimal (exhaustive)	662	17	94.166974
Click Only	1680	0	0.000005
Cheapest Building	1548	20	0.000082
Best Ratio	1438	16	0.000076

Fig 7: *A screenshot depicting the results the different algorithms produced for a goal of 1400 Cookies*

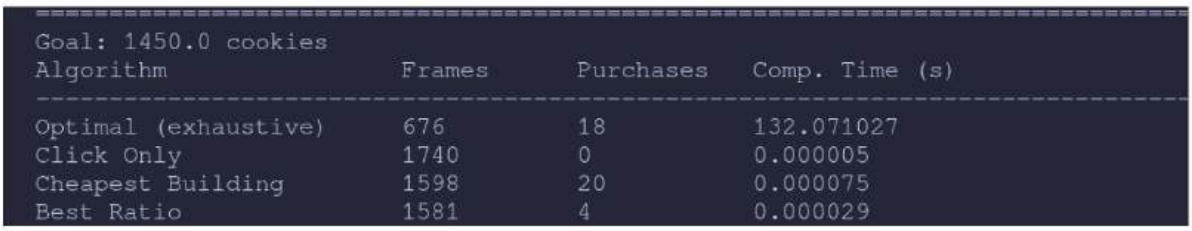

Goal: 1450.0 cookies

Algorithm	Frames	Purchases	Comp. Time (s)
Optimal (exhaustive)	676	18	132.071027
Click Only	1740	0	0.000005
Cheapest Building	1598	20	0.000075
Best Ratio	1581	4	0.000029

Fig 8: *A screenshot depicting the results the different algorithms produced for a goal of 1450 Cookies*

To no surprise the results clearly demonstrate that the optimal algorithm consistently achieves the goal in fewer frames than all greedy baselines. The magnitude of improvement varies depending on the goal value.

As shown in figure 4, at smaller goals like 250 cookies the time difference is marginal with 74 frames, or roughly two and a half seconds. However, as shown in figure 7, in larger goal simulations like 1400 Cookies the optimal path is more than twice as fast as the next best alternative. Quite predictably though this comes with a heavy cost, as it took a whopping 95 seconds to compute whereas the next best algorithm which was the Best Ratio algorithm took less than one-ten-thousandth of a second.

The "Click Only" algorithm serves as a baseline as it consistently performs the worst out of all the algorithms the deeper into the late game the simulation progresses and as such highlights the importance of building and upgrading purchases in Cookie Clicker.

Unfortunately despite all the optimizations made I was never able to get any results for goals of 1500 Cookies or higher within a reasonable amount of time (<15 mins), demonstrating that ultimately there will always be a huge drawback to optimality in the form of its computational time, even when the version of NP-Hard optimization problem present has nuances that can be exploited for additional efficiency.

5.3 Pruning Effectiveness

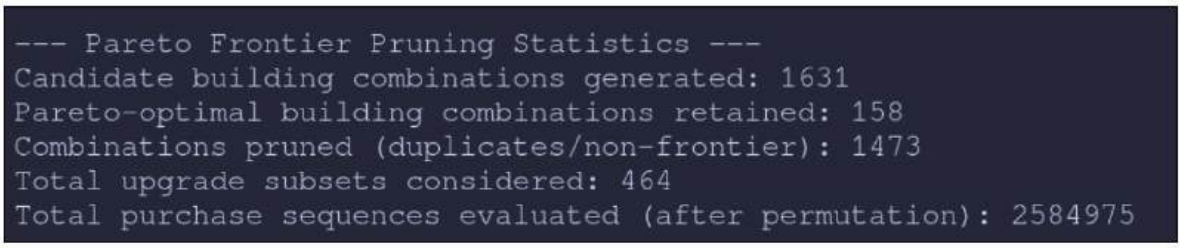

Fig 9: *A screenshot depicting the Pareto efficiency of the optimal algorithm for a goal of 1400 Cookies*

The Pareto frontier pruning technique is the key to making this exhaustive search algorithm computationally feasible. As shown in Figure 9, for a goal of 1400 cookies, the naive brute-force approach would need to consider all possible combinations of building purchases across three building types. Even with all the optimization the program still considered over 2.5 million candidates so without it the search likely wouldn't have been computationally feasible at all even for this goal that is so small it only allows 3 of Cookie Clicker's 20 building types within the scope.

The pruning effectiveness increases as the goal grows, because the proportion of non-maximal combinations also grows; however the worst case complexity scenario still remains exponential time rather than polynomial.

5.4 Optimal Sequence Details

The optimal algorithm also outputs the specific purchase sequence it identified, along with the final building counts and upgrade distribution. This provides insight into what optimal play actually looks like. The algorithm includes a tie-breaking rule: if multiple sequences achieve the same frame count, the sequence with less purchases is preferred, as it represents a strategy with less direct control and yet the same results, showing a higher degree of efficiency and strategy that would likely surge past the sequences it tied with should the goal be extended.

6. Discussion

6.1 Computational Complexity Analysis

The proposed algorithm achieves significant practical improvements over naive brute-force search through the application of Pareto frontier pruning. However at the end of the day, the worst-case time complexity remains exponential.

Let *m* be the number of building types within the scope of the goal G. For each building combination on the Pareto frontier, the algorithm must:

- Generate all valid upgrade subsets using generate_upgrade_subsets, which iterates over all combinations of unlockable upgrades. This is bounded by 2^u, where *u* is the number of unlockable upgrades within the goal's scope.
- Generate all valid permutations of the combined purchase set using the recursive

backtracking function generate_sequences. This is bounded by *k*!, where *k* is the total number of purchases (buildings plus upgrades) in the combination.

● Evaluate each permutation via evaluate_sequence, which simulates the game frame-by-frame. This requires O(k) time per sequence to process each purchase event.

The total number of Pareto-optimal combinations itself grows combinatorially with the number of building types in scope. This is because the number of ways to distribute a budget across *m* categories follows the stars-and-bars combinatorial formula, which is exponential in *m* for a fixed budget.

Therefore, the overall worst-case time complexity is:

$$T(G) = O(P(G) \cdot 2u \cdot k! \cdot k)$$

Where:

P(G) = the number of Pareto-optimal combinations for goal G.

This is consistent with the problem's structural similarity to other NP-hard combinatorial optimization problems. The results from Section 5 confirm this analysis: for a goal of 1400 cookies, the algorithm required 95 seconds of computation time and evaluated over 2.5 million candidate sequences, while the greedy baselines completed in under one-ten-thousandth of a second.

6.2 Practical Effectiveness of Pruning

Despite the exponential worst-case complexity, the Pareto frontier pruning proves highly effective in practice for the tested goal range. The effectiveness of this pruning technique grows as the goal increases because the proportion of non-maximal combinations in the total search space also grows. At small goals, most combinations are naturally maximal because the budget is tight and there are few ways to leave cookies unspent. At larger goals, the combinatorial explosion of possible combinations includes an increasing proportion of non-maximal ones, which the Pareto frontier efficiently eliminates.

Unfortunately this is still not enough for the program to be able to handle incredibly large goals because as shown it wasn’t able to compute an answer for a goal of 1500 within a reasonable amount of time.

6.3 Comparison to Traveling Salesman Problem

The problem of finding the optimal purchase sequence in Cookie Clicker bears structural similarity to the Traveling Salesman Problem. Both problems involve finding an optimal ordering of actions from a combinatorially large set. In TSP, the goal is to find the shortest cycle in a complete graph. In Cookie Clicker, the goal is to find the shortest path through the state space that reaches the endpoint.

However, a key difference exists that makes the Cookie Clicker problem more tractable in practice. In TSP, all nodes must be visited exactly once, meaning the solution is always a

permutation of the complete node set. In Cookie Clicker, the set of purchases is not fixed in advance, it is constrained by the goal budget. This constraint is what enables the Pareto frontier pruning technique to exploit the budget constraint to determine which purchase compositions are worth considering and in the process significantly lower the amount of combinations needing to be considered.

6.4 Limitations

Several limitations of this investigation should be acknowledged:

- Restricted Game Model: The simulation excludes several mechanics present in the full version of Cookie Clicker, including golden cookies, seasonal events, and the prestige system. These mechanics introduce random elements that would fundamentally alter the nature of the optimization problem from deterministic to probabilistic.
- Scalability Ceiling: Despite the effectiveness of Pareto pruning, the algorithm's exponential complexity means it cannot scale to arbitrarily large goals. The factorial growth in permutation generation becomes restrictive as the number of purchases in a combination exceeds roughly 15-20. This is reflected in the 95-second computation time for the 1400-cookie goal and the fact that the 1500 goal could not be computed within a reasonable length of time.
- Frame-Perfect Assumptions: The model assumes perfect information and frame-perfect execution. In practice, human players cannot achieve this level of precision, though the theoretical optimal remains valuable as a benchmark for what is achievable.

7. Conclusion

This investigation successfully modeled the game Cookie Clicker as a computational optimization problem and developed a systematic algorithm for finding the optimal sequence of purchase actions. The key contributions of this work are:

- Problem Formulation: The game was modeled as a state space search problem, with the objective of minimizing the time to reach a user-defined end goal, which in combination with the game’s deterministic mechanisms which produce precise and reproducible results, allowed for the time values to serve as a direct measurement of algorithmic efficiency.
- Pareto Frontier Pruning: Pareto pruning was implemented and greatly narrowed the search space by eliminating any combination that left cookies unspent, saving magnitudes of efficiency by doing so, due to the great cost of evaluation all of those combinations would’ve borne.
- Upgrade Handling: The algorithm integrates upgrade purchases by checking unlockability and affordability for each Pareto-optimal building combination, then generating all valid permutations that respect upgrade preconditions. This two-stage approach takes advantage of the game’s typically unnoticed upgrade requirements to once again narrow the search space and thus dampen the rate of combinatorial explosion

that would result from treating upgrades as unrestricted purchases.

- Dynamic Programming Optimizations: Memoization techniques were employed to cache intermediate time calculations across overlapping permutations requiring evaluation, lowering the cost of it by eliminating redundant computation.
- Empirical Evaluation: The algorithm was benchmarked against three greedy heuristics, demonstrating consistent performance improvements. The results show that the optimal algorithm achieves the goal in significantly fewer frames than any baseline, with the gap widening as the goal increases. For a goal of 1400 cookies, the optimal algorithm was more than twice as fast as the next best alternative.

The results validate the effectiveness of the algorithmic approach while highlighting the fundamental computational complexity of the problem. The worst-case time complexity remains exponential, consistent with the problem's structural similarity to NP-hard problems such as the Traveling Salesman Problem. The Pareto frontier pruning technique serves as an effective heuristic that takes advantage of the unique nuances of Cookie Clicker to narrow the search space but does not fundamentally alter the problem's complexity class.

The presence of combinatorial explosion leads to the developed optimal algorithm being impractical for usage of large goals, however the results produced do raise the question of whether, with time and exploitation of the other various nuances of Cookie Clicker over traditional NP-hard problems like the Traveling Salesman Problem, such definitively optimal algorithms can be practically applied in real time for relatively large goals in such contexts.

Regardless, simply having a definitively optimal algorithm to use as reference is of great value, since it can serve as a definitively correct, quantifiable measurement of accuracy for other algorithms.

Works Cited

Cormen, Thomas H., Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. Introduction to Algorithms. 3rd ed. Cambridge, MA: MIT Press, 2009.

Russell, Stuart, and Peter Norvig. Artificial Intelligence: A Modern Approach. 4th ed. Hoboken, NJ: Pearson, 2020.

Wikipedia, s.v. "Travelling salesman problem." Accessed September 3, 2026.https://en.wikipedia.org/wiki/Travelling_salesman_problem.

"Cookie Clicker Wiki." Cookie Clicker Wiki. Accessed September 2, 2026. https://cookieclicker.wiki.gg/wiki/Cookie_Clicker_Wiki.

"File:Pareto Efficient Frontier 1024x1024.png." Wikipedia. Accessed September 1, 2026. https://en.wikipedia.org/wiki/Pareto_front#/media/File:Pareto_Efficient_Frontier_1024x1024 .png.

Code availability

https://github.com/bakaboyplaystoomuchstarrail-dotcom/Cookie-Clicker-Projekt

Acknowledgment

The author acknowledges the use of AI-based tools to assist in the development of this paper. Specifically:

- Writing Assistance: Deepseek was used to help proofread and give suggestions on tone and focus for certain sections of this paper, including the data analysis and implementation; however all AI-generated text was reviewed and edited by the author for accuracy, clarity, and consistency with the paper's technical content.
- Data Analysis: No AI tools were used for data analysis. All results were generated by running the author's Python simulation code.

The author developed all formulae, algorithms and optimization strategies put forth in this paper, reserving the use of AI tools to exclusively to aid with execution of the project, more specifically AI-generated advice on the essay was taken as general feedback that the author used to direct the tone and key focus of certain sections rather than exemplar work to be copied directly, and the author takes full responsibility for the accuracy and integrity of all content in this paper.

What Shapes Entrepreneurial Intent? Active Learning and Multicultural Experiences among Japanese High School Students By Toma Ihara

Abstract

Japan has some of the lowest rates of early-stage entrepreneurial activity among developed economies, yet research on its causes has focused on university students rather than high schoolers. This study examined whether exposure to active learning environments and multicultural experiences is associated with entrepreneurial intent among high school students. A cross-sectional online survey collected 71 responses in August 2026 from respondents aged 15 to 19 or older in Japanese, International Baccalaureate, American AP, and other international curricula. The analysis compared curriculum groups on item means and standard deviations, then correlated each type of exposure with entrepreneurial intent and tested both in one regression.

Results indicated that entrepreneurial intent was lower among Japanese curriculum students than among American/AP students but similar to IB students. Multicultural exposure predicted entrepreneurial intent after controlling for active learning exposure, which showed no association. These findings suggest that curriculum reform focused on teaching methods, such as Japan's 2022 inquiry-based study period, may need to include intercultural contact to shift entrepreneurial orientation. Larger school-based samples and longitudinal designs are needed to test whether this relationship is causal.

Keywords Entrepreneurial intent, active learning, multicultural exposure, high school education, Japan

Introduction

The capacity to identify problems, generate original solutions, and act on them independently has shifted from the periphery of education policy toward its center, driven by labor markets in which routine analytical work is increasingly automated. Japan occupies an instructive position within this shift. Despite consistently strong academic outcomes, the Global Entrepreneurship Monitor has characterized entrepreneurial activity in Japan as disappointingly low, noting that the country's total early-stage entrepreneurial activity rate did not exceed 5.4 percent across a fourteen-year period, and attributing this primarily to cultural factors of conformity and failure avoidance (Yokoyama and Birchley). Such explanations, however, rest on an assumption that has rarely been tested directly: that the constraint lies in approaches to learning rather than in cultural expectations. If the constraint is embedded in approaches to learning, particularly in classroom experiences that emphasize textbook knowledge at the expense of independent inquiry, it may be amenable to change. Alternatively, if the constraint arises from cultural norms of conformity, students may be less willing to participate in innovative entrepreneurial activities even when the opportunities are presented. This distinction determines whether low entrepreneurial engagement is treated as an educational design problem to be addressed or a cultural condition to be accepted.

Research on entrepreneurial intention has largely converged on the Theory of Planned Behavior, which models intention as a function of attitude, subjective norm, and perceived behavioral control. Moriano et al., drawing on 1,074 students across six countries, found that attitude and perceived behavioral control predicted entrepreneurial career intention consistently across cultural contexts, while the effect of subjective norm varied significantly between them. The findings suggest that cultural differences may shape whether individuals feel socially supported in pursuing entrepreneurship, rather than whether they possess the capability to do so. Kim, Huruta, and Lee extended this framework to South Korean high school students, another exam-intensive system, confirming that perceived entrepreneurial capacity, perceived social norm, and attitude each predicted intention. In Japan, Kurata et al. surveyed 275 engineering students and found attitude and self-efficacy to be significant predictors, concluding that project-based and theory-based learning should be combined rather than substituted for one another. Collectively, this body of work identifies the learning environment as a candidate mechanism, though findings on its role are mixed, and largely measures it as a self-reported attribute within a single system rather than as a condition that varies.

Three major gaps emerge from the existing literature. First, research on entrepreneurial intention in Japan has concentrated on university and technical college populations, despite evidence that attitudes toward risk-taking, initiative, and self-directed action may become consolidated during secondary education. Second, the educational environment is typically treated as one construct measured inside one system, which limits the comparative analysis that would isolate its effect. Third, Japan introduced the Period for Inquiry-Based Cross-Disciplinary Study into the national high school curriculum in 2022, a reform intended to develop students' ability to find and solve problems through inquiry (MEXT), yet its relationship to entrepreneurial orientation remains unexamined. Japan provides an especially suitable context to address these gaps, because traditional national-curriculum schools, international programs such as Stanford e-Japan, and IB or international schools operate concurrently within the same country, producing sharply different learning conditions while holding nationality and broader social context relatively constant.

This study examines whether exposure to active learning environments and multicultural experience is associated with differences in entrepreneurial intent among Japanese high school students. Survey data were collected from 71 respondents in four curriculum groups (Japanese national, International Baccalaureate, American/AP, and other international curricula), and exposure to active learning and multicultural experience was measured for each respondent. Rather than relying on self-assessed confidence, which may be susceptible to social-desirability bias or overestimation, the research instrument prioritized behavioral indicators. Respondents were asked to report activities they had initiated and opportunities they had pursued. Encouragement from family, teachers, and peers was recorded separately. The intended contribution is to test whether entrepreneurial intent varies with students' exposure to active learning and multicultural experience, which bears on whether low entrepreneurial engagement in exam-centered systems reflects classroom design or cultural expectations. This question has

implications not only for Japan but also for other education systems whose institutional structures remain organized primarily around examination performance.

Research Design

This study used cross-sectional survey data to examine whether exposure to active learning environments and multicultural experiences is associated with entrepreneurial intent among high school students, with Japan as the primary case. The survey method was selected over interviews and classroom observation because it generates standardized responses from a geographically dispersed sample, measures several variables within a single instrument, and produces ordinal data suitable for group comparison. The instrument was constructed so that pedagogical exposure and social encouragement could be measured separately, since the study's central distinction is between constraints located in classroom practice and constraints located in the expectations of those around the student.

Participants were 71 respondents aged 15 to 19 or older (15: n=4; 16: n=20; 17: n=28; 18: n=14; 19 or older: n=5), drawn from the researcher's own school, other secondary schools in Japan, and the participant networks of Stanford e-Japan, Stanford e-Entrepreneurship Japan, the Reischauer Scholars Program, and LIFT, a youth nonprofit organization. By school location, 49 respondents reported a regional city, 17 a major city, and 5 a rural area. The survey did not ask for the country of residence. The survey was open to students who could read Japanese or English and had access to the online form. Current enrollment in secondary education was not verified: five respondents selected “19 or older,” and at least one reported currently attending university. No responses were excluded.

Respondents reported exposure to four curriculum types: the Japanese national curriculum (n=45), the International Baccalaureate (n=25), the American/AP curriculum (n=11), and other international curricula (n=4; two Korean, one British Columbia, one unspecified). These categories are not mutually exclusive: 11 respondents selected more than one curriculum, so the 71 respondents produced 85 curriculum selections in total, and each respondent appears in every group selected. Curriculum background served as the primary grouping variable, and out-of-school program participation was recorded separately as a secondary dimension of exposure.

The survey consisted of 25 items organized into five bilingual sections (Japanese and English), using multiple choice, checkboxes, five-point frequency scales, and one open-ended question.

Section one collected demographic information, including age, school location, and curriculum background. Section two measured the educational environment through two general items on classroom format and program participation, followed by five frequency items on specific active learning behaviors during a typical month: small-group discussion, defending opinions with evidence, designing or leading class activities, receiving peer feedback, and investigating before teacher explanation. These were written as frequency questions rather than agreement statements to reduce self-assessment bias. The five items operationalize the practices

that the Period for Inquiry-Based Cross-Disciplinary Study, introduced into the Japanese curriculum in 2022, was designed to promote; reported frequency of these practices was treated as a more stable indicator than participation in that subject, which schools implement with considerable variation.

Section three measured multicultural exposure through three items: exposure to people from different backgrounds (Q10), approach to learning a new topic (Q11), and international experiences (Q13). Q11 describes information-seeking behavior and may reflect curiosity rather than intercultural contact, so the multicultural composite is reported two ways: with all three items, and with Q10 and Q13 only. Section four measured entrepreneurial intent through four behavioral items: generating an original idea, initiating a project outside school, pursuing learning beyond the curriculum, and acting independently on a difficult problem. A fifth item asked whether family, teachers, or peers had encouraged the respondent to pursue entrepreneurship. This item describes what others do toward the student rather than the student's own behavior, and it is likely to vary with the type of school attended. It was therefore recorded as a separate measure of social encouragement and excluded from the entrepreneurial intent composite. This section prioritized past behavior over self-rated confidence, on the reasoning that self-assessed confidence is susceptible to social-desirability bias and that reported action is a more defensible indicator. Section five captured the perception of whether entrepreneurship had been presented as a viable path, whether interest changed after a global program, and one open-ended question.

Construct definitions and indicators were adapted from Kim, Huruta, and Lee, who applied the theory of planned behavior to South Korean high school students. Their educational environment, perceived social norm, and entrepreneurial intention constructs informed Sections two and four. The present instrument departs from that source by asking respondents about their actual behaviors and experiences rather than their perceived confidence.

The survey was administered through Google Forms and was designed to take three to four minutes. It remained open from 4 August to 18 August 2026. Participants were recruited through direct Instagram messages using the researcher's personal and program networks, with several recipients asked to circulate the link within their communities. No institution distributed the survey officially. Before beginning, participants were shown a bilingual disclaimer explaining the study's purpose, its supervision by a mentor at the University of Pennsylvania, and the confidentiality of responses. Participation was voluntary and no incentive was offered.

Responses were exported to Google Sheets and each item was assigned a numerical value on a one-to-five scale, with higher values indicating greater frequency, independence, or engagement. The three multicultural items (Q10, Q11, Q13) had three response levels coded 1, 3, and 5; Q13 was coded by the number of experience types selected (none = 1, one or two = 3, three or more = 5). Four composite scores were calculated as the arithmetic mean of their constituent items: active learning exposure (five items), multicultural exposure (three items), multicultural exposure without Q11 (two items), and entrepreneurial intent (four items, excluding encouragement). Encouragement was reported separately and entered into no composite. On that item only, the response "No, they discourage it" was coded 0 (n=3), so the item ranges from 0 to

5. Item means and standard deviations were computed for each curriculum group. Because curriculum groups overlap, all correlation and regression analyses used the 71 individual respondents, each counted once.

The analysis had three steps. First, Pearson correlations tested the association between each exposure composite and entrepreneurial intent. Second, Williams' t-test for dependent correlations compared the two correlations, because both share the same outcome. Third, a multiple regression entered active learning and multicultural exposure together to estimate each one's unique contribution. Spearman correlations were also computed as a check on treating ordinal items as interval data and showed the same pattern (ρ=.03 for active learning; ρ=.48 and .46 for the three- and two-item multicultural composites).

Results

This section presents the survey findings in two parts. The first examines the relationship between active learning environments and entrepreneurial intent. The second, addressed in section 3.2, examines multicultural experiences and entrepreneurial intent. All data are drawn from 71 responses collected between 4 August and 18 August 2026.

Active learning environments & entrepreneurial intent

Table 1. Educational systems and exposure to active learning

Curriculum	Small group discussion	Defending opinions with evidence	Student-led class activities	Investigation prior to teacher explanation	Peer feedback during class	Composite (5 items)
Japanese (n=45)	3.16±1.21	2.98±1.16	1.91±0.90	2.64±1.26	2.49±1.10	2.64±0.89
International Baccalaureate (IB), n=25	4.40±0.65	3.88±0.93	3.08±1.00	3.40±1.04	3.56±1.00	3.66±0.63
American Advanced Placement (AP), n=11	3.73±0.90	3.64±0.67	2.82±0.75	2.82±1.08	3.27±0.79	3.25±0.54
Other (n=4)	4.50±0.58	4.25±0.96	3.50±1.29	3.25±1.50	3.75±0.96	3.85±0.93

Note. Cells show M±SD on a five-point frequency scale. Composite is the mean of the five items. 11 respondents selected more than one curriculum and were counted in each group they selected. The Other group (n=4) is shown for completeness; estimates from so few respondents are imprecise and are not interpreted.

As shown in Table 1, the difference in active learning exposure between curricula is not equal across all five classroom behaviors. All four groups scored above the midpoint on small-group discussion, including Japanese-curriculum students (M=3.16, SD=1.21). However, Japanese curriculum students scored lowest on the behaviors in which students take a more active role. They reported a mean of only 1.91 (SD=0.90) for designing or leading class activities, compared with IB students (M=3.08, SD=1.00) and American/AP students (M=2.82, SD=0.75). Their scores for receiving peer feedback (M=2.49, SD=1.10) and investigating before teacher

explanation (M=2.64, SD=1.26) were also lower than their scores for small-group discussion (M=3.16) and defending opinions with evidence (M=2.98). Among the three larger groups, IB students scored highest on all five items (for example, M=4.40, SD=0.65 for small-group discussion).IB and AP students also reported more multicultural exposure than Japanese-curriculum students (Table 3), and individual active learning and multicultural exposure scores were positively correlated (r=.27, p=.025). The two exposures therefore overlap, which is why Section 3.2 tests them together.

Table 2. Educational systems and entrepreneurial engagement

Curriculum	Original idea generation	Initiative beyond school	Self-directed learning	Response to challenges	Intent composite (4 items)	Encouragement toward entrepreneurship (reported separately)
Japanese (n=45)	3.49±1.49	3.22±1.80	3.73±1.30	4.11±0.88	3.64±1.02	3.07±1.25
International Baccalaureate (IB), n=25	3.68±1.25	3.84±1.52	3.96±1.31	3.60±1.19	3.77±0.88	3.68±0.95
American Advanced Placement (AP), n=11	4.18±0.40	4.64±0.50	4.55±0.52	4.00±0.77	4.34±0.39	3.45±1.04
Other (n=4)	4.50±0.58	4.75±0.50	4.75±0.50	3.50±1.73	4.38±0.52	4.50±0.58

Note. Cells show M±SD. The intent composite is the mean of the first four items. Encouragement is reported separately and is not part of the composite. 11 respondents selected more than one curriculum and were counted in each group they selected. The Other group (n=4) is shown for completeness; estimates from so few respondents are imprecise and are not interpreted.

According to Table 2, Japanese curriculum students, who reported the least active learning in Table 1, scored lowest on three of the four behavioral items. The other groups did not follow the same order: IB students reported more active learning than American/AP students, but American/AP students scored higher on initiative outside school (M=4.64, SD=0.50 vs. M=3.84, SD=1.52) and self-directed learning (M=4.55, SD=0.52 vs. M=3.96, SD=1.31). Japanese curriculum students also reported the lowest encouragement to pursue entrepreneurship (M=3.07, SD=1.25), an item reported separately because it describes what others do toward the student. This suggests the barrier is less about individual ability and more about the expectations of those around them.

Two results stand out. First, Japanese-curriculum students scored highest on the challenge item (M=4.11, SD=0.88), higher than every other group. This means the Japanese system does develop persistence. What does not appear to develop is the willingness to start something independently without being told to. Second, American/AP students scored high on personal entrepreneurial behaviors but only moderately on encouragement (M=3.45, SD=1.04).

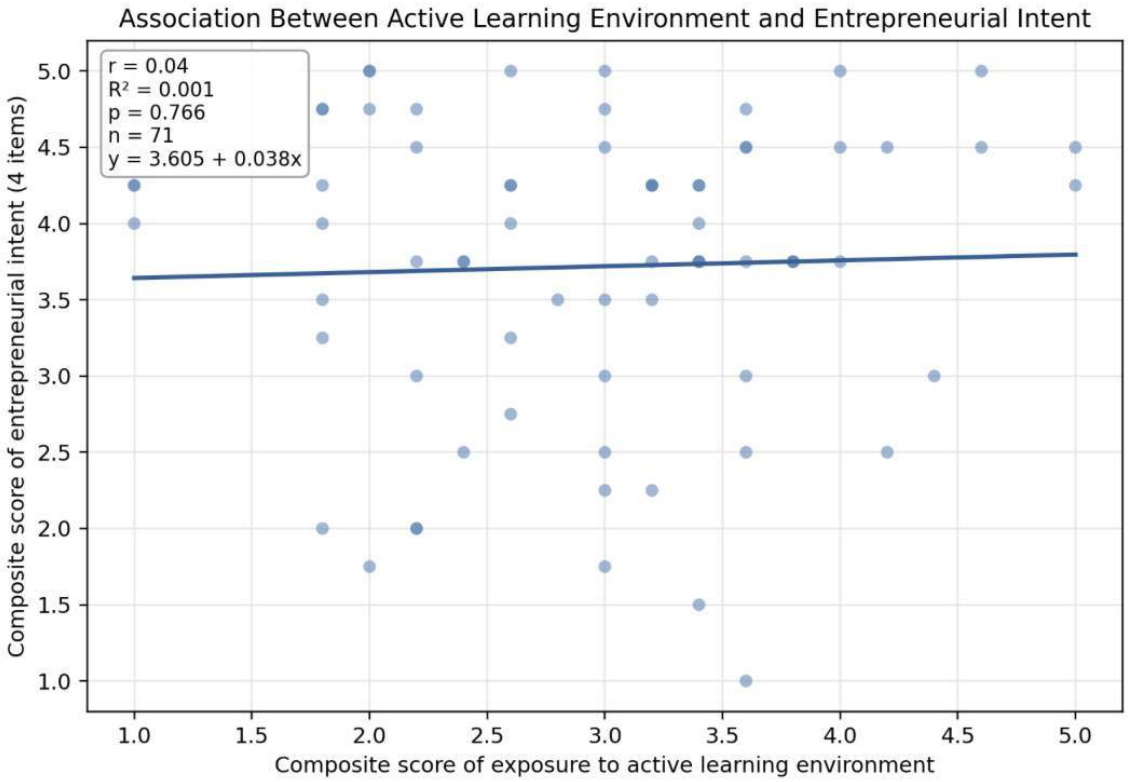

Figure 1. Association Between Active Learning Environment and Entrepreneurial Intent

No meaningful relationship was observed between exposure to an active learning environment and entrepreneurial intent, r=.04, p=.766, n=71. The X axis of Figure 1 represents the composite score of active learning exposure across the five indicators in Table 1, and the Y axis represents the composite score of entrepreneurial intent across the four behavioral indicators in Table 2 (encouragement excluded). The regression model indicated that a one-point increase in active learning score was associated with a 0.04-point increase in entrepreneurial intent, and the model explained 0.1% of the variance (R^2=.001). The fitted line is nearly flat, and students with similar active learning scores reported very different levels of entrepreneurial intent. Active learning exposure on its own, therefore, did not predict entrepreneurial intent in this sample. Section 3.2 examines whether multicultural exposure does.

3.2 Multicultural experiences & entrepreneurial intent

Table 3. Educational systems and multicultural experiences

Curriculum	Exposure to multicultural perspectives	Approach to new topic	International experiences	Composite (3 items)	Composite (2 items, without approach to new topic)
Japanese (n=45)	2.73±1.57	3.18±1.59	3.22±1.43	3.04±1.18	2.98±1.23

International Baccalaureate (IB), n=25	4.60±0.82	3.96±1.43	3.72±1.14	4.09±0.79	4.16±0.75
American Advanced Placement (AP), n=11	4.45±0.93	3.73±1.01	3.36±1.21	3.85±0.60	3.91±0.70
Other (n=4)	5.00±0.00*	3.50±1.91	4.50±1.00	4.33±0.54	4.75±0.50

Note. Cells show M±SD. Each of the three items had three response levels coded 1, 3, and 5. Composite (3 items) is the mean of all three items; Composite (2 items) excludes the approach to the new topic. 11 respondents selected more than one curriculum and were counted in each group they selected. The Other group (n=4) is shown for completeness; estimates from so few respondents are imprecise and are not interpreted.
(All four Other-group respondents chose the highest response, so SD=0.00 reflects a ceiling effect, and this cell cannot be interpreted.)

Table 3 shows that the clearest difference between curricula was on exposure to multicultural perspectives (Q10). Q10 had three response levels (minimal = 1, moderate = 3, high = 5). Japanese-curriculum students averaged M=2.73 (SD=1.57), just below "moderate." IB students (M=4.60, SD=0.82) and American/AP students (M=4.45, SD=0.93) averaged close to "high." All four students in the Other group chose "high" (M=5.00); with n=4 and no variation at the ceiling, this cell is shown but not interpreted. The Japanese-IB gap on Q10 (1.87 points) was larger than any Japanese-IB gap in Table 1 (largest: 1.24 points, small-group discussion). On the approach-to-new-topic question (Q11), the gap was smaller: Japanese-curriculum students scored 3.18 (SD=1.59) and IB students 3.96 (SD=1.43), a difference of 0.78 points compared with 1.87 points on Q10. This item describes how students look for information, which is closer to curiosity than to intercultural contact, so it is also analyzed in a composite without it. On the international-experience question (Q13), Japanese-curriculum students scored 3.22 (SD=1.43) and IB students 3.72 (SD=1.14).

Japanese curriculum students also reported the lowest encouragement in Table 2 (M=3.07, SD=1.25). At the individual level, however, encouragement was not clearly related to multicultural exposure ($r=.12$, $p=.314$), whereas it was positively related to active learning exposure ($r=.48$, $p<.001$) and to entrepreneurial intent ($r=.39$, $p<.001$). Encouragement, therefore, appears to go with classroom practice more than with intercultural contact.

At the group level, IB students combined high multicultural exposure (Table 3) with the highest encouragement among the three larger groups (M=3.68, SD=0.95), whereas American/AP students combined high multicultural exposure (M=4.45, SD=0.93) with only moderate encouragement (M=3.45, SD=1.04). Because the individual-level association was weak, these group patterns should not be read as an effect of multicultural exposure on encouragement.

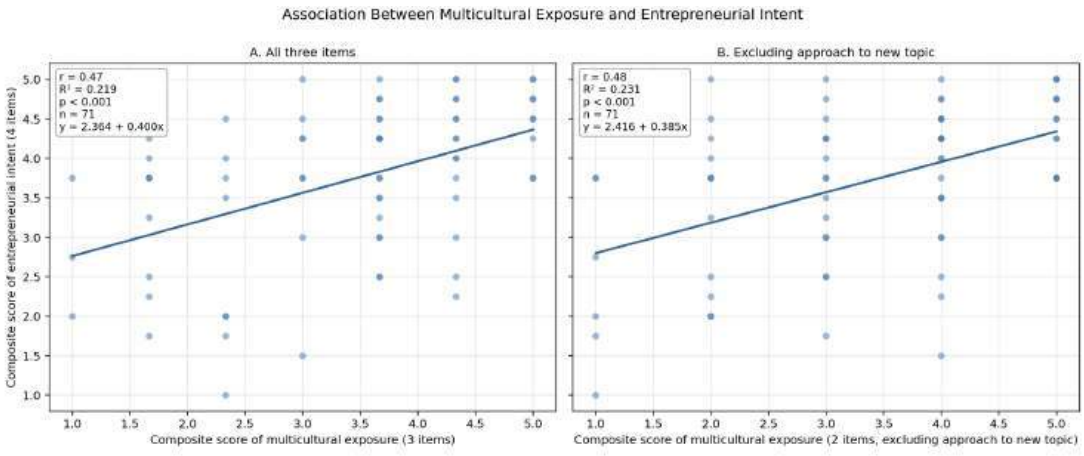

Figure 2. Association between multicultural exposure and entrepreneurial intent. (A) Three-item multicultural composite. (B) Two-item composite without an approach to a new topic.

A moderate positive relationship was observed between multicultural exposure and entrepreneurial intent, r=.47, p<.001, n=71 (Figure 2A). The X axis of Figure 2A represents the composite score of multicultural exposure across the three indicators in Table 3, and the Y axis represents the composite score of entrepreneurial intent across the four behavioral indicators in Table 2 (encouragement excluded). The regression model indicated that a one-point increase in multicultural exposure score was associated with a 0.40-point increase in entrepreneurial intent, and the model explained 21.9% of the variance (R^2=.219).

Because Q11 reflects curiosity more than intercultural contact, the composite was also calculated without it (Q10 and Q13 only; Figure 2B). The result was almost identical, r=.48, p<.001, R^2=.231, so the association does not depend on that item. Q11 alone correlated with intent at r=.28 (p=.017), weaker than Q10 (r=.42) and Q13 (r=.37).

Table 4. Regression of entrepreneurial intent on active learning and multicultural exposure (N=71)

Predictor	b	SE	β	p
Model 1 (three-item multicultural composite)				
Active learning	−0.10	0.12	−.10	.392
Multicultural exposure	0.42	0.09	.49	<.001
R^2 = .227, F(2, 68) = 10.01, p < .001				
Model 2 (two-item multicultural composite)				
Active learning	−0.07	0.12	−.06	.570
Multicultural exposure	0.40	0.09	.49	<.001
R^2 = .234, F(2, 68) = 10.42, p < .001				

Note. Entrepreneurial intent is the four-item composite (encouragement excluded). b = unstandardized coefficient; β = standardized coefficient.

To test whether the two relationships differ, and to estimate each exposure's unique contribution, both exposures were entered into one regression (Table 4). Multicultural exposure remained a significant predictor of entrepreneurial intent (β=.49, p<.001), while active learning did not (β=−.10, p=.392). The two predictors together explained 22.7% of the variance, and multicultural exposure accounted for almost all of it (unique R^2=.226, versus .008 for active learning). The correlation for multicultural exposure (r=.47) was significantly larger than for active learning (r=.04), t(68)=3.32, p<.01. The two-item composite gave the same result.

Discussion

In this study, students in IB, AP, and other international curricula generally reported greater exposure to active learning practices and to multicultural contact than students in the Japanese curriculum. Entrepreneurial intent was lower in the Japanese curriculum than in the American/AP curriculum (M=3.64 vs. 4.34), but similar to the IB curriculum (M=3.77). However, active learning exposure showed no association with entrepreneurial intent (r = .04, p = .766). Interestingly, Japanese-curriculum students reported the highest score for responding to challenges, suggesting that their lower levels of entrepreneurial engagement were not simply attributable to a lack of persistence or problem-solving capacity. A study of Korean high school students by Kim et al. reported a related pattern: the educational environment did not significantly mediate the effects of perceived social norm, attitude, or perceived entrepreneurial capacity on entrepreneurial intention, while perceived social norm was directly associated with intention (p < .10). In the present study, Japanese-curriculum students reported the lowest level of encouragement from people around them (M = 3.07), and encouragement was positively related to both active learning exposure (r = .48) and entrepreneurial intent (r = .39). Together, these findings suggest that the social environment around students may matter alongside classroom pedagogy in shaping entrepreneurial intent.

A substantially stronger pattern emerged for multicultural exposure. Japanese-curriculum students reported markedly lower exposure to people from different backgrounds than students in IB, American/AP, and other curricula. Multicultural exposure was moderately and positively associated with entrepreneurial intent (r = .47, p < .001) and explained 21.9% of its variance. When both exposures entered one regression, multicultural exposure remained a significant predictor (β = .49, p < .001) and active learning did not (β = −.10, p = .392). Overall, the findings suggest that exposure to diverse perspectives and international experiences is more strongly associated with entrepreneurial intent than active-learning pedagogy. The two exposures were themselves correlated (r = .27), so the design cannot rule out that multicultural exposure stands in for other differences between schools, such as family background.

These findings carry implications for schools and curriculum designers. The data suggest that introducing active learning practices alone may be insufficient to account for differences in entrepreneurial intent. Japanese-curriculum students reported some exposure to small-group

discussion but very little to student-led class design, and even among students with relatively high active learning scores, entrepreneurial intent varied widely. By contrast, students with higher multicultural exposure tended to report higher entrepreneurial intent. This suggests that curriculum reform efforts focused solely on pedagogical methods, such as the Period for Inquiry-Based Cross-Disciplinary Study introduced in 2022, may have limited effects on entrepreneurial orientation if they are not accompanied by opportunities for students to interact with peers from different cultural and linguistic backgrounds. Programs such as Stanford e-Japan (Stanford Program on International and Cross-Cultural Education) illustrate one possible model for providing such cross-cultural exposure.

For educators, the results point toward specific classroom practices that may support entrepreneurial thinking. Japanese curriculum students scored lowest on the items that require student agency: designing class activities, receiving peer feedback, and investigating before teacher explanation. These are practices that teachers can introduce without restructuring an entire curriculum. At the same time, the multicultural findings suggest that exposing students to diverse perspectives, whether through international guest speakers, cross-cultural project partnerships, or multilingual learning materials, may be as important as changing the format of instruction. Teachers who combine active learning methods with opportunities for students to encounter unfamiliar viewpoints may be addressing both dimensions that this study found relevant to entrepreneurial orientation.

For students, the findings suggest that entrepreneurial intent is shaped not only by what happens inside the classroom but by the range of perspectives students encounter outside it. Japanese-curriculum students who reported high multicultural exposure also tended to report higher entrepreneurial intent (among the 38 respondents who selected only the Japanese curriculum, $r=.49$, $p=.002$), even when their classroom environment remained lecture-based. This suggests that extracurricular and informal multicultural experiences may represent one possible avenue through which students encounter diverse perspectives. The data also show that Japanese-curriculum students scored lower than IB students on the approach-to-new-topic item ($M=3.18$ vs. 3.96), but by a smaller margin than on exposure to multicultural perspectives ($M=2.73$ vs. 4.60), which suggests that the desire to explore is less limited than the opportunity for intercultural contact. The lower encouragement scores may indicate that some students have less perceived social support for acting on this curiosity. Students who actively seek diverse environments may be creating for themselves the conditions that their formal education does not yet provide.

Conclusion

The study produced one central finding. Multicultural exposure correlated with entrepreneurial intent ($r=.47$, $p<.001$), whereas active learning exposure did not ($r=.04$, $p=.766$), and in a regression containing both, only multicultural exposure remained a significant predictor ($\beta=.49$, $p<.001$). Because the two conditions often appear together in the same curricula, they function as combined inputs in most students' lives. The present data suggest, however, that they

do not contribute equally. Multicultural exposure may matter for entrepreneurial intent, although this survey cannot show the mechanism. In exam-centered systems like Japan, where both exposures are structurally limited, the multicultural dimension may be the more pressing gap to address.

These findings should be interpreted in light of several limitations. The sample of 71 students is too small to support strong statistical generalization, particularly for the smaller groups (American/AP n=11, Other n=4). Recruitment relied on the researcher's personal and program networks, which skews toward students already engaged with international education; a random school-based sample would likely produce different baseline scores. All data are self-reported, introducing social desirability bias and recall error. Composite reliability was modest (Cronbach's α=.67 for entrepreneurial intent, .64 for multicultural exposure). The survey also did not record the country of residence or verify current enrollment in secondary education, so the assumption that respondents share one national context could not be checked. Finally, the curriculum item allowed multiple selections, producing 85 curriculum selections from 71 respondents, which prevents a clean separation between groups and limits the inferences that can be drawn from between-group comparisons.

Future research could address these limitations in several ways. A larger sample drawn from schools rather than program networks would better capture how students outside the international-education pipeline respond to the same measures. Comparative research in other exam-centered systems would test whether the Japanese pattern reflects something specific to Japan or something common to this education model. A long-term design measuring entrepreneurial intent before and after multicultural intervention, such as participation in a program like Stanford e-Japan, would provide stronger evidence for causation rather than correlation. Qualitative interviews could complement the quantitative data by examining what, specifically, students take from multicultural exposure that translates into entrepreneurial orientation.

Works Cited

Kim, Min-Sun, Andrian Dolfriandra Huruta, and Cheng-Wen Lee. “Predictors of Entrepreneurial Intention among High School Students in South Korea.” Sustainability, vol. 14, no. 21, 2022, article 14168. https://doi.org/10.3390/su142114168.

Komatsu, Hikaru, Jeremy Rappleye, and Iveta Silova. “Student-Centered Learning and Sustainability: Solution or Problem?” Comparative Education Review, vol. 65, no. 1, 2021. https://doi.org/10.1086/711829.

Kurata, Karin, et al. “Entrepreneurial Intention Among Engineering Students: The Moderating Role of Entrepreneurship Education in Japan.” Behavioral Sciences, vol. 15, no. 5, 2025, article 663. https://doi.org/10.3390/bs15050663.

Ministry of Education, Culture, Sports, Science and Technology (MEXT). Courses of Study. 2018, https://jp-cos.github.io/en/84M/0000000000000.

Mori, Akiko. “Facilitating Collaboration between Japanese High Schools and Universities: A Qualitative Exploration of the Role of Education Outreach Coordinators.” Frontiers in Education, vol. 9, 2024, article 1393183. https://doi.org/10.3389/feduc.2024.1393183.

Moriano, Juan A., et al. “A Cross-Cultural Approach to Understanding Entrepreneurial Intention.” Journal of Career Development, vol. 39, no. 2, 2012, pp. 162–85. https://doi.org/10.1177/0894845310384481.

Mutakinati, L., I. Anwari, and Y. Kumano. “Analysis of Students’ Critical Thinking Skill of Middle School through STEM Education Project-Based Learning.” Jurnal Pendidikan IPA Indonesia, vol. 7, no. 1, 2018. https://doi.org/10.15294/jpii.v7i1.10495.

Sakurai, Riho. “Impacts of Recent Education Reforms in Japan: Voices from Junior High Schools in Japan.” Journal of International Cooperation in Education, vol. 18, no. 2, 2016, pp. 55–65. https://www.academia.edu/63760720/Impacts_of_Recent_Education_Reforms_in_Japan_Voices_from_Junior_High_Schools_in_Japan

Shimahara, Nobuo K., and Akira Sakai. Learning to Teach in Two Cultures: Japan and the United States. Routledge, 2018.

Stanford Program on International and Cross-Cultural Education. “Stanford e-Japan.” Stanford University, https://spice.fsi.stanford.edu/fellowship/stanford-e-japan. Accessed 8 Oct. 2026.

Yokoyama, Kazuko, and Sarah Louisa Birchley. “Entrepreneurship in the Japanese Context.” Transnational Entrepreneurship in South East Asia, Springer, 2020, pp. 17–39. https://doi.org/10.1007/978-981-32-9252-9_2.

Education-Labor Mismatch And Brain Drain: A Political Economy Analysis Of Youth Exclusion In Jordan And Lebanon By Aarohi Bhattacherjee

Abstract

This paper explores the structural drivers of unemployment among educated youth and skilled emigration in Jordan and Lebanon. Beyond traditional supply-side "skills mismatch" explanations, this study applies a critical political economy framework to evaluate how macroeconomic structures, financialization, and state policies condition labor demand. Through a qualitative comparative literature synthesis of 24 core institutional and academic sources, the paper contrasts Jordan's chronic economic stagnation and public-sector reliance against Lebanon's post-2019 financial collapse and banking insolvency. The findings show that expanding higher education without complementary industrial job creation converts educated labor into an export commodity. While Jordan experiences a steady, career-driven outflow of professionals to Gulf Cooperation Council (GCC) states, Lebanon exhibits acute, widespread global emigration. Remediating youth exclusion requires structural macroeconomic reforms, industrial policy, and financial stabilization rather than exclusive reliance on curricular adjustments.

Keywords Political Economy, Youth Unemployment, Jordan, Lebanon, Brain Drain, Labor Economics.

1. Introduction

Countries across the Middle East and North Africa (MENA) consistently exhibit some of the highest youth unemployment rates globally. This structural reality presents a persistent paradox: tertiary education enrollment continues to expand rapidly across the region, yet formal employment opportunities remain severely constrained. Rather than facilitating upward social mobility, higher education frequently leads to prolonged labor market queueing—where graduates remain voluntarily unemployed while waiting for secure public sector jobs—underemployment, or high-skilled emigration (World Bank, "*MENA Economic Update"* 14).

Mainstream development policy often frames this dynamic as a supply-side "skills mismatch," asserting that educational institutions fail to impart market-relevant competencies (Fakih and Marrouch 1440). However, focusing exclusively on human capital deficits overlooks the broader macroeconomic conditions that govern firm formation and labor demand.

This study addresses the following research question: *How do the political-economic structures of Jordan and Lebanon generate education-labor mismatch and skilled youth emigration, and why do these processes assume distinct forms in each context?*

By comparing Jordan's chronic, low-growth economic stagnation with Lebanon's post-2019 financial implosion, this paper evaluates how national accumulation regimes—the macro-level framework through which capital is generated, invested, and circulated within an economy—condition patterns of youth marginalization. Ultimately, this paper contends that while

supply-side mismatches may exist in specific technical sectors, they cannot account for graduate exclusion without addressing structural constraints on domestic labor demand.

2. Theoretical Framework and Literature Review

Human Capital Theory posits that investments in education increase individual productivity, yielding higher wages and lower unemployment risks. When applied to developing or rentier economies, however, this linear relationship frequently falters. Under Signalling and Credentialism theories, in contexts where productive labor demand is constrained, university degrees function primarily as sorting mechanisms for congested applicant pools rather than direct indicators of workplace capability.

While empirical studies on MENA labor markets highlight real deficits in technical competencies, foreign language proficiency, and specialized vocational training (Economic Research Forum, "*Youth Unemployment"* 12), attributing graduate unemployment solely to curricular deficiencies presents major analytical limitations. Supply-side frameworks overlook why private enterprise fails to generate high-value positions capable of absorbing tertiary graduates.

To rigorously evaluate the supply-side counterfactual, this synthesis directly engages with empirical studies advocating the skills mismatch thesis. Fakih and Marrouch present quantitative evidence from school-to-work transition surveys in Lebanon, demonstrating significant horizontal and vertical skills mismatches among university graduates, alongside perceived deficits in employer-demanded soft and technical skills (1440–1448). Their findings confirm that curricular misalignments and informational asymmetries do increase initial job-search durations for young degree holders.

However, evaluating the relative explanatory power of both frameworks reveals that supply-side factors alone cannot account for the macro-level volume of youth exclusion. If curricular mismatch were the primary driver of graduate unemployment, targeted vocational training, firm-level apprenticeships, and university curriculum revisions would yield proportional expansions in high-skilled formal employment. Yet empirical evidence across both Jordan and Lebanon demonstrates that even in sectors where graduates possess high technical proficiency—such as healthcare, pharmacy, civil engineering, and software development—the domestic economy fails to absorb them due to macroeconomic structural ceilings (World Bank, "Promoting Youth Employment" 22; World Bank, "Lebanon Economic Monitor 2021" 11). Thus, while supply-side frictions function as secondary exacerbating mechanisms that dictate *who* gets hired first, macroeconomic capital allocation determines the total *number* of productive positions available.

To address these analytical gaps, this study integrates Push-Pull Migration Models with regional political economy frameworks. Structural push factors—such as low real wages, limited career progression, clientelism, and macroeconomic volatility—drive skilled labor outward. Neoliberal restructuring and deindustrialization have further exacerbated these dynamics. As Adam Hanieh (115-125) and Gilbert Achcar (78-85) demonstrate, privatization and trade

liberalization destabilized domestic manufacturing across the region, redirecting capital toward low-productivity services, real estate, and financial speculation.

Furthermore, Ali Kadri (62-70) highlights how the absence of productive domestic investment creates a dual labor market: exporting educated citizens while importing low-wage foreign labor. State-business networks and welfare retrenchment solidify these structural bottlenecks. Pete W. Moore (102-118) documents how state-crony networks in Jordan protected incumbent firms, suppressing competitive job creation. Similarly, Anne Marie Baylouny (42-52) demonstrates how state welfare retrenchment forced households onto informal coping strategies, while Hannes Baumann (85-98) shows how Lebanon's post-war reconstruction relied on real estate financialization and sovereign debt accumulation, prioritizing elite rent extraction over productive employment generation.

3. Methodology

This study utilizes a qualitative comparative literature synthesis to evaluate the structural drivers of youth exclusion in Jordan and Lebanon. Rather than conducting an empirical primary data collection, the paper synthesizes evidence across 24 core peer-reviewed academic monographs, journal articles, and multilateral policy publications published between 2004 and 2025.

Document Selection & Corpus Assembly

To construct the 24-document research corpus, academic databases (Google Scholar, JSTOR) and policy repositories (World Bank Open Knowledge Repository, International Labour Organization Knowledge Gateway, and the Economic Research Forum) were searched using explicit thematic keywords: "youth unemployment Jordan", "Lebanon financial crisis labor demand", "MENA political economy", and "skills mismatch Arab world".

Initial search queries yielded 68 potential documents. These were screened through a three-stage inclusion protocol:

1. Thematic Alignment: Selection of sources explicitly examining youth labor dynamics, macroeconomic structure, or skilled migration in the MENA region.
2. Institutional Authority & Rigor: Restricting non-academic policy literature to peer-reviewed studies or publications from multilateral developmental entities (e.g., World Bank, ILO, Economic Research Forum).
3. Theoretical Relevance: Prioritizing studies grounded in political economy, accumulation regimes, or structural labor demand analysis.

Following this protocol, 24 core documents were finalized for detailed qualitative content analysis (Hanieh; Achcar; Kadri; Moore; Baylouny; Baumann; Assaad and Krafft; Diwan and Haidar; Economic Research Forum, "Youth Unemployment"; Economic Research Forum, "Labor Market Dynamics"; ILO, "Transition to the Labor Market"; ILO, "Lebanon Follow-up"; ILO, "Global Employment Trends"; Kasparian; Tzannatos; World Bank, "Promoting Youth Employment"; World Bank, "Jordan Economic Monitor 2020"; World Bank, "Lebanon Economic

Monitor 2023"; World Bank, "Jordan Economic Monitor 2024"; World Bank, "MENA Economic Update"; ILO, "Labour Market Supply"; RAND Corporation; Fakih and Marrouch; World Bank, "Lebanon Economic Monitor 2021").

4. Case Study 1: Jordan – Stagnation and Public Sector Queueing

Jordan's labor market dynamics are fundamentally shaped by its macroeconomic growth model, which is characterized by chronic low GDP growth, fiscal austerity, and heavy reliance on foreign aid and worker remittances (World Bank, "*Jordan Economic Monitor 2024"* 8). Capital inflows into the Jordanian economy remain heavily concentrated in real estate, retail, and low-value service sectors. Consequently, high-tech, high-value-added services, and exportable manufacturing industries—the primary sectors capable of absorbing university graduates—remain structurally underdeveloped (Economic Research Forum, "*Labor Market Dynamics"* 15).

This structural imbalance directly dictates the pattern of sectoral labor demand. Historically, public sector employment served as a cornerstone of Jordan's state-citizen social contract. Despite decades of fiscal contraction and state hiring freezes, queuing for public sector employment remains a rational economic choice for many graduates. The domestic formal private sector is overwhelmingly dominated by micro- and small enterprises that offer low starting wages, negligible career progression, and minimal job security (Assaad and Krafft 112).

Gendered labor market exclusion is particularly pronounced in Jordan, where female labor force participation remains among the lowest in the world at approximately 14 percent (ILO, "*Labour Market Supply"* 22). For educated female graduates, waiting for civil service appointments is an intensely rational decision. The public sector provides mandatory maternity benefits, predictable working hours, formal wage parity, and social legitimacy that are largely absent in unregulated or informal private sector workplaces (RAND Corporation 45).

Stagnant real private sector wages create a severe mismatch between graduate earnings expectations and market realities, resulting in widespread occupational underemployment. Faced with limited upward mobility at home, educated Jordanians turn to emigration as a structural outlet. Jordan exhibits a structured, career-driven migration flow directed primarily toward the Gulf Cooperation Council (GCC) economies (Tzannatos 238). While the resulting remittance inflows sustain household consumption and domestic demand, this continuous exit drains public institutions and domestic industries of top-tier technical expertise, reinforcing state reliance on external aid and foreign capital.

5. Case Study 2: Lebanon – Rentier Collapse and Systemic Exodus

Prior to its catastrophic financial crash in 2019, Lebanon operated under a highly financialized rentier growth model. This system was sustained by high interest rates designed to attract commercial bank deposits, deep sovereign debt accumulation, and intense real estate speculation (Baumann 134; World Bank, "*Lebanon Economic Monitor 2021"* 5). Capital allocation systematically favored non-productive financial services, luxury real estate, and

tourism, creating a severe structural deficit in high-skilled engineering, agricultural, and industrial labor demand long before the collapse (Diwan and Haidar 12; Fakih and Marrouch 1445).

The post-2019 financial implosion, hyperinflation, and complete banking insolvency destroyed Lebanon's formal employment landscape. According to the official Lebanon Follow-up Labor Force Survey conducted by the Central Administration of Statistics (CAS) and the International Labour Organization, the youth unemployment rate surged to **47.8 percent** in January 2022 (ILO, "Lebanon Follow-up" 12; CAS 2022). Unlike Jordan, where public sector queueing remains a persistent structural feature, Lebanon's sovereign default and public sector wage collapse eliminated state employment as a viable safety net or queueing destination (ILO, "*Lebanon Follow-up"* 18).

The crisis has severely compounded gendered exclusion. Hyperinflation and the collapse of public infrastructure—including public transit and childcare services—disproportionately elevated unpaid domestic care burdens, forcing female university graduates out of the formal workforce or accelerating their push toward international migration (ILO, "*Transition to the Labor Market"* 34).

Consequently, Lebanon's migration dynamics transformed overnight. What was historically a managed stream of professional circular migration shifted post-2019 into an emergency, unrestricted exodus (Kasparian 56). High rates of emigration among medical doctors, university professors, and specialized engineers reflect an acute institutional breakdown. Unlike Jordan's GCC-bound temporary employment flows, Lebanon's post-crisis migration represents a multi-continental, permanent flight of human capital toward Europe, North America, and Australia.

6. Comparative Synthesis

A side-by-side analysis of Jordan and Lebanon reveals that while both nations suffer from high graduate exclusion and brain drain, the underlying macroeconomic mechanisms differ significantly.

Comparative Dimension	**Jordan**	**Lebanon**
Primary Economic Driver	Chronic stagnation, fiscal austerity, rentier aid dependency	Banking collapse, sovereign default, hyperinflation
Labor Market Mechanics	Rational public sector queueing; weak private sector absorption	Acute private sector contraction; complete state insolvency
Education-Demand Fit	Credentials function as sorting mechanisms for scarce roles	Severe degradation of white-collar professional labor market
Gender Dynamics	Low participation (~14%); public sector preference for structural benefits	Forced displacement due to care burdens and collapsing real wages
Emigration Pattern	Structured, career-driven outflow to GCC economies	Emergency, multi-continental high-skilled flight

To evaluate these dynamics with precision, labor market exclusion must be disaggregated into four distinct operational outcomes:

1. **Unemployment:** The complete absence of formal work, prevalent among Jordanian graduates queuing for scarce civil service positions.
2. **Underemployment:** Working part-time or in low-productivity informal roles below one's capacity, widespread across both countries.
3. **Occupational Mismatch:** Employment in roles completely unrelated to degree qualifications due to weak private sector high-skilled demand.
4. **Emigration:** Skilled outward mobility chosen as a structural response to domestic labor market foreclosure.

7. Conclusion and Strategic Policy Frameworks

Comparing Jordan and Lebanon demonstrates that graduate unemployment and brain drain are not merely supply-side educational failures; they are the direct structural outputs of distinct macroeconomic regimes. In both cases, expanding tertiary education without active industrial job creation converts educated human capital into an export commodity. Remediating youth exclusion requires moving beyond simple curricular revisions toward structural macroeconomic and industrial policy interventions:

- ****For Jordan:**** Policy must focus on active industrial strategy targeting high-value sectors such as renewable energy, software engineering, and pharmaceuticals. Simultaneously, enforcing strict labor standards, mandatory social security, and maternity protections in private enterprise is vital to close the job-quality gap with the civil service.
- ****For Lebanon:**** The immediate priority must be macro-financial restructuring, resolving commercial banking insolvency, and stabilizing the currency. Restoring baseline productive capacity requires directing capital into public infrastructure, power grids, and higher education institutions.
- ****Regional Initiatives:**** MENA economies should establish formal diaspora engagement frameworks to leverage remote knowledge networks, turning brain drain into virtual "brain gain."

Limitations and Future Research

This study is limited by the scarcity of disaggregated post-2023 labor market data in Lebanon following state institutional collapse. Future empirical research should utilize longitudinal household surveys to track graduate career trajectories across socioeconomic strata and assess the long-term impact of remote diaspora work on local capital formation.

Works Cited

Achcar, Gilbert. The People Want: A Radical Exploration of the Arab Uprising. University of California Press, 2013.

Assaad, Ragui, and Caroline Krafft. The Jordanian Labor Market in the Wake of the Arab Spring. Oxford University Press, 2015.

Baumann, Hannes. Citizen Hariri: Lebanon's Neoliberal Reconstruction. Hurst, 2016.

Baylouny, Anne Marie. Privatizing Welfare in the Middle East: Kin Mutual Aid Associations in Jordan and Lebanon. Indiana University Press, 2010.

Diwan, Ishac, and Jamal Ibrahim Haidar. "Political Connections Reduce Job Creation in Lebanon." Journal of Development Studies, vol. 57, no. 8, 2021, pp. 1373–1396.

Economic Research Forum. Labor Market Dynamics and Structural Transformation in Jordan. ERF Research Report No. 441, Economic Research Forum, 2022.

Youth Unemployment and the Transition from School to Work in the MENA Region. ERF Working Paper Series No. 1320, Economic Research Forum, 2019.

Fakih, Ali, and Walid Marrouch. "Employment and Skill Mismatch Among Youth in Lebanon." International Journal of Manpower, vol. 40, no. 8, 2019, pp. 1438–1457.

Hanieh, Adam. Lineages of Revolt: Issues of Contemporary Capitalism in the Middle East. Haymarket Books, 2013.

International Labour Organization. Global Employment Trends for Youth: Arab States Regional Brief. ILO, 2023.

Labour Market Supply and Demand in Jordan: Structural Imbalances and Policy Challenges. ILO Regional Office for Arab States, 2020.

Lebanon Follow-Up Labor Force Survey: Rapid Assessment of Employment Shifts. ILO and Central Administration of Statistics, 2022.

Transition to the Labor Market: Youth School-to-Work Transition Survey in Lebanon. ILO Regional Office for Arab States, 2017.

Kadri, Ali. Arab Development Denied: Dynamics of Accumulation by Wars of Encroachment. Anthem Press, 2014.

Kasparian, Choghig. The Circular Migration of Lebanese Graduates: Flows and Reintegration. European University Institute, 2014.

Moore, Pete W. Doing Business in the Middle East: Politics and Economic Crisis in Jordan and Kuwait. Cambridge University Press, 2004.

RAND Corporation. Youth in Jordan: Transitions from Education to Employment. RAND Research Report Series, RAND Corporation, 2014.

Tzannatos, Zafiris. "Labor Market Performance and Policies in the GCC and Wider MENA Region." Middle East Development Journal, vol. 8, no. 2, 2016, pp. 232–253.

World Bank. Jordan Economic Monitor: Navigating Turbulent Times. World Bank Group, 2024.

Jordan Economic Monitor: Weathering the Storm. World Bank Group, 2020.

Lebanon Economic Monitor: Lebanon Sinking (To the Top 3). World Bank Group, 2021.

Lebanon Economic Monitor: The Normalization of a Crisis Economy. World Bank Group, 2023.

MENA Economic Update: Growth, Labor Absorption, and Structural Reform. World Bank Group, 2025.

Promoting Youth Employment in Jordan: Structural Constraints and Options. World Bank Group, 2018.

Evaluating ASPM as a Genetic Modifier of Congenital Zika-Associated Microcephaly

By Aanya Sood

Abstract

Prenatal Zika virus (ZIKV) infection does not affect every fetus the same; only some exposed fetuses develop Zika syndrome, including severe microcephaly, suggesting that genetic differences may influence susceptibility. The purpose of this review is to specifically examine whether variants in *ASPM*, the gene most commonly associated with inherited primary microcephaly, could alter how prenatal ZIKV infection affects the developing brain. *ASPM* helps neural progenitor cells divide correctly and produce the cells necessary for brain growth. Any harmful variants in the *ASPM* gene can lead to autosomal recessive primary microcephaly. Overall, the findings indicate that ZIKV disrupts the growth and survival of developing brain cells and alters pathways associated with microcephaly. Past studies have shown harmful health effects of ZIKV. However, no study has directly compared *ASPM* variants in fetuses exposed to ZIKV with microcephaly versus those without. Although not confirmed, current evidence supports a possible gene-environment interaction in which mild or regulatory *ASPM* variants may make neural progenitor cells more susceptible to damage from ZIKV infection. Identifying this gap in research could contribute to significant progress in this field by guiding future genetic sequencing, explaining why fetal outcomes differ after ZIKV exposure, and helping researchers distinguish inherited microcephaly from Zika-associated microcephaly.

Keywords

ASPM; Zika virus; congenital Zika syndrome; microcephaly; p53 stress pathway; genetic susceptibility; genetic sequencing; neural progenitor cells

Introduction

Prenatal exposure to the same virus can produce very different outcomes. While some fetuses exposed to Zika virus (ZIKV) develop congenital Zika syndrome, including severe microcephaly, others are born without major congenital abnormalities. This variation suggests that ZIKV exposure alone may not determine whether microcephaly develops and that fetal genetics could influence susceptibility (Marques et al., 2025).

ZIKV can cross the placenta and infect cells that are essential for fetal brain development. ZIKV targets radial glial cells, which are neural progenitor cells responsible for producing many of the neurons in the cerebral cortex. Neural progenitor cells are crucial for brain development. Consequently, the infection reduces the neural progenitor population, preventing cortical development, and decreases the expression of several genes associated with microcephaly (Wu et al., 2016). The gene most commonly associated with microcephaly is *ASPM*, which helps neural progenitor cells divide symmetrically, allowing sufficient cells to be produced for normal growth. Even without ZIKV exposure, harmful *ASPM* variants can also lead to primary microcephaly. Thus, inherited *ASPM*-related microcephaly and Zika-associated microcephaly have different causes, but both can interfere with the growth and division of neural progenitor cells.

Various human studies depict that genetic background may influence the effects of prenatal ZIKV infection. A separate study by Caires-Júnior et al. (2018) studied three pairs of dizygotic, or fraternal, twins in which only one twin developed congenital Zika syndrome. Neural progenitor cells from the affected twins showed greater viral replication and more severe reductions in cell growth than cells from their unaffected siblings. Because the twins shared the same pregnancy but differed genetically, the results suggest that genetic differences may influence cellular susceptibility to ZIKV. Another case-control study compared 40 children with congenital Zika syndrome with 48 ZIKV-exposed children without congenital abnormalities. Variants in the immune-related genes *NOS2* and *TNF* were associated with the risk or severity of congenital Zika outcomes (Gomes et al., 2021). Although these studies did not investigate *ASPM*, they demonstrate that host genetic variants may contribute to differences in fetal outcomes.

Despite this evidence, no study has directly compared *ASPM* variants in ZIKV-exposed fetuses with microcephaly against those in similarly exposed fetuses without congenital abnormalities. A recent scoping review identified several possible maternal and fetal genetic modifiers of prenatal ZIKV-related brain injury but did not establish *ASPM* directly as the susceptible gene (Marques et al., 2025). This matters because ZIKV infection and reduced *ASPM* function affect similar cells and developmental processes, yet their possible interaction remains unclear. Identifying a genetic modifier could help explain why fetal outcomes differ after ZIKV exposure, distinguish inherited microcephaly from Zika-associated microcephaly, and guide future genetic studies. However, current evidence is insufficient to support the use of *ASPM* variants for clinical risk prediction.

This paper examines whether variants in the *ASPM* gene contribute to the reason why only some fetuses exposed to ZIKV develop microcephaly, while others don't. It argues that *ASPM* is a biologically plausible genetic modifier because *ASPM* dysfunction and ZIKV infection both disrupt neural progenitor-cell division and brain growth. However, the evidence currently available is indirect and does not prove that *ASPM* variants increase susceptibility to congenital Zika-associated microcephaly.

The review compares evidence from human genetic studies, neural-cell experiments, cerebral organoids, and animal models. It first examines how *ASPM* supports neural progenitor-cell division and how ZIKV disrupts similar developmental processes. It then evaluates evidence that genetic background influences how ZIKV affects the fetus. It considers the indirect link between ZIKV and reduced *ASPM* expression and compares this proposed interaction with alternative factors, such as infection timing and placental damage, that could influence susceptibility.

Body

I - Normal *ASPM* Function and Consequences of Disruption

The *ASPM* gene is critical for cell division, especially in the developing brain. The cells that divide are known as neural progenitor cells. These cells divide a limited number of times

before developing into neurons and other brain cells. These cells eventually form the cerebral cortex, the brain's outer layer, which is involved in functions such as movement, memory, language, and reasoning (i.e., higher-level thinking). Any decrease in the number of successful progenitor cell divisions can be significant and lead to apoptosis (programmed cell death). To examine *ASPM's* role in this process, a specific study done by Higgins et al. (2010) reduced *ASPM* activity in cultured human cells. The paper originally established that this specific gene is required for normal spindle orientation and for the efficient completion of cytokinesis during cell division. The spindles pull chromosomes apart, and cytokinesis is the final step in which the original cell physically divides into two daughter cells. One cannot occur without the other, which can disrupt cell division. When *ASPM* activity is reduced, the researchers observed abnormalities in spindle organization, orientation, and cytokinesis. The reduction led to frequent failures to complete these processes properly. Their experiments also indicated that the C-terminus — the final section of the *ASPM* protein — is required for *ASPM* to reach the correct location and perform the proper functions. This finding helps explain why mutations that produce shortened *ASPM* proteins may be harmful. While the protein may still be produced, it may not reach the spindle. However, the study used general cultured human cells rather than neural progenitor cells inside the developing brain. It therefore provides direct evidence of *ASPM's* cellular function, but cannot show how those defects affect brain size. To expand on this finding, a 2016 study found that *ASPM* interacts with CITK, another protein involved in cell division, to regulate astral microtubules (Gai et al., 2016). These are specific microtubules that stem from the mitotic spindle and help position it correctly during cell division. In neural progenitor cells, spindle orientation can influence whether they continue dividing or begin differentiating into specialized brain cells. This is what results in an early stop to brain development.

Animal studies connect *ASPM's* cellular function to reduced brain growth and the loss of developing neural cells. One specific study created mice with mutations that resulted in shortened *ASPM* proteins due to a loss of genetic instructions. The mutant mice developed mild microcephaly, linking abnormal *ASPM* to reduced brain size in living organisms (Pulers et al. 2010). A similar study in mice demonstrates that when researchers removed *ASPM* from the developing mouse brain, they observed increased DNA damage and apoptosis in neural progenitor cells. While some apoptosis is normal, excessive apoptosis can reduce the number of cells available to form the cortex (Tonosaki et al., 2023). Overall, the *ASPM* gene supports brain growth by maintaining both accurate division and survival of neural progenitor cells. It helps determine the correct number of cells in a developing cerebral cortex.

II - How ZIKV Damages the Developing Brain

ZIKV can reduce fetal brain growth by interfering with the reproduction and survival of neural progenitor cells. Garcez et al. (2016) exposed human neural stem cells, neurospheres, and cerebral organoids to ZIKV. Neurospheres and cerebral organoids both originate from neural stem cells and serve as laboratory models of brain tissue development. The infected neurospheres showed abnormal growth, cell detachment, and a risk of reduced survival. More importantly, over

the course of 11 days, ZIKV-exposed organoids grew approximately 40% less than uninfected cells.

A similar pattern emerged in another system: Li et al. (2016) studied prenatal ZIKV infection in mouse embryos and found that the infection concentrated in brain regions containing neural progenitor cells, where it was associated with increased apoptosis, disrupted cell division, and thinning of the cerebral cortex. The study further found reduced expression of several genes associated with microcephaly, including *ASPM*. Together, these studies show that ZIKV negatively affects the neural progenitor-cell population required for normal cortical growth.

Further evidence explains how ZIKV disrupts neural progenitor cell division. A 2016 study examined human neuroepithelial stem cells. These are stem cells that are crucial for maintaining and building up the nervous system. They later develop into radial glial cells, which produce neurons important for the cerebral cortex. Along with these stem cells, the study also observed fetal brain tissue maintained in a laboratory and tissue from one fetus with ZIKV-associated microcephaly. After conducting each observation, ZIKV primarily infected radial glial cells in the ventricular and subventricular zones, where progenitor cell division occurs. Following exposure to ZIKV, an activated protein involved in viral immunity and cell division, known as TBK1, moved away from the centrosomes. Centrosomes organize the spindle fibers that separate chromosomes during cell division. Without properly functioning centrosomes, cells may not divide correctly. The infected cells consequently showed abnormal numbers of centrosomes and disrupted mitosis, the process in which a cell divides its chromosomes between two cells (Onorati et al., 2016). A subsequent study by Gabriel et al. (2017) investigated the developmental effects of similar damage in cerebral organoids. The researchers found that ZIKV disrupted centrosome organization and altered the direction of radial glial cell division. Once again, the progenitor cells differentiated into brain cells too early, leaving fewer progenitors available to continue dividing. As a result, the developing brain tissue became thinner. Onorati et al. identify a possible cause of disrupted cell division, while Gabriel et al. show how this disruption can limit brain growth.

III - Genetic Background Influences ZIKV Outcomes

To determine whether genetic background contributes to different outcomes after prenatal ZIKV exposure, Caires-Júnior et al. (2018) studied twins whose outcomes differed despite sharing the same pregnancy. The researchers identified nine twin pairs: seven dizygotic, or fraternal, pairs and two monozygotic, or identical, pairs. The results indicated that six out of seven of the fraternal pairs were discordant, meaning one twin developed congenital ZIKV syndrome while the other did not, whereas both identical twin pairs were concordant because both twins were affected. This pattern suggested a genetic influence because identical twins share nearly all of their DNA, while fraternal twins share about half. However, the number of twins was too small to establish this relationship on its own; the twin pattern raises the possibility of genetic susceptibility, but it cannot show whether genes directly alter how cells respond to ZIKV. The researchers therefore performed an experiment using cells from three discordant fraternal twin

pairs.

Blood cells from each twin were reprogrammed into stem cells and then developed into neural progenitor cells. Before ZIKV exposure, cells from affected and unaffected twins grew at similar rates. However, cells from the affected twin showed significantly reduced growth compared to those from the unaffected twin. The distinction appearing only after infection is important because it supports a gene interaction: the decrease in growth indicated that the affected twins were unable to withstand ZIKV. Ultimately, genetic susceptibility could alter the extent to which fetal neural cells are damaged by infection.

Similarly, Borda et al. (2021) sequenced 40 unrelated children with confirmed prenatal ZIKV exposure: 29 with congenital Zika syndrome and 11 without any related prenatal complications. 10 out of the 29 affected children had at least one harmful variant in a gene linked to the features of congenital Zika syndrome. The researchers also found evidence that changes in two genes, *DISP3* and *IL12RB2*, can reduce the risk of severe effects. This matters because not every child reacted the same when exposed to ZIKV, and genetic differences can help explain why some children developed congenital Zika syndrome while others did not. Additionally, the study examined when the mothers were infected during pregnancy. Infection during the first trimester occurred in 65.5% of those affected with congenital Zika syndrome, but in only 27% of unaffected children. Because infection during this period may increase the risk of birth defects, the researchers included the timing of exposure in their analysis. Fundamentally, Borda et al. (2021) supported the idea that many factors may influence ZIKV outcomes and, more importantly, that one such factor may be genetic differences.

Overall, these studies suggest that a person's biological background may influence how ZIKV affects them. Caires-Júnior et al. (2018) reflected on how differences in biological background can lead to different responses to the same ZIKV infection. Moreover, Borda et al. (2021) found that changes in several genes may also affect the risk of severe outcomes. These findings suggest that *ASPM* could influence ZIKV-related microcephaly, although neither study establishes a relationship between inherited ASPM variants and ZIKV susceptibility.

IV - Indirectly Evaluating the Proposed *ASPM* and ZIKV interaction

The existing evidence makes an *ASPM*-ZIKV interaction plausible because reduced *ASPM* function and ZIKV infection disrupted several of the same developmental processes. Studies examined in Section I showed that major reductions in *ASPM* function can interfere with spindle organization, progenitor cell division, and brain growth and development. Furthermore, studies examined in Section II similarly showed that ZIKV can disrupt the mitotic spindle, alter the direction of progenitor cell division, and cause progenitors to develop early on. These findings are significant because both *ASPM* dysfunction and ZIKV interaction can reduce the progenitor cell population crucial for the construction of the cerebral cortex.

Karuppan et al. (2020) provided a more detailed example of how a genetic weakness can alter the effects of prenatal ZIKV infection. The researchers exposed both groups of mice to ZIKV: mice with normal *BECN1* expression and mice with one working copy of the *BECN1*

gene. These mice produced about 60% less Beclin1 protein. This specific protein helps cells break down cellular components. This process, called autophagy, can affect viral replication, inflammation, and cell growth in response to ZIKV exposure.

About 25% of the ZIKV-exposed mice with reduced *BECN1* expression developed poor growth and undeveloped brains. These effects were most noticeable two to three weeks after birth. This reflects how the same ZIKV exposure caused different outcomes in the two genetic backgrounds. Reduced Beclin1 exacerbated the effects of the infection. This supports the idea that a person's genes can affect the severity of ZIKV-related brain abnormalities and that another developmental gene can also modify the response.

The researchers also measured four genes linked to microcephaly, including *ASPM*. Even without ZIKV, mice with reduced *Becn1* expression presented lower *ASPM* expression than mice with normal *Becn1* expression. The reduction was even more significant after ZIKV exposure, reporting an overall difference of about 2.5- to 3-fold. This suggests a possible link between ZIKV susceptibility and reduced *ASPM* activity, but it does not establish whether inherited ASPM variants influence infection severity.

However, it is important to note that a reduction in *ASPM* expression is not the same as inheriting an *ASPM* variant, since the DNA sequence does not change. Therefore, the study provides indirect support for the *ASPM* hypothesis, but it does not show that inherited *ASPM* variants are responsible for the outcomes. Nevertheless, the strongest conclusion is that *ASPM* is a possible but unproven modifier.

V. Alternative Explanations for Unequal ZIKV Outcomes

While *ASPM*-ZIKV interaction is plausible, differences in infection timing, placental health, and maternal immunity could also explain why exposed fetuses experience unequal outcomes. Coutinho et al. (2021) studied the timing of infection in 511 women with confirmed ZIKV infection, representing 513 exposed fetuses. 35.9% of fetuses exposed during the first trimester faced adverse outcomes, compared with 2.8% exposed during the second trimester and 1.9% exposed during the third trimester. These outcomes included pregnancy loss, microcephaly, and other brain abnormalities.

The researchers found that infection before 11 completed weeks of pregnancy was linked to a 14-fold higher risk of an adverse outcome compared to later in the pregnancy. This explains that fetuses with similar genetic backgrounds may have different outcomes simply because infection occurs at different stages of development. Early in pregnancy, brain-forming cells multiply rapidly, so disruption during this period may have a greater effect on brain growth. Similarly, Miner et al. (2016) provided a possible explanation for how prenatal ZIKV infection can affect fetal development through the placenta. The study examined two mouse models and found that ZIKV infected trophoblasts and fetal blood-vessel cells in the placenta. Trophoblasts are placental cells that help form the barrier between the mother's and the fetus's blood. The infection led to cell death, damage to placental blood vessels, and thinning of the placental layer that exchanges materials between the mother and fetus.

The placenta normally provides the fetus with oxygen and proper nutrients while helping protect it from infection. Damage to its cells and blood vessels can reduce blood flow, weaken the protective barrier, and limit fetal growth. Additionally, the study found that placental infection was linked to restricted growth inside the uterus and could lead to fetal loss. The researchers also detected ZIKV in fetal tissue, which can actually affect brain development directly by entering fetal brain tissue and indirectly by damaging the placenta. Therefore, reduced brain growth does not require an *ASPM* variant.

Despite there being other alternative factors that could affect prenatal ZIKV outcomes, these two studies are especially important. Coutinho et al. (2021) show that early infection is strongly linked to adverse outcomes, while Miner et al. (2016) show how placental damage can impair fetal growth.

Discussion

Although current studies provide no evidence of a direct association between *ASPM* variants and congenital Zika-associated microcephaly, the findings make such a relationship biologically plausible. *ASPM* supports spindle organization and neural progenitor-cell division, whereas ZIKV disrupts centrosomes, mitotic spindles, and neural progenitor cells. Both factors can therefore reduce the number of progenitor cells required for the proper development of the cerebral cortex. This overlap suggests that a mild *ASPM* variant could leave these cells less able to withstand ZIKV-related damage.

The genetic-background studies further support this possibility. Caires-Júnior et al. (2018) found that cells from affected and unaffected twins grew similarly before infection, but cell growth in affected twins decreased after infection. This suggests that the genetic difference between the twins, specifically fraternal twins, increases vulnerability only when infection occurs. Borda et al. (2021) also found evidence that several developmental and immune-related genes may influence ZIKV outcomes. However, neither study identified *ASPM*, but together they emphasize that genetic background can influence susceptibility.

Nevertheless, Karuppan et al. (2020) provided the closest connection to *ASPM*. Approximately 25% of ZIKV-exposed offspring with reduced *BECN1* expression developed smaller, more underdeveloped brains, while several microcephaly genes, including *ASPM*, showed reduced expression. This corroborates the idea that weakening a host gene can worsen ZIKV-related effects and affect the *ASPM* gene. However, reduced *ASPM* expression is not the same as inheriting an *ASPM* variant. The study therefore provides indirect support for the hypothesis rather than evidence that *ASPM* is what specifically caused the abnormal outcome. Because this evidence remains indirect, the possible contribution of *ASPM* is more likely to be considered alongside other factors. Coutinho et al. (2021) found that infection before 11 completed weeks of pregnancy was associated with a 14-fold higher risk of an adverse outcome, while another study done by Miner et al. (2016) showed that ZIKV can damage the placenta and restrict fetal growth. These findings provide more direct evidence for infection timing and placental damage than currently exists for *ASPM*. However, together, these explanations suggest

that the *ASPM* gene could act as a modifier in a multi-step process.

Moreover, the research being compared currently has animal and human models, and does not directly identify *ASPM* as a genetic modifier. However, future research should first compare a large group of ZIKV-exposed children with microcephaly to exposed children without abnormalities. A genome-wide association study (GWAS) could first identify common genetic differences connected to significant ZIKV outcomes. Researchers could then use genome sequencing to read almost all of the DNA of affected and unaffected children and their parents. Studying the parents would show whether a child inherited an *ASPM* variant from one or both of their parents. This is important because *ASPM*-related effects may be recessive, meaning a child may need harmful variants in both copies of the gene to develop microcephaly. Additionally, this method could identify rare variants and changes in regulatory DNA that control the strength of *ASPM* activation and may be missed by GWAS. Researchers can then test whether these variants are more common in ZIKV-exposed children with microcephaly than in ZIKV-exposed children without abnormalities such as microcephaly. The groups should be matched for ancestry and account for infection timing, maternal health, and other factors that could independently influence ZIKV severity.

Researchers could create two groups of neural progenitor cells with the same genetic background, except that one group would carry an *ASPM* variant while the other would have normal *ASPM*. Each group would then be divided again so that half of the cells are exposed to ZIKV and half are not. A comparison can be made based on how quickly the cells divide, how many survive, and how much the groups grow. If the variant cells developed normally without ZIKV but showed greater damage after infection than the normal cells, this would support an interaction between the variant and the virus.

Conclusion

This review examined whether *ASPM* variants could help explain why only some fetuses exposed to ZIKV develop microcephaly. The evidence shows that *ASPM* dysfunction and ZIKV infection affect many of the same developmental processes. Human studies show that genetic background can influence responses to ZIKV, and animal research provides indirect evidence linking ZIKV susceptibility to reduced *ASPM* expression. However, current evidence doesn't examine *ASPM* variants and instead focuses on alternative factors, such as infection timing, that can influence ZIKV susceptibility. The available research therefore supports *ASPM* as a biologically plausible but unconfirmed modifier. Establishing whether inherited *ASPM* variants contribute to ZIKV susceptibility will require direct genetic comparisons and experiments that distinguish the effects of these variants from changes in gene expression caused by infection. Distinguishing inherited microcephaly from Zika-associated microcephaly matters because similar outcomes may result from different causes. Harmful variants affecting both copies of *ASPM* may independently cause inherited primary microcephaly, whereas a milder variant might only increase vulnerability when combined with ZIKV. Understanding this difference is what allows us to further understand the topic of microcephaly. Clarifying the relationship would

support more accurate diagnoses, improve the interpretation of sequencing results, and guide research into why similarly exposed fetuses experience different outcomes.

Work Cited

Borda, Victor, et al. “Whole-Exome Sequencing Reveals Insights into Genetic Susceptibility to Congenital Zika Syndrome.” PLOS Neglected Tropical Diseases, edited by Helene Dutartre, vol. 15, no. 6, June 2021, p. e0009507. DOI.org (Crossref), https://doi.org/10.1371/journal.pntd.0009507.

Caires-Júnior, Luiz Carlos, et al. “Discordant Congenital Zika Syndrome Twins Show Differential in Vitro Viral Susceptibility of Neural Progenitor Cells.” Nature Communications, vol. 9, no. 1, Feb. 2018, p. 475. DOI.org (Crossref), https://doi.org/10.1038/s41467-017-02790-9.

Coutinho, Cm, et al. “Early Maternal Zika Infection Predicts Severe Neonatal Neurological Damage: Results from the Prospective Natural History of Zika Virus Infection in Gestation Cohort Study.” BJOG: An International Journal of Obstetrics & Gynaecology, vol. 128, no. 2, Jan. 2021, pp. 317–26. DOI.org (Crossref), https://doi.org/10.1111/1471-0528.16490.

Fish, Jennifer L., et al. “Aspm Specifically Maintains Symmetric Proliferative Divisions of Neuroepithelial Cells.” Proceedings of the National Academy of Sciences, vol. 103, no. 27, July 2006, pp. 10438–43. DOI.org (Crossref), https://doi.org/10.1073/pnas.0604066103.

Gabriel, Elke, et al. “Recent Zika Virus Isolates Induce Premature Differentiation of Neural Progenitors in Human Brain Organoids.” Cell Stem Cell, vol. 20, no. 3, Mar. 2017, pp. 397-406.e5. DOI.org (Crossref), https://doi.org/10.1016/j.stem.2016.12.005.

Gai, Marta, et al. “ASPM and CITK Regulate Spindle Orientation by Affecting the Dynamics of Astral Microtubules.” The EMBO Reports, vol. 17, no. 10, Oct. 2016, pp. 1396–409. DOI.org (Crossref), https://doi.org/10.15252/embr.201541823.

Garcez, Patricia P., et al. “Zika Virus Impairs Growth in Human Neurospheres and Brain Organoids.” Science, vol. 352, no. 6287, May 2016, pp. 816–18. DOI.org (Crossref), https://doi.org/10.1126/science.aaf6116.

Gomes, Julia A., et al. “Association between Genetic Variants in NOS2 and TNF Genes with Congenital Zika Syndrome and Severe Microcephaly.” Viruses, vol. 13, no. 2, Feb. 2021, p. 325. DOI.org (Crossref), https://doi.org/10.3390/v13020325.

Higgins, Julie, et al. “Human ASPM Participates in Spindle Organisation, Spindle Orientation and Cytokinesis.” BMC Cell Biology, vol. 11, no. 1, Dec. 2010, p. 85. DOI.org (Crossref), https://doi.org/10.1186/1471-2121-11-85.

Karuppan, Mohan Kumar Muthu, et al. “Reduced-Beclin1-Expressing Mice Infected with Zika-R103451 and Viral-Associated Pathology during Pregnancy.” Viruses, vol. 12, no. 6, June 2020, p. 608. DOI.org (Crossref), https://doi.org/10.3390/v12060608.

Li, Cui, et al. “Zika Virus Disrupts Neural Progenitor Development and Leads to Microcephaly in Mice.” Cell Stem Cell, vol. 19, no. 1, July 2016, pp. 120–26. DOI.org (Crossref), https://doi.org/10.1016/j.stem.2016.04.017.

Marques, Fernanda J. P., et al. “Genetic Modifiers of Prenatal Brain Injury After Zika Virus Infection: A Scoping Review.” Genetic and Genomic Medicine, 2 Jan. 2025. DOI.org

(Crossref), https://doi.org/10.1101/2025.01.02.25319896.
Miner, Jonathan J., et al. “Zika Virus Infection during Pregnancy in Mice Causes Placental Damage and Fetal Demise.” Cell, vol. 165, no. 5, May 2016, pp. 1081–91. DOI.org (Crossref), https://doi.org/10.1016/j.cell.2016.05.008.
Onorati, Marco, et al. “Zika Virus Disrupts Phospho-TBK1 Localization and Mitosis in Human Neuroepithelial Stem Cells and Radial Glia.” Cell Reports, vol. 16, no. 10, Sept. 2016, pp. 2576–92. DOI.org (Crossref), https://doi.org/10.1016/j.celrep.2016.08.038.
Pulvers, Jeremy N., et al. “Mutations in Mouse Aspm (Abnormal Spindle-like Microcephaly Associated) Cause Not Only Microcephaly but Also Major Defects in the Germline.” Proceedings of the National Academy of Sciences, vol. 107, no. 38, Sept. 2010, pp. 16595–600. DOI.org (Crossref), https://doi.org/10.1073/pnas.1010494107.
Tonosaki, Madoka, et al. “Loss of Aspm Causes Increased Apoptosis of Developing Neural Cells during Mouse Cerebral Corticogenesis.” PLOS ONE, edited by Carlos Oliva, vol. 18, no. 11, Nov. 2023, p. e0294893. DOI.org (Crossref), https://doi.org/10.1371/journal.pone.0294893.
Wu, Kong-Yan, et al. “Vertical Transmission of Zika Virus Targeting the Radial Glial Cells Affects Cortex Development of Offspring Mice.” Cell Research, vol. 26, no. 6, June 2016, pp. 645–54. DOI.org (Crossref), https://doi.org/10.1038/cr.2016.58.

Feasibility of Nuclear Electric Propulsion for Interplanetary Transport

By Hibiki Narita

Abstract

Electric propulsion (EP) uses electrical energy to accelerate propellant to high exhaust velocities, resulting in high specific impulse and reduced propellant requirements. This paper evaluates the technical feasibility of high-power Nuclear Electric Propulsion (NEP) for crewed Mars transport targeted for the 2050 launch window, mapping the core architectural performance limits and engineering trade-offs required to realize rapid interplanetary transit. This review synthesizes published estimates of the technical requirements and limitations for the viability of a manned mission by examining trajectory dynamics, system-specific mass scaling, dynamic power conversion loops, and high-power thruster mechanics. Previous mission studies indicate that a baseline NEP architecture can achieve a one-way Mars transfer of approximately 90 days, while an advanced architecture targeting much higher power and lower specific mass can reduce the reported transfer time to approximately 39 days. The analysis also considers dynamic specific-impulse operation, modular thruster configurations, thruster lifetime, megawatt-scale heat rejection, and high-power power-processing systems. Ultimately, the paper proposes a phased technology-development roadmap from ground testing to orbital demonstrations and identifies power-system specific mass, thermal management, thruster lifetime, and high-power electronics as the primary barriers to rapid crewed Mars transportation.

Keywords

Electric Propulsion, Nuclear Electric Propulsion, Specific Mass, Interplanetary Transport, Mars Transit, Feasibility Analysis

1. Introduction

For decades, human spaceflight has been limited by the physics of chemical rocketry. To go faster, a spacecraft must carry more propellant, yet carrying more propellant increases the vehicle's mass, limiting its acceleration. Chemical engines generally have lower specific impulse than electric propulsion systems, meaning that they require more propellant to achieve a given total change in velocity. Therefore, they require massive initial mass in Low Earth Orbit to provide the propellant needed for Earth departure and interplanetary flight. Such traditional ballistic transfers require transit times of months, exposing human crews to severe galactic cosmic radiation.[1,2] Unlike chemical combustion, Electric Propulsion (EP) harnesses electrical energy to accelerate propellants to exceptionally high exhaust velocities, which drastically reduces propellant mass requirements. However, as EP systems generate very low thrust per unit of input power, complex engineering trade-offs must be made between electrical power generation, spacecraft dry mass, and total flight time.

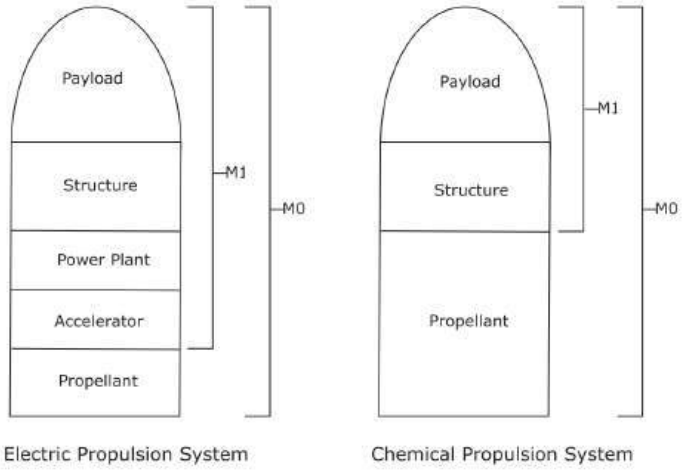

Figure 1. Recreation of mass comparison diagram[3]

The diagram above shows a mass comparison between an electrical propulsion system and a chemical propulsion system. M0 is the total initial mass and M1 is the total final mass, after the propellant is used up. Electric propulsion has higher amounts of M1, and meaning more mass remains after the acceleration.[3]

This paper explores the technical feasibility of using Nuclear Electric Propulsion (NEP), a type of electric propulsion, to support rapid crewed interplanetary transport, specifically targeting a Mars transfer by 2050, and the critical technological bottlenecks that must be resolved to enable such a mission. NEP addresses this challenge by coupling a nuclear fission reactor with high-power electric thrusters. Instead of relying on solar energy, an NEP system converts nuclear thermal energy into megawatt-scale electricity through dynamic conversion systems to electrically accelerate ionized gas propellants to extreme exhaust velocities.[4] Because NEP operates independently of solar flux, it provides continuous acceleration throughout deep-space transits without the power drop-offs characteristic of Solar Electric Propulsion (SEP). However, replacing solar arrays with heavy reactor cores, crew radiation shielding, and power conditioning units adds significant dry mass to the spacecraft.[5,6] Consequently, determining the viability of NEP requires balancing the continuous, high-efficiency acceleration it provides against the strict mass and thermal penalties introduced by onboard nuclear power generation.

Given recent literature in deep-space trajectory design and space nuclear power, evaluating this question is critical for defining the roadmap of human interplanetary exploration. While low-power EP (tens of kilowatts) has been thoroughly demonstrated on robotic platforms for station-keeping and deep-space science, existing studies often focus on uncrewed, long-duration cargo delivery rather than high-acceleration human transit.[7] Short-duration human missions require power levels of megawatts to achieve acceptable thrust levels. By evaluating the performance gaps between current state-of-the-art electric propulsion technologies and the multi-megawatt requirements of human interplanetary travel, this paper addresses an important

technological gap in current space architecture planning, identifying the exact specific mass (α) and thermal thresholds necessary to make rapid transit a reality.

This paper investigates what power-system specific mass could plausibly be achieved by nuclear electric propulsion (NEP) architectures by 2050, given current technology readiness levels and projected advances in reactor design, power conversion, and thermal management. It hypothesizes that reaching the approximately 1 kg/kW target associated with a 39-day Earth–Mars transfer will remain a significant technological challenge, as current NEP technologies have not yet demonstrated this level of system-specific mass. This analysis evaluates the gap between this target and the capabilities of existing technologies to assess its plausibility by 2050.[8–10] To evaluate this question, this paper reviews published studies on electric propulsion performance, reference mission architectures, technology readiness, and engineering constraints. Section 2 compares electric propulsion technologies using reported performance metrics. Section 3 examines reference spacecraft architectures and their associated power requirements, specific masses, and transit times. Section 4 assesses subsystem maturity using historical flight experience and reported technology readiness levels. Finally, Section 5 identifies the principal technical bottlenecks and outlines a phased development roadmap toward a potential crewed Mars mission by 2050.

2. Fundamentals and Comparative Review of Electric Propulsion

In order to evaluate the viability of multi-megawatt NEP for a manned mission to Mars, one must first analyze the fundamental physics that distinguish electric propulsion (EP) from traditional chemical combustion. Chemical propulsion has long been the standard for space transportation. However, its dependence on internally stored propellant results in substantial mass penalties on missions requiring high Δv.[9] On the other hand, EP separates energy generation from propellant storage, which enables superior exhaust velocities at the cost of lower thrust-to-weight ratios.[11] Understanding the physics behind EP, recognizing the different types of EP thrusters and their power sources, and comparing the performance of several EP systems provides the necessary fundamentals to assess why high-power NEP represents a critical architectural shift for deep space transportation.[11]

2.1 Physics and System Architecture of Electric Propulsion

The fundamental operational mechanism of EP relies on using electrical energy to accelerate propellant particles, which distinguishes it directly from conventional chemical propulsion systems. Chemical propulsion derives its thrust from the chemical reactions of propellants by expanding hot gases through a physical nozzle. Therefore, it remains strictly bound by the thermodynamic energy density of the fuel. In contrast, electric propulsion decouples the energy source from the propellant mass entirely. [11] Electric propulsion achieves exhaust velocities between 10 km/s and over 100 km/s, exceeding traditional chemical rockets, which translates directly into a dramatically higher specific impulse. (I_{sp}). [11] Exhaust velocity is

the effective speed at which propellant mass exits the back of a rocket nozzle or thruster to produce motion.

For this analysis, the EP architecture consists of five primary subsystems, or Critical Technology Elements (CTE). CTE-1: Power Source/Reactor, CTE-2: Power Conversion, CTE-3: Power Management and Distribution, CTE-4: Electric Propulsion Subsystem, and CTE-5: Primary Heat Rejection.[12] By replacing high-thrust, short-duration chemical burns with continuous, low-thrust acceleration, EP reduces required propellant mass ratios for space missions. Understanding how these systems achieve higher exhaust velocities requires examining the different physical mechanisms used to accelerate the propellant.[12]

2.2 Categorization and Acceleration Mechanisms of Electric Thrusters

EP systems are broadly categorized into electrothermal, electrostatic, and electromagnetic architectures depending on the specific mechanism used to convert electrical energy into propellant kinetic energy.[13] Electrothermal thrusters, such as arcjets and resistojets, represent the most straightforward approach: an electrical power source directly heats a gaseous propellant using high-temperature resistive heating elements or high-current electric arcs. The thermalized gas is then expanded through a conventional physical nozzle to generate thrust. While electrothermal designs offer a reliable step up from chemical combustion, their performance remains constrained by material temperature limits and nozzle erosion.[11] In contrast, electrostatic systems, such as gridded ion engines like NASA's NSTAR and NEXT, completely bypass thermal limits by using electrical fields to accelerate charged particles. These engines first ionize neutral propellants (noble gases like xenon) and then apply high-voltage electrostatic grid potentials to extract and shoot the positive ions out at extreme exhaust velocities, reaching high operational efficiencies at low propellant mass flow rates.[11]

Electromagnetic systems represent the high-power frontier of electric propulsion, relying on the interaction of electric currents and magnetic fields to generate Lorentz forces that accelerate plasma plumes. This category includes Hall-Effect Thrusters (HET), Magnetoplasmadynamic Thrusters (MPDT), Pulsed Plasma Thrusters (PPT), and advanced magnetoplasma systems like the Variable Specific Impulse Magnetoplasma Rocket (VASIMR). By accelerating plasma rather than ion beams, electromagnetic architectures overcome space-charge density limits, allowing them to process significantly higher mass flow rates and deliver Newton-level thrust. However, achieving these thrust levels requires multi-kilowatt to megawatt power throughputs, introducing severe operational challenges around electrode degradation, high-heat plasma containment, and thermal management. Scaling these electromagnetic technologies is essential for human-rated deep space architectures, as they provide the raw thrust density and high exhaust velocity necessary to shorten crewed interplanetary travel times.[11]

2.3 Power Sources and Energy Conversion Architectures

The continuous operation of electric propulsion systems requires dedicated and continuous electrical power generation, which is primarily derived from either solar or nuclear energy sources. Solar Electric Propulsion (SEP) converts solar flux into electricity via photovoltaic arrays, eliminating the need to carry power-generating fuel onboard. While SEP eliminates the mass penalty of carrying power-generating fuel, its major limitation is the inverse-square law of solar intensity.[4,13] As a spacecraft travels farther from the Sun toward Mars or the outer planets, available solar flux drops dramatically, causing a corresponding drop in electrical power output and thrust capability. The following diagram illustrates the sequential energy conversion flow within a standard SEP system:

[Solar Flux] → [Photovoltaic Arrays] → [PMAD Network] → [Electric Thruster Array]

Nuclear Electric Propulsion (NEP) bypasses solar flux dependencies by utilizing an onboard nuclear fission reactor, where the controlled splitting of uranium atoms (^{235}U) produces immense thermal energy. This process moves through four distinct operational stages to turn thermal energy into usable thrust.

1. Thermal Generation: The fission reactor core heats a circulating high-pressure working fluid (such as a helium-xenon gas mixture or liquid metal).[4,12]
2. Dynamic Energy Conversion: The heated fluid is routed through closed-loop dynamic conversion machinery, such as a Brayton cycle gas turbine or Stirling engine, to spin a generator, converting thermal energy into high-voltage electricity.[4,12]
3. Power Management and Distribution (PMAD): The PMAD network conditions, steps down, and routes thousands of volts to match the exact current and voltage requirements of the engines. [4,12]
4. Plasma Acceleration: High-power electromagnetic or electrostatic thruster arrays use this electricity to ionize and accelerate propellant to extreme exhaust velocities.[4,12]

2.4 Performance Comparison of Propulsion Types

Evaluating EP performance across different acceleration mechanisms highlights clear tradeoffs between specific impulse, thrust capacity, and electrical power demands. Gridded Ion Engines maximize propellant efficiency with specific impulses, but operate at low power levels and deliver modest thrust output, as it is constrained by the complex beam neutralization. HET and PPT systems offer dependable performance for orbital logistics, but channel erosion and lower efficiency limit their capabilities for deep space travel. High-power electromagnetic architectures, including MPDT and VASIMR, reach hundreds of kilowatts and megawatts of power, and they deliver Newton-level thrust and variable continuous operations at the expense of massive thermal management and power-handling difficulty.[11,12]

Propulsion Type	Specific Impulse (s)	Thrust Range	Power Range (kW)
Ion Engine	3000-10000	10-100 mN	1-10
HET	1500-3000	50-200 mN	0.5-10
PPT	500-1200	μN-mN	1
MPDT	1000-5000	0.1-5 N	100-1000
VASIMR	1000-5000	mN-N	100-500+

Figure 2. Comparison of Advanced Electric Propulsion Systems[11]

2.5 Conclusion

High-power electric propulsion represents a critical architectural shift for deep space transportation, offering superior propellant efficiency through electrostatic, electromagnetic, and electrothermal acceleration mechanics.[13] While SEP serves as an effective baseline for lower power orbital logistics, NEP emerges as the superior power architecture for ambitious Earth-to-Mars transit profiles and deep space missions.[4,13] By decoupling energy generation from solar flux availability, an onboard fission reactor provides the high-density electrical output required to drive heavy electromagnetic thruster arrays continuously, drastically reducing total deep space transit time.[4,12] Ultimately, balancing the system specific mass (α) of these nuclear power loops against the operational lifespans and erosion limits of high thrust electric engines will determine the technical viability of next-generation crewed and cargo architectures traversing the solar system.[2,11,12]

3. Reference Mission Architecture

To evaluate the technical feasibility of a 2050 crewed Mars mission using NEP, this section reviews reference mission architectures proposed in previous studies. Long-duration spaceflight exposes human crews to continuous cosmic radiation and microgravity-induced physiological degradation. Consequently, minimizing interplanetary transit times is the primary objective of this crewed architecture. Achieving rapid transit requires synthesizing numerical trajectory modeling, multi-megawatt reactor power plant scaling, dynamic propellant management, and fault-tolerant thruster array configurations into a unified mission framework.

3.1 Trajectory Dynamics and Multi-Phase Mission Execution

The reference mission architecture utilizes continuous low-thrust propulsion operating across three distinct trajectory phases. Unlike chemical missions that rely on short, impulsive maneuvers, low-thrust trajectories apply continuous acceleration over extended operational durations. This continuous thrust profile alters orbital energy gradually, requiring precise numerical integration techniques to model orbital escape, heliocentric cruise, and target capture phases.

Phase 1 begins with an Earth escape spiral originating from a circular Low Earth Orbit (LEO) at an altitude of 200km.[14] The NEP system executes continuous forward numerical integration burns to gradually climb out of Earth's gravity well toward the outer edges of Earth's Sphere of Influence (SOI). Simulation data establishes that this spiraling departure requires a total phase duration of approximately 42.1 days, involving 36.6 days of continuous active thrusting.[14]

Phase 2 initiates as the spacecraft crosses the boundary of the Earth SOI and enters the heliocentric transfer phase. Deep-space trajectory management utilizes direct optimal control algorithms to solve the non-linear transfer problem. These algorithms adjust vehicle orientation, acceleration vectors, and power allocations to maximize velocity while managing total propellant consumption. Continuous acceleration along optimized heliocentric vectors allows the vehicle to achieve higher average speeds than unpowered chemical coast trajectories.[8,14]

Phase 3 completes the transit upon reaching the Mars SOI. The system executes a retro-burn capture spiral to descend into a Low Mars Orbit (LMO) at an altitude of 400km. Trajectory optimization ensures precise orbital insertion into LMO without over-consuming remaining propellant reserves.[14]

3.2 Reactor Scaling and Power Density Performance Limits

The overall performance of a crewed NEP spacecraft depends on the specific mass (α) of its nuclear power plant, defined as total system mass per kilowatt of electrical power generated ($\alpha = m_{power}/P_{energy}$). Because NEP reactors operate independently of solar distance, the system provides constant power output throughout the entire trajectory, reducing overall transit time by approximately 10% compared to solar electric options.[4,13] The performance capabilities of the architecture scale directly with power density and specific mass benchmarks.

A baseline architecture using approximately 12MW of electrical power and a specific mass of $\alpha = 4$kg/kW is associated with a reported one-way Mars transfer time of approximately 90 days. This baseline provides a significant reduction in deep-space radiation exposure compared to a 6-to-9-month one-way chemical transfer to Earth to Mars.[1,4]

An advanced architecture considered in the literature increases electrical power to approximately 200 MW while reducing specific mass to approximately 1 kg/kW. The referenced architecture reports a one-way Earth–Mars transfer time of approximately 39 days.[1,11] Achieving this lower specific mass requires developing high-temperature gas-cooled reactor cores and lightweight power conversion units capable of operating at high thermal efficiencies.

However, the reported 39-day transfer should be interpreted as an advanced architecture-level projection rather than a demonstrated mission capability. The study's authors were affiliated with the organization developing VASIMR, a relevant consideration when evaluating the independence of its performance projections.[1,11] Furthermore, the associated specific-mass target of approximately 1 kg/kW represents an ambitious design objective rather than a demonstrated capability of an operational space nuclear-electric power system.

The figure below shows the rough relationship between transit time and system specific mass. However, the transit time is dependent on other variables including available electrical power, power-to-mass ratio and specific impulse, and not solely on specific mass.

Specific Mass (kg/kW)	Power Level (MW)	One-Way Transit Time (days)
4.0	12	90
1.0	200	39

Figure 3. Comparison of Two Published NEP Mission Architectures

3.3 Thrust Metrics, Payload Trade-Offs, and Dynamic I_{sp} Operations

System thrust allocation dictates a direct trade-off between total transit speed and delivered payload mass. One study reports that a continuous thrust of 4 N yields a final payload mass ratio of 46.45% over a 295.59-day flight, while increasing thrust to 20 N reduces the payload mass ratio to 20.58% and shortens the transit time.[14] These results illustrate the trade-off between transit duration and payload delivery within the mission scenario analyzed in that study. However, the 295.59-day result differs substantially from the 90-day and 39-day reference architectures discussed in Section 3.2. Because these estimates may rely on different assumptions about spacecraft design, power, and trajectory, they should not be directly compared without examining the underlying mission parameters.

To optimize this performance trade-off, high-power electromagnetic thrusters like VASIMR® utilize dynamic variable specific impulse (I_{sp}) modulation. By varying I_{sp} between 4,000 s and 30,000 s at constant input power, the propulsion system adapts its acceleration mechanics to match mission phase requirements.[1,11] High mass flow rates and lower I_{sp} are applied during planetary escape and capture spirals to maximize thrust output. During heliocentric cruise, the system shifts to high I_{sp} to maximize exhaust velocity. Dynamic I_{sp} tuning saves between 6 and 7 metric tons of propellant compared to constant I_{sp} baselines, preserving critical mass allowance for life support systems and habitat structures.[1]

3.4 Modular Thruster Arrays and Reliability Boundaries

Human rating a multi-megawatt propulsion system requires built-in redundancy to prevent single-point failures during interplanetary transit. Rather than deploying a single large engine, the reference architecture utilizes a clustered array of modular thrusters. For example, a target total system thrust of 16N can be delivered by four operational 4N thrusters working in parallel.[14]

This multi-thruster layout provides essential operational fail-safe capabilities. If a single thruster unit experiences a power processing or hardware failure, the system redistributes power across the remaining engines to complete the mission safely.[10,15] Additionally, distributing total thrust output lowers individual engine lifespan requirements to 6,384 hours (including a 2 times

safety factor). Spreading operational wear across multiple units ensures that component erosion stays within manageable hardware boundaries throughout the mission.

3.5 Conclusion

Evaluating the feasibility of a multi-megawatt NEP vehicle for deep-space transport requires a comprehensive assessment of subsystem maturity, heritage hardware, and operational reliability. While low-power EP has achieved total flight qualification on robotic spacecraft, scaling these architectures to megawatt-class regimes introduces severe qualification gaps across thermal, electrical, and structural domains. Analyzing Technology Readiness Levels (TRLs), historical flight heritage, failure modes, and critical engineering margins provides the framework necessary to establish a realistic development roadmap for high-power interplanetary transportation.

4. Mission Feasibility & Heritage Analysis

4.1 Technology Readiness Levels and Critical Technology Elements

Establishing overall mission feasibility begins with classifying the system architecture into discrete Critical Technology Elements (CTEs) and assessing their respective Technology Readiness Levels (TRLs) based on flight heritage and ground testing metrics.

The primary power source, designated as CTE-1, consists of a high-temperature fission reactor assembly utilizing enriched Uranium (^{235}U) coupled with refractory metal alloy cladding.[4,12] While these alloys are structurally capable of surviving extended thermal stresses, megawatt-class space reactors remain at TRL 3-4 due to a persistent lack of integrated full-power vacuum ground testing.[2,15] Closely integrated with the reactor is CTE-2, the dynamic power conversion system. Closed-loop Brayton and Potassium-Rankine dynamic power conversion cycles offer high thermal-to-electric conversion efficiency exceeding 25-30%.[4,16] Although small-scale Brayton units have demonstrated extensive operational hours in simulated vacuum loops, multi-megawatt turbomachinery, high-temperature gas bearings, and working-fluid containment loops still require system-level flight qualification to advance beyond TRL 4-5.[4,12,17]

Routing this generated energy requires CTE-3, the Power Management and Distribution (PMAD) architecture. High-voltage PMAD systems must step down kilovolt-level outputs from dynamic generators to match the impedance and current requirements of high-power thruster arrays. Processing megawatt-scale throughput requires lightweight, high-temperature silicon carbide (SiC) or gallium nitride (GaN) power electronics alongside active cooling channels to manage parasitic thermal dissipation, keeping the technology currently at TRL 4.[7,10,15] The actual thrust generation is handled by CTE-4, the electric propulsion subsystem. High-power electromagnetic thrusters, such as Magnetoplasmadynamic Thrusters (MPDT) and the Variable Specific Impulse Magnetoplasma Rocket (VASIMR), have demonstrated hundred-kilowatt-class operation in ground vacuum chambers at TRL 5-6.[1,9,11]. However, scaling these engines into

operational megawatt arrays requires validating long-duration electrode erosion rates, cathode life limits, and high-density plasma containment under continuous operations.[9,17,18]

Finally, CTE-5 encompasses the primary heat rejection radiators, which represent the dominant mass driver for megawatt NEP architectures. Operating dynamic power cycles requires rejecting low-temperature waste heat via large-area liquid metal heat-pipe radiator panels.[4,19] Mitigating single-point failure modes from micrometeoroid and orbital debris (MMOD) impacts requires advanced structural designs, including composite self-healing materials and armored fluid loops, placing current radiator designs at TRL 3-4.[20]

4.2 Flight History and Operational Heritage

High-power NEP draws foundational operational concepts from decades of flight-proven SEP and low-power nuclear space power systems. Evaluating historical antecedents highlights both key milestones achieved and the structural failure modes that must be addressed in future deep-space architectures.

The flight data from historical platforms demonstrates that long-duration electric propulsion is operationally viable, but it highlights distinct vulnerability points in power conditioning and mechanical endurance. Early space nuclear demonstrations like SNAP-10A confirmed that orbital reactor startup and fluid thermal loops operate as predicted in microgravity, though electrical control circuitry proved to be a single point of failure.[4,15] Subsequent robotic missions such as Deep Space 1 and Dawn shifted focus toward electric thruster longevity. Deep Space 1 successfully logged over 16,000 hours of continuous thrust using its NSTAR gridded ion engine, but revealed critical failure mechanisms including grid erosion and high-voltage short circuits caused by conductive flake deposition. Dawn expanded on this heritage by validating multi-thruster cross-strapping over nearly six years of operation, while also demonstrating that active thrust-vectoring could compensate for mechanical failure of primary attitude control hardware.[13,21,22]

Ground testing platforms like NASA's Kilopower (KRUSTY) project have further bridged the gap by proving the dynamic stability of nuclear cores integrated with Stirling converters. To bridge the remaining gap between historical low-power SEP and high-power NEP, system risk mitigations rely on incorporating automated high-voltage arc-clearing circuitry, redundant dynamic conversion loops, active thrust-vectoring gimbals, and modular thruster switching networks.[4,15]

4.3 Performance Margins and Feasibility Trade Studies

Mission feasibility is bound by strict system-level performance margins to accommodate structural mass growth, thermal degradation, and orbital dynamics throughout all phase transitions. Mass margin is an important consideration in spacecraft design because it accounts for potential increases in dry mass as the design develops. The following relationship expresses mass margin as a percentage of dry mass:

$$\text{Mass Margin } (\%) = (M_{capacity} - M_{dry}) / M_{dry} \times 100 \geq 20\ \%$$

In accordance with standard space system engineering practices, deep-space NEP architectures maintain a minimum 20–30% dry mass growth margin at Phase A/B design states to account for heavy PMAD shielding and heat rejection structures.[6,20,22]

Beyond structural mass growth, the feasibility of rapid Earth–Mars transfers depends heavily on achieving a low system specific mass, defined as $\alpha = M_{power} / P_{electric}$. The literature considered in this paper reports specific-mass values on the order of several kg/kW for megawatt-class NEP concepts, while the rapid-transfer architecture requires a substantially lower value of approximately 1 kg/kW. Values exceeding 15kg/kW significantly increase vehicle inertia, ultimately negating the flight-time advantages over advanced chemical-aerocapture options.[2,8,14]

Environmental degradation also dictates hardware oversizing. Radiator surfaces experience degradation from atomic oxygen in low Earth orbit and space weathering in interplanetary space. Consequently, primary heat rejection loops incorporate a 15% area margin to maintain nominal cycle efficiencies over a 5-year mission lifespan.[16,19] Furthermore, power electronics and control modules are heavily shielded against total ionizing dose (TID) effects to withstand the cumulative neutron and gamma flux generated by the primary reactor core throughout continuous operation.[5,23]

5. Technical Bottlenecks and Future Roadmaps

Translating multi-megawatt NEP from conceptual mission architectures into operational reality for a 2050 crewed Mars transfer requires resolving critical engineering bottlenecks. While the theoretical performance metrics of high-power electric propulsion offer significant advantages over chemical systems, current TRLs remain constrained by hardware limits across power density, thermal management, component erosion, and high-power electronics.[10,15] Addressing these challenges requires a structured development roadmap that maps subsystem maturation from present research baselines to flight qualification.

5.1 Primary Technical Bottlenecks

Bottleneck 1: Lowering Power Plant Specific Mass (α)

The defining metric for rapid interplanetary NEP transit is system specific mass (α), expressed as total power plant mass per kilowatt of electrical output (kg/kW).[4] The reference architectures considered in this paper associate a 90-day transfer with approximately 12 MW of power and a specific mass of 4 kg/kW, while the reported 39-day architecture uses approximately 200 MW and a specific mass near 1 kg/kW.[1] Current space nuclear reactors and energy conversion systems operate at specific mass values far higher than these targets. Reducing α requires transitioning from heavy shielding and legacy materials to lightweight, high-temperature refractory metal alloys and oxide dispersion-strengthened (ODS) steels capable of enduring long-term nuclear radiation and thermal stress.[5] Furthermore, dynamic thermal-to-electric conversion systems—such as Brayton or Stirling cycles—must achieve significantly higher

thermal efficiencies to reduce total reactor core size and mass.[4,16] The 1 kg/kW benchmark should be treated as a development target whose feasibility remains to be established, rather than as a validated specification for a future flight system.

Bottleneck 2: Megawatt-Class Thermal Heat Rejection

Energy conversion systems are inherently limited by thermodynamic efficiency, meaning a multi-megawatt NEP reactor must reject massive quantities of waste heat into the vacuum of space.[15] For a 12-200MW electrical system operating at typical conversion efficiencies, hundreds of megawatts of thermal waste must be radiated continuously. Because radiation is the sole heat dissipation mechanism in space, the required radiator surface area scales dramatically. Current radiator technologies are too heavy and susceptible to micrometeorite impacts to meet mission constraints. Breakthroughs in ultralight carbon-composite heat pipes, liquid metal loops, and self-healing materials are essential to construct large-scale radiator structures without exceeding spacecraft mass budgets.[15]

Bottleneck 3: Thruster Lifetime and Channel Erosion

High-power electromagnetic and electrostatic thrusters subjected to continuous operation over thousands of hours suffer severe structural degradation due to plasma erosion.[9,10] In gridded ion and Hall-effect thrusters, high-energy ion bombardment gradually erodes acceleration grids and discharge channel walls, leading to efficiency losses and premature hardware failure. To achieve the required continuous operational lifetime—such as the 6,384-hour single-engine lifespan target identified for clustered 4N arrays—thruster designs must shift toward electrodeless plasma architectures (e.g., VASIMR® or inductive plasma thrusters) or integrate magnetic shielding configurations that keep high-energy ions from contacting physical channel surfaces.[11,18]

Bottleneck 4: High-Power Electronics and Power Processing Units (PPUs)

PMAD networks must process and route high-voltage, megawatt-scale electrical energy from the reactor conversion system to the thruster arrays.[7,23] Current space-grade Power Processing Units (PPUs) are engineered for kilowatt-class applications and exhibit high mass-to-power ratios. Scaling PPUs to multi-megawatt operational levels while maintaining high electrical efficiency (>95%) requires advanced wide-bandgap semiconductors (such as silicon carbide or gallium nitride), lightweight high-voltage cabling, and dedicated active thermal management to prevent electronic overheating.[15]

5.2 Technology Readiness Level (TRL) Maturation Roadmap

To bridge the gap between present laboratory demonstrations and a human-rated flight system by 2050, key subsystems must progress systematically through established TRL milestones.[15]

Figure 4. NEP Subsystem TRL Maturation Table

Subsystem Element	Current TRL	2035 Target	2050 Target
High-Temp Space Fission Reactors	TRL 3-4	TRL 6	TRL 9
Megawatt Lightweight Radiators	TRL 3	TRL5-6	TRL 9
High-Power Thrusters (VASIMR / MPDT)	TRL4-5	TRL6-7	TRL 9
Megawatt Power Processing Units (PPU)	TRL3-4	TRL 6	TRL 9

Phase I: Ground Qualification and Material Testing (2026–2035)

Focuses on validating high-temperature core materials, testing liquid metal heat pipe radiator prototypes in vacuum chambers, and conducting long-duration endurance burns of high-power plasma thrusters. Subsystems must reach TRL 6 through rigorous ground testing in simulated space environments.[12]

Phase II: Orbital Demonstrations and Scaling (2035–2045)

Involves launching sub-scale in-space technology demonstrations (e.g., 100kW to 1MW NEP precursor flights) to test reactor deployment, active thermal rejection, automated power processing, and dynamic I_{sp} modulation in Low Earth Orbit and cislunar space.[4,12]

Phase III: System Integration and Human Rating (2045–2050)

Integrates full-scale multi-megawatt (12-200MW) reactor units, lightweight radiator arrays, and clustered thruster configurations onto the primary crew transit vehicle. Comprehensive operational validation establishes system redundancy and safety margins required for TRL 9 human-rated flight readiness ahead of the 2050 Mars launch window.[1,2]

6. Conclusion

High-power Nuclear Electric Propulsion (NEP) has the potential to enable rapid crewed transportation between Earth and Mars, but its feasibility depends on several unresolved technological challenges.[4,5,8] By decoupling energy generation from traditional chemical propellant limits and solar flux constraints, multi-megawatt NEP provides the continuous, high-efficiency acceleration required to drastically reduce transit times and shield human crews from prolonged deep-space radiation and microgravity degradation.[1,6,21] However, achieving operational feasibility depends directly on overcoming critical technical bottlenecks: specifically driving power plant specific mass down to $\alpha \leq 1$kg/kW, scaling megawatt-class thermal rejection radiators, and extending thruster operating lifespans.[9,14,15]

Putting the core body sections into conversation illustrates how physics, trajectory design, heritage data, and hardware limits intersect to define this mission architecture. The high exhaust velocities of 10-100km/s and superior specific impulse (I_{sp}) established in Section 2 provide the

physical mechanism that enables Section 3's low-thrust, continuous trajectory profiles.[18,21] The trade-offs between thrust density and efficiency directly dictate mission dynamics.[6,22] Modulating I_{sp} dynamically, such as via VASIMR architectures operating between 4,000s and 30,000s, allows the spacecraft to balance the high thrust required for planetary escape spirals with the high exhaust velocity needed for heliocentric cruise, preserving critical propellant mass for life support and payload budgets.[1,11,19]

Evaluating past flight history in Section 4 contextualizes the subsystem TRLs and engineering bottlenecks analyzed in Section 5.[12,15] Historical missions like Deep Space 1 and Dawn proved that electric propulsion is operationally viable over thousands of hours, yet they highlighted single-point vulnerability modes such as grid short-circuits, channel erosion, and attitude control degradation.[13] Bridging the gap between low-power heritage systems under 10kW and human-rated megawatt architectures requires addressing these exact failure mechanisms through magnetoplasma shielding, electrodeless designs, and modular thruster clustering.[10]

Ultimately, the literature reviewed in this paper identifies power-system specific mass, defined as $\alpha = m_{power} / P_{energy}$, as a major factor influencing the feasibility of rapid interplanetary transit.[4,14] The reference architectures illustrate this relationship: a 12 MW system with a specific mass of approximately 4 kg/kW is associated with a reported 90-day transfer, while a separate advanced architecture using approximately 200 MW and approximately 1 kg/kW is associated with a reported 39-day transfer.[2] Values exceeding 15kg/kW introduce excessive vehicle inertia that negates the transit time advantages over chemical-aerocapture options.[8]

Several critical questions arise from this synthesis that define the research roadmap toward 2050. First, can megawatt-class space nuclear reactors and dynamic conversion loops be fully validated on the ground given current facility and environmental limitations?[16] Second, what advanced structural materials can simultaneously meet the extreme demands of nuclear radiation, thermal stress, and micrometeoroid impact shielding without exceeding strict spacecraft mass growth margins?[5]

To answer these questions, future scholarly study should prioritize two specific methodologies. Scholars and space agencies should utilize sub-scale orbital testbeds ranging from 100kW to 1MW in cislunar space during Phase II between 2035 and 2045, as operating full nuclear-thermal-electric loops in real microgravity and radiation environments provides empirical data that ground vacuum chambers cannot fully replicate.[4,12,15] Additionally, researchers should employ accelerated particle-in-cell erosion modeling combined with high-temperature materials testing on refractory metal alloys, ODS steels, and carbon-composite self-healing heat pipes, using advanced computational modeling validated by targeted proton beam experiments to accelerate the qualification of TRL 6–9 hardware required for human-rated flight.[5,7,11]

Acknowledgements

The author wishes to thank Dr. Devin Carroll and Mr. Kieran Tait for their extensive support.

Works Cited

(1)Ilin, A. V.; Cassady, L. D.; Glover, T. W.; Diaz, F. R. C. VASIMR® Human Mission to Mars. 2011.

(2)Rollins, S. The Requirements of a Nuclear Propulsion System for the Human Exploration of Mars.

(3)Saltzer, C.; Craig, R. T.; Fetheroff, C. W. Comparison of Chemical and Electric Propulsion Systems for Interplanetary Travel. Proc. IRE 1960, 48 (4), 465–476. https://doi.org/10.1109/JRPROC.1960.287394

(4)Mason, L.; Oleson, S.; Jacobson, D.; Schmitz, P.; Qualls, L.; Smith, M.; Ade, B.; Navarro, J. Nuclear Power Concepts and Development Strategies for High-Power Electric Propulsion Missions to Mars. Nucl. Technol. 2022, 208 (sup1), 6. https://doi.org/10.1080/00295450.2022.2045180

(5)El-Genk, M. S.; Tournier, J.-M. A Review of Refractory Metal Alloys and Mechanically Alloyed-Oxide Dispersion Strengthened Steels for Space Nuclear Power Systems. J. Nucl. Mater. 2005, 340 (1), 93–112. https://doi.org/10.1016/j.jnucmat.2004.10.118

(6)Sankaran, K.; Cassady, L.; Kodys, A. D.; Choueiri, E. Y. A Survey of Propulsion Options for Cargo and Piloted Missions to Mars. Ann. N. Y. Acad. Sci. 2004, 1017 (1), 450–467. https://doi.org/10.1196/annals.1311.027

(7)Janhavi Kulkarni. Electric Propulsion for Space Exploration: Principles, Technologies, Challenges, and Future Directions. Acceleron Aerosp. J. 2026, 6 (1), 1653–1659. https://doi.org/10.61359/11.2106-2602

(8)Bérend, N.; Moreno, E. C.; Ruault, J.-M.; Epenoy, R. Feasibility Assessment of Rapid Earth–Mars Transfers Using High-Power Electric Propulsion. J. Spacecr. Rockets 2014, 51 (3), 946–957. https://doi.org/10.2514/1.A32560

(9)Dale, E.; Jorns, B.; Gallimore, A. Future Directions for Electric Propulsion Research. Aerospace 2020, 7 (9), 120. https://doi.org/10.3390/aerospace7090120

(10)Subodh Avinash Wagh; Pooja Pawar. Challenges and Innovations in Electric Propulsion Systems for Space Transportation: A Comprehensive Review. Acceleron Aerosp. J. 2025, 4 (2), 929–944. https://doi.org/10.61359/11.2106-2511

(11)Nadarajan Assari Syamala, S. R.; Kareem Navaz, F. M. Advancements in Plasma Propulsion Technologies: VASIMR and Emerging Systems for Interplanetary Exploration; 2025; pp 1–9. https://doi.org/10.52202/080554-0006

(12)Martin, A. K.; Polzin, K. A.; Curran, F. M.; Myers, R. M.; Rodriguez, M. A. A Technology Maturation Plan for the Development of Nuclear Electric Propulsion; Huntsville, AL.

(13)Casanova-Álvarez, M.; Navarro-Medina, F.; Tommasini, D. Feasibility Study of a Solar Electric Propulsion Mission to Mars. Acta Astronaut. 2024, 216, 129–142. https://doi.org/10.1016/j.actaastro.2024.01.001

(14)Choi, S.-J.; Choi, M. Time-Optimal Earth–Mars Transfer Using Nuclear Electric Propulsion. Nucl. Eng. Technol. 2025, 57 (10), 103685. https://doi.org/10.1016/j.net.2025.103685

(15)Polzin, K. A.; Martin, A. K.; Curran, F. M.; Myers, R. M.; Rodriguez, M. A. Strategy for Developing Technologies for Megawatt-Class Nuclear Electric Propulsion Systems.
(16)Tirila, V.-G.; Munro-O'Brien, T.; Negre, B.; Wittig, A.; Ryan, C.; Dinan, R.; Malley, D.; Syed, T.; Reed, N.; Schwageraus, E. Development of a High Power Nuclear Electric Propulsion System for Interplanetary Missions. Th Int. Astronaut. Congr.
(17)Karkadakattil, A. Nuclear Electric Propulsion for Deep-Space Missions: Technologies, Challenges, and Future Outlook. Nucl. Sci. Eng. 2026, 0 (0), 1–25. https://doi.org/10.1080/00295639.2025.2607945
(18)Bathgate, S. N.; Bilek, M. M. M.; Mckenzie, D. R. Electrodeless Plasma Thrusters for Spacecraft: A Review. Plasma Sci. Technol. 2017, 19 (8), 083001. https://doi.org/10.1088/2058-6272/aa71fe
(19)Miao, X.; Zhang, H.; Wang, Q.; Xia, Y.; Sun, W. Optimum Design of Nuclear Electric Propulsion Spacecraft for Deep Space Exploration. Energy Rep. 2022, 8, 9629–9641. https://doi.org/10.1016/j.egyr.2022.07.146
(20)Alexandrovich, S. A. Interplanetary Transport System "Horizon."
(21)Martinez-Sanchez, M.; Pollard, J. E. Spacecraft Electric Propulsion-An Overview. J. Propuls. Power 1998, 14 (5), 688–699. https://doi.org/10.2514/2.5331
(22)Landau, D.; Chase, J.; Randolph, T.; Timmerman, P.; Oh, D. Electric Propulsion System Selection Process for Interplanetary Missions. J. Spacecr. Rockets 2011, 48 (3), 467–476. https://doi.org/10.2514/1.51424
(23)Dudukovich, R.; Fraire, J. A. Evolution of Interplanetary Internet: Technology Demonstrations for Space Network Operations. Th Int. Conf. Space Oper.

Understanding Inflation in Uzbekistan, 2016–2025 By Akbarxon Nasirov

Abstract

Between 2016 and 2025, inflation in Uzbekistan ranged from 7.3% to 18.8% and peaked after the 2017 exchange-rate liberalization. This paper asks which of four factors (GDP growth, trade openness, the exchange rate and the policy rate) moved with these changes. Using 10 annual observations and Pearson correlation, it finds that the year-on-year change in the som–dollar rate is positively correlated with inflation ($r = +0.80$, $p = 0.01$, $n = 9$), but that this result depends heavily on 2017. GDP growth (−0.34), trade openness (−0.09) and the policy rate (+0.39) are not statistically significant. The positive policy-rate correlation most likely reflects the central bank raising rates in response to inflation. With only 10 annual observations the results are indicative rather than causal; quarterly data would allow a regression.

1 Introduction

Uzbekistan experienced large swings in inflation, with rates at times nearly doubling before falling sharply in later years (2016-2025), raising a question: what were the main factors behind these fluctuations? To contextualize this analysis, it is first necessary to define inflation. Inflation is the general increase in the prices of goods and services over time, which reduces the buying power of money, and possible factors for changes in inflation can be trade openness and gross domestic product (GDP). The period between 2016 and 2025 is chosen because inflation was moderate in 2016, spiked in 2017 after economic liberalization, then declined toward 2025. Trade openness is the degree to which a country integrates with global markets through exports, imports, and capital flows. It is commonly measured by taking the sum of a country's total exports and imports and dividing that value by its gross domestic product (GDP). Furthermore, in open economies, the country's production possibilities increase by trade expansion, resulting in higher productive output and low inflation (Binici et al.). Trade openness can affect inflation in two different ways. It can reduce inflation by exposing domestic producers to foreign competition, which limits their ability to raise prices. At the same time, it can increase inflation by making the economy more dependent on imported goods, whose prices rise when the domestic currency weakens, as occurred in Uzbekistan after 2017, when the som lost roughly half its value and imported goods became far more expensive in local terms. Therefore, which effect dominates depends on a country's trade composition. In September 2017, the Central Bank unified the official and black-market exchange rates and allowed the Uzbek som to float freely (Russell). After the trade liberalization, Uzbekistan's foreign trade increased by 18% that year, and trade openness (measured by exports plus imports divided by GDP) rose from roughly 26% in 2016 to its peak value of 63% by 2023, later decreasing to around 57% by 2024-25 (Russell; "Trade"). Gross Domestic Product (GDP) measures the total monetary value of all goods and services produced within a country over a given period, such as a year or a quarter. It operates as a primary indicator of an economy's size and overall health. In 2017 Uzbekistan's GDP growth decreased to 4.7%, recovered to 6.8% by 2019, and fell sharply to just 1.6% in 2020 because of

COVID-19, before rebounding to 8.2% the following year ("GDP Growth"). In 2025, GDP growth in Uzbekistan reached 7.7%, its highest rate since 2021. In 2017, GDP growth decreased to 4.7% while inflation almost doubled to 18.8%. In 2024, GDP growth increased slightly from 6.3% to 6.7%, but inflation also increased from 8.8% to 9.8%. Conversely, when growth comes from productivity improvements rather than rising demand, output expands and price pressure eases. So GDP growth and inflation do not always move in opposite directions. Other factors, such as exchange rates or the monetary policy regime, may also matter. Uzbekistan's economy has grown substantially over the last decade. While trade openness changes gradually in most countries, Uzbekistan's change was unusually rapid. It can be seen that it went from one of the most closed economies in the region — trade openness of just 26% in 2016, with a fixed exchange rate and heavy import restrictions that gave rise to a flourishing black market — to 63% by 2023 after the 2017 liberalization. The 2017 reform changed several things simultaneously. It is true that it led to higher inflation that year, but it shifted the exchange rate, trade policy, and the monetary policy framework, which contributed to lower inflation in later years. So that's part of what makes Uzbekistan an interesting case to study. This paper studies the extent to which trade openness, GDP growth, exchange rates, and monetary policy explain changes in inflation in Uzbekistan from 2016 to 2025, using descriptive analysis and correlation. The following section reviews the existing literature on these relationships; the paper then presents the data and methodology, reports the results, and concludes with a discussion of the findings.

2 Literature Review

2.1 Trade openness and inflation

Binici et al. found that trade openness affects inflation not directly. It is affected through market competition and productivity; once these two factors are controlled for, trade openness itself becomes statistically insignificant. This is because if trade openness brings increased competition and productivity, then it can lower inflation, but if there is trade openness and no competition and productivity, it does not reduce inflation.

Embergenov et al. found that trade openness increases inflation in Uzbekistan specifically. Within this overall effect, export openness was linked to lower inflation, while import openness was linked to higher inflation. This suggests that the effect of trade openness on inflation depends on a country's trade composition — the balance between exports and imports — rather than on its overall level alone. Two countries with similar levels of trade openness could see very different inflation outcomes depending on how that openness is composed.

2.2 Exchange rates and inflation

Poghosyan studied seven Caucasus and Central Asia (CCA) countries, including Uzbekistan, using monthly data from 1995 to 2020, focusing on the speed of exchange rate pass through (ERPT). The study found that pass-through averages about 10 percent on impact and about 25 percent after 12 months. Pass-through was lower by roughly 2 percentage points on

impact, and 5 percentage points after 12 months, in energy-exporting countries — a group that includes Uzbekistan — compared with energy-importing ones. Although the relationship between exchange rate changes and inflation is positive, the direction of causality is not obvious from correlation alone. A Granger causality test showed that exchange rate changes cause inflation, while inflation does not cause exchange rate changes, supporting the use of inflation as the dependent variable and exchange rate changes as the explanatory variable. The study also found no significant difference in pass-through between fixed and floating exchange rate regimes, suggesting that a transition to floating does not itself impose additional inflationary costs.

Prior to September 2017, Uzbekistan's official exchange rate had been tightly pegged, a system that benefited state-owned and well-connected enterprises able to buy artificially cheap foreign currency, as well as criminal networks running a black market with officials bribed to look the other way. Ordinary Uzbeks, by contrast, were forced to exchange money illegally on the black market, and the restrictive system deterred foreign investment by making it difficult to move currency out of the country. When the Som was floated in September 2017, the effect was immediate: the currency devalued by almost half overnight, and the black market disappeared. Inflation surged in response, nearly doubling to reach close to 20% by early 2018. In the longer term, however, this liberalization was expected to attract more investment and make the economy more open and competitive (Russell).

2.3 Economic growth and inflation

Khan and Senhadji investigated whether inflation affects growth differently above and below a certain level. They found evidence of a threshold above which inflation has a negative effect on growth. The threshold is lower for industrial than for developing countries. For industrial countries the estimates are 1-3% and for developing countries 11-12%. Uzbekistan's inflation rate exceeded the 11-12% developing-country threshold in 2017, 2018, 2019, and 2022, while falling below it in other years — which suggests the inflation-growth relationship may not have been constant across the study period (2016-2025).

Beknazarov examined Uzbekistan's key macroeconomic indicators using annual data from 2010 to 2025, treating GDP growth as the dependent variable. The study found that investment activity and industrial production have a statistically significant positive effect on GDP growth, confirming the central role of capital accumulation and real sector expansion in Uzbekistan's economic development. The services sector contributes positively but to a lesser degree, while inflation shows a negative correlation with economic growth — indicating that periods of higher inflation in Uzbekistan have tended to coincide with slower growth.

2.4 Monetary policy and inflation

In September 2017, the Central Bank of Uzbekistan (CBU) adopted a new approach in which monetary policy focused on achieving price stability. Restrictions on access to foreign exchange were abolished, unifying the exchange rate and resulting in a roughly 50 percent depreciation. Importantly, the CBU was also removed from the list of state bodies under

government control, opening the way for greater operational independence. This shift was formalized in 2019, when new legislation established price stability explicitly as the primary objective of monetary policy and the CBU began transitioning toward an inflation targeting regime. Operationally, bank reserves initially served as the CBU's target, later replaced by the policy rate and interbank rate. Although the 2017 depreciation triggered a temporary surge in inflation, inflation subsequently trended downward under a consistently tight monetary policy stance, though it remains relatively high and shows inertia (Al Rasasi and Cabezon).

2.5 Synthesis and research gap

Taken together, these studies show that the determinants of inflation are not settled, particularly for transition economies. Binici et al. and Embergenov et al. come to opposite conclusions on trade openness, but their contexts differ substantially: Binici et al. examined OECD countries, where markets are already highly competitive, while Embergenov et al. studied Uzbekistan. This distinction matters, because the competition channel through which openness is expected to reduce inflation may operate weakly in an economy still dominated by state-owned enterprises. The exchange rate is the least contested of the factors, with Poghosyan establishing both the magnitude of pass-through and its causal direction for a sample including Uzbekistan. However, no existing study examines trade openness, GDP growth (economic growth), exchange rate, and monetary policy regime together for Uzbekistan over the 2016-2025 period. Embergenov et al. explicitly note this limitation, stating that no monetary policy variable could be included in their analysis because interest rate data were unavailable. This paper addresses that gap by incorporating the Central Bank of Uzbekistan's policy rate alongside the other factors: GDP growth, trade openness, and exchange rates.

2.6 Hypotheses

Based on the studies reviewed above, this paper expects the following relationships: Hypothesis 1: Trade openness is expected to increase inflation in Uzbekistan, because the country imports more than it exports (Embergenov et al.). Hypothesis 2: A weaker som is expected to increase inflation, because imported goods become more expensive (Poghosyan). Hypothesis 3: The effect of GDP growth on inflation is unclear, because the existing evidence mostly studies the opposite direction, from inflation to growth. Hypothesis 4: A higher policy rate is expected to reduce inflation, but the opposite pattern may appear because the central bank raises rates when inflation is already high.

3 Data and Methodology

3.1 Data

The study covers a decade from 2016 to 2025, giving 10 annual observations. Inflation figures come from the Central Bank of the Republic of Uzbekistan ("Annual Inflation Indicators"). GDP growth (annual percentage change), trade openness (calculated as exports plus

imports divided by GDP, expressed as a percentage) and the exchange rate come from the World Bank's World Development Indicators ("GDP Growth"; "Trade"; "Official Exchange Rate"). The exchange rate is the official rate, some per US dollar, averaged over the year. It is measured as the percentage change from the previous year, calculated from Table 1 as (rate in year t ÷ rate in year t−1 − 1) × 100, so it covers 2017 to 2025 (9 observations). The policy rate comes from the Central Bank of the Republic of Uzbekistan's published policy rate history ("Policy Rate Dynamics"). Because the CBU publishes rates by date range rather than by year, the rate in effect at the end of each year was used. The World Bank publishes trade openness only annually, so annual frequency was used to match all variables.

Table 1: Macroeconomic indicators for Uzbekistan, 2016–2025

Year	Inflation (%)	GDP growth (%)	Trade openness (% of GDP)	Exchange rate (UZS/USD)	Policy rate (%)
2016	9.7	6.7	25.69	2,965.25	9.0
2017	18.8	4.7	41.04	5,113.88	14.0
2018	14.3	5.7	61.40	8,069.61	16.0
2019	15.2	6.8	62.30	8,836.79	16.0
2020	11.1	1.6	53.11	10,054.26	14.0
2021	10.0	8.2	54.75	10,609.46	14.0
2022	12.3	6.1	60.05	11,050.15	15.0
2023	8.8	6.3	63.10	11,734.83	14.0
2024	9.8	6.7	57.53	12,652.29	13.5
2025	7.3	7.7	57.40	12,579.39	14.0

Sources: Central Bank of the Republic of Uzbekistan (inflation, policy rate); World Bank, World Development Indicators (GDP growth, trade openness, exchange rate).

3.2 Method

This study calculates Pearson correlation coefficients, which range from −1 to +1, between inflation and each of the four explanatory variables. Positive means the two move together; negative means they move in opposite directions; zero means no linear relationship. Because the sample contains only 10 observations, a regression with four explanatory variables would not produce reliable estimates, so pairwise correlations are used. P-values come from the t-test for a correlation coefficient, and 95% confidence intervals from the Fisher z-transformation. Poghosyan found through Granger causality testing that exchange rate changes cause inflation rather than the reverse, which supports treating inflation as the outcome being explained.

3.3 Limitations

The sample of 10 observations is small, so results should be read as indicative rather than conclusive. Correlation shows that two variables move together; it does not establish that one causes the other. The exchange-rate result rests heavily on 2017, the year of the devaluation:

without that year the correlation falls from +0.80 to +0.57 and is no longer statistically significant (p = 0.14). The Central Bank sets the policy rate partly in response to inflation, so any observed relationship may reflect the bank reacting to inflation rather than influencing it.

4 Results

Table 2 below reports the correlations between inflation and the four explanatory variables – GDP growth, trade openness, the exchange rate, and the policy rate – for the period 2016 to 2025 (2017 to 2025 for the exchange rate, which is measured as the change from the previous year).

Table 2: Correlation between annual inflation and explanatory variables, Uzbekistan, 2016–2025

Variable	r	n	p-value	95% interval
GDP growth	−0.336	10	0.342	−0.80 to +0.37
Trade openness	−0.088	10	0.809	−0.68 to +0.57
Exchange rate (annual % change)	+0.800	9	0.010	+0.29 to +0.96
Policy rate	+0.390	10	0.265	−0.32 to +0.82

Note: Pearson correlations. With n = 10 a correlation must exceed 0.632 in size to be significant at the 5% level (0.666 for n = 9).

With 10 annual observations a correlation must exceed 0.632 in size to be significant at the 5% level, and GDP growth (−0.34, p = 0.34), trade openness (−0.09, p = 0.81) and the policy rate (+0.39, p = 0.27) do not. The exchange rate measured in levels (−0.51, p = 0.13) is also not significant, and it is misleading because the level rises almost every year. Measured as the year-on-year change in som per US dollar, the exchange rate is positively correlated with inflation (r = +0.80, p = 0.01, n = 9). This depends on 2017, when som per dollar rose 72.5% and inflation rose from 9.7% to 18.8%: without 2017 the correlation is +0.57 (p = 0.14).

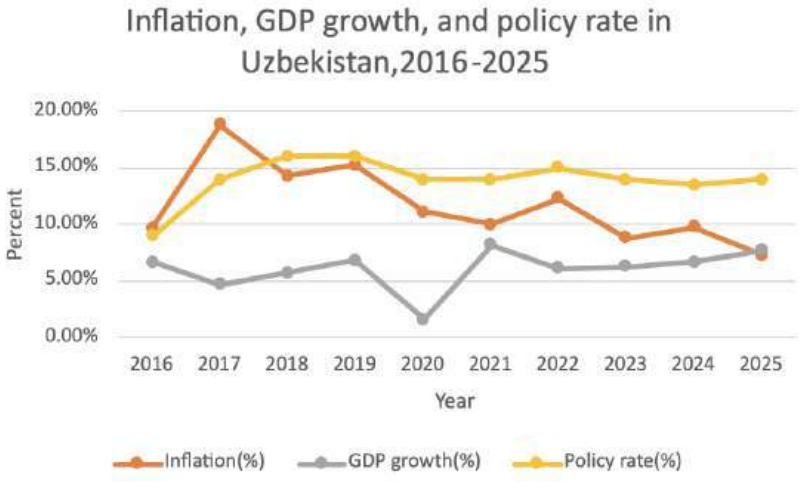

Figure 1: Inflation, GDP growth and policy rate in Uzbekistan, 2016–2025

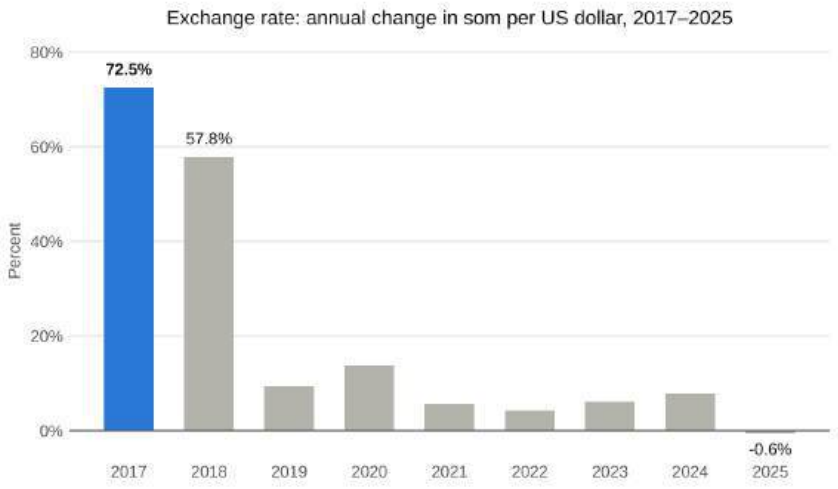

Figure 2: Exchange rate: annual change in som per US dollar, 2017–2025

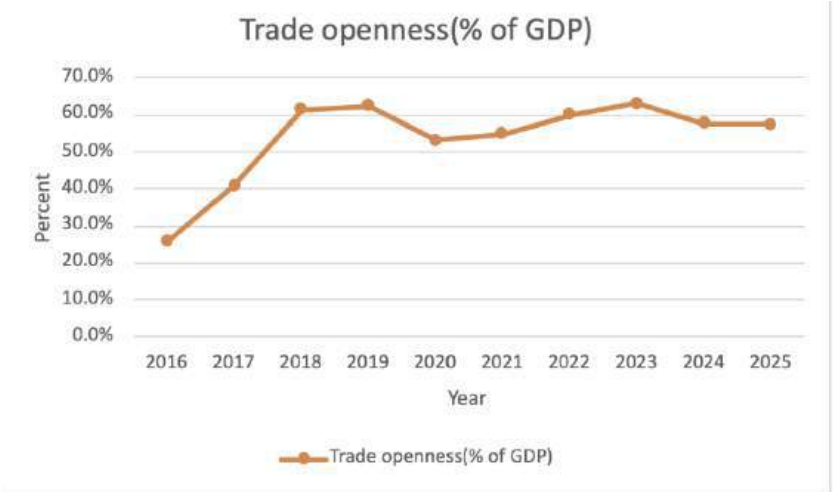

Figure 3: Trade openness (% of GDP), 2016–2025

5 Discussion

Hypothesis 1: Trade openness is associated with higher inflation. This idea is not supported. At −0.09, there is essentially no relationship. Embergenov et al. used quarterly data from 1997-2018 and a more sophisticated method over a different period, so this relationship is not visible in a simple annual correlation.

Hypothesis 2, that a weaker som is associated with higher inflation, is supported in direction: the year-on-year rise in som per dollar correlates +0.80 with inflation ($p = 0.01$, $n = 9$). This agrees with Poghosyan, whose pass-through result is also about changes, not its level. The result rests heavily on 2017; without that year the correlation falls to +0.57 and is no longer significant ($p = 0.14$). The earlier level-based result (−0.51) reflected the som's steady decline while inflation fell after 2018, not a failure of pass-through.

Hypothesis 3 expected the link between GDP growth and inflation to be unclear. The correlation is −0.34 ($p = 0.34$), which is not statistically significant, and it is consistent with

Beknazarov, who finds inflation and growth inversely related. The COVID year 2020 combines the lowest growth (1.6%) with moderate inflation (11.1%) and pulls the correlation toward zero; without 2020 it is −0.74 ($p = 0.02$, $n = 9$). This is a post-hoc check, not a tested hypothesis. Khan and Senhadji study the opposite direction, the effect of inflation on growth, so their threshold cannot by itself explain the sign found here.

Hypothesis 4 stated that a higher policy rate would go with lower inflation, but the result shows +0.39 – positive. The reason is reverse causality: the central bank raises rates because inflation is high, so the two rise together even though the policy is meant to bring inflation down. Figure 1 shows this plainly – the rate sat at 9% until inflation spiked in 2017, then jumped to 14-16% and stayed elevated. As a further check, last year's policy rate correlates −0.51 with this year's inflation ($n = 9$, $p = 0.16$): the sign reverses once the rate is lagged, which fits the central bank reacting to inflation, although the result is not statistically significant.

6 Conclusion

What explains changes in inflation in Uzbekistan, 2016-2025? To test this, annual data for each factor (GDP growth, trade openness, exchange rates, policy rate) were collected, and a correlation analysis was run. Over 2016–2025 the clearest pattern is that inflation moved with the pace of som depreciation ($r = +0.80$ for the year-on-year change, $p = 0.01$), but this result rests largely on the 2017 devaluation. GDP growth (−0.34), trade openness (−0.09) and the policy rate (+0.39) show no statistically significant link to inflation in 10 annual observations. The positive policy-rate correlation most likely reflects the central bank responding to inflation rather than causing it. Correlation is not causation, and ten annual observations are too few for regression. Quarterly data, giving about 40 observations, would allow a regression that separates effects that all changed together in 2017.

Works Cited

Al Rasasi, Moayad, and Ezequiel Cabezon. Uzbekistan's Transition to Inflation Targeting. IMF Working Paper no. 2022/229, International Monetary Fund, 2022, https://doi.org/10.5089/9798400224584.001.

"Annual Inflation Indicators." Central Bank of the Republic of Uzbekistan, 2026, https://cbu.uz/en/monetary-policy/annual-inflation/indicators/.

Beknazarov, B. "Analysis of Uzbekistan's Main Economic Indicators and GDP Growth." Muhandislik va Iqtisodiyot, vol. 4, no. 2, 2026, pp. 7–11, https://doi.org/10.5281/zenodo.18622829.

Binici, Mahir, et al. Trade Openness, Market Competition, and Inflation: Some Sectoral Evidence from OECD Countries. CESifo Working Paper no. 3690, CESifo, 2011.

Embergenov, Bakhitbay, et al. "Analyzing the Impact of Trade Openness on Inflation: A Time Series Study for Uzbekistan." Proceedings of the 6th International Conference on Future Networks and Distributed Systems, Association for Computing Machinery, 2022, pp. 381–86, https://doi.org/10.1145/3584202.3584257.

"GDP Growth (Annual %) – Uzbekistan." World Development Indicators, World Bank, 2026, https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=UZ.

Khan, Mohsin S., and Abdelhak S. Senhadji. "Threshold Effects in the Relationship between Inflation and Growth." IMF Staff Papers, vol. 48, no. 1, 2001, pp. 1–21.

"Official Exchange Rate (LCU per US$, Period Average) – Uzbekistan." World Development Indicators, World Bank, 2026, https://data.worldbank.org/indicator/PA.NUS.FCRF?locations=UZ.

Poghosyan, Tigran. Exchange Rate Pass-Through in the Caucasus and Central Asia. IMF Working Paper no. 2020/154, International Monetary Fund, 2020, https://doi.org/10.5089/9781513552958.001.

"Policy Rate Dynamics." Central Bank of the Republic of Uzbekistan, 2026, https://cbu.uz/en/monetary-policy/refinancing-rate/.

Martin Russell. Uzbekistan Comes in from the Cold: A New Era of Reforms. EPRS Briefing PE 630.344, European Parliamentary Research Service, 2018.

"Trade (% of GDP) – Uzbekistan." World Development Indicators, World Bank, 2026, https://data.worldbank.org/indicator/NE.TRD.GNFS.ZS?locations=UZ.

Consumer Behavior Towards Sustainable Cosmetic Products: A Literature Review

By Ava Tan Li

Abstract

The cosmetic industry is one of the fastest growing industries worldwide, contributing significantly to economic growth while simultaneously raising concerns about its environmental impact. During the last decade this market sector has grown by 3.8% per year on average. At the same time, as sustainability becomes increasingly important in consumer markets, demand for eco-friendly cosmetic products has continued to grow. However, consumers' perspectives towards sustainability are not always homogeneous, and purchasing decisions are often influenced by a variety of factors. Therefore, this narrative literature review investigates the factors that influence consumer behavior toward sustainable cosmetic products, with a specific focus on sustainability perceptions, trust, and willingness to pay as key drivers of purchase intention. Thus, a structured literature review was conducted using fifteen studies published between 2014 and 2025 from regions including Asia, Europe, North America, and Africa. The findings reveal that sustainability generally increases consumer's willingness to purchase cosmetic products and that they will likely pay higher prices for the ones mentioned as environmentally friendly. However, concerns regarding product performance, higher prices, greenwashing, and uncertainty around sustainability claims often reduce consumers' willingness to purchase these products. The review also found that consumer understanding of sustainability varies considerably across individuals and regions, influencing how sustainability claims are interpreted and understood. In markets such as Hungary, for example, consumers support environmentally friendly products in principle, but perceived product effectiveness remains the main barrier: consumers are reluctant to use natural cosmetics that perform less well than conventional alternatives. While in markets like Romania, by contrast, the limiting factor is not indifference but a lack of exposure to sustainability education. These findings suggest that an increase in transparency, consumer education, and the credibility of environmental claims may strengthen consumer trust and encourage greater adoption of sustainable cosmetic products.

Keywords sustainable cosmetics; sustainability perceptions; consumer behavior; purchase intention; willingness to pay; consumer trust; greenwashing; perceived value; attitudes.

Introduction

The cosmetic industry is one of the fastest-growing industries worldwide. During the past decade, this market has grown by over 3.78% a year on average. The expectation for this industry is to exceed a market value of 380 billion USD by the end of 2026 from 267.7 billion USD in 2021 (Alviri et al.). Consequently, this rapid expansion highlights the economic importance of the cosmetics industry and its significant contribution to the global consumer market. Beyond its market value, the cosmetics industry contributes through multiple channels, including direct manufacturing of cosmetic products and indirect supply-chain operations. These effects extend further still: the industry improves individual quality of life through personal care

products, strengthens competitiveness, and contributes to broader economic prosperity in global markets (Turcu and Brancu).

As the industry continues to grow its environmental impact also becomes more important to examine. Therefore, understanding how consumers respond to sustainability in cosmetics is increasingly relevant for both companies and researchers interested in the future of the cosmetic industry. Consequently, researchers on this matter have paid more attention to products described as eco-friendly or labeled as green or sustainable. However, definitions of these concepts vary across studies. For example, Rhee and Lee describe eco-friendly products as "those with ingredient claims that were conveyed visually through a certification logo and text disclosure". By contrast, Joung et al. define them more broadly as products that consumers consider in terms of their environmental characteristics. Beyond differences in wording, definitions of eco-friendly products also vary in scope, ranging from narrow label-based classifications to broader ones that consider the full product life cycle, production processes, and packaging. Some researchers consider certifications or specific ingredient standards (Rhee and Lee), whereas others rely on consumer perceptions (Joung et al.), which definitions are based on consumer perceptions, which may not always reflect the true environmental impact of the product. Additionally, some of these definitions focus on natural or non-toxic components while others incorporate ethical and sustainability considerations, including sourcing practices, waste reduction, and cruelty-free testing (Lin et al.; Amberg and Fogarassy).

Despite the growing body of research on sustainable cosmetics, there are still questions that remain unanswered. One major limitation is that much of the existing research examines consumers' intentions to purchase sustainable products rather than their actual buying behavior. As a result, it is unclear whether people's positive attitudes towards sustainability consistently reflect real purchasing decisions. Furthermore, many studies rely on surveys, self-reported responses, or online experiments, which can introduce bias because participants may exaggerate their environmental concerns or willingness to spend more on sustainable products (Joung et al.; Rhee and Lee). Another challenge is the inconsistent use of terms such as *sustainable, green, natural, organic, and ecofriendly.* Since researchers often define these concepts differently, comparing findings across studies becomes difficult (Lin et al.; Amberg and Fogarassy). In addition, many of the current literature focuses on specific demographic groups—such as women, university students, or Generation Z— or is conducted within a limited number of countries, reducing the extent to which the findings can be generalized to broader populations. Finally, there is relatively little research comparing consumer behavior across different product categories, including skincare versus makeup, luxury versus mass-market brands, and products at varying price points. Addressing these gaps would provide a more comprehensive understanding of how sustainability influences consumer purchasing decisions in the cosmetics industry.

To help synthesize the available literature, this paper examines the key factors that influence consumer behavior toward sustainable cosmetic products. It focuses on variables such as consumer's perceptions of sustainability, purchase intentions, willingness to pay, trust in brands, product pricing, knowledge of sustainability, and perceived product quality and

performance. The review seeks to explain why consumers express interest in sustainable cosmetics while also exploring the obstacles that discourage them from making actual purchases. By synthesizing findings from studies conducted in different countries and using a variety of research methods, this paper highlights both consistent trends and areas where the evidence remains inconclusive. It also distinguishes between purchase intentions and willingness to pay, recognizing that positive attitudes toward sustainability do not always result in corresponding purchase behavior. Ultimately, this review provides a clearer understanding of the current state of research on sustainable cosmetic consumption while identifying important directions for future studies.

This review is complementary to the systematic review by Limbu et al., which synthesizes the determinants of green cosmetics purchase intention and behavior. There are three major differences that define its contribution. First, Limbu et al. synthesize factors associated with both purchase intention and purchase behavior, whereas this review keeps purchase intention and willingness to pay separate, which matters because studies such as Joung et al. and Rhee and Lee show that stated support does not equal unlimited willingness to pay: Joung et al. found a premium of roughly 15-28% using a dichotomous-choice method, and that the premium shrank as the price rose, while Rhee and Lee find that sustainability can lower perceived effectiveness in makeup. Second, whereas Limbu et al. synthesize the literature across countries as one body of evidence, this review separates out the European studies enough to show that the main barriers differ by market. Third, this review turns the findings into a conceptual model rather than a list of factors. Because this review is narrative and covers fifteen sources, it provides a targeted perspective to the existing literature, highlighting the distinction between intention and willingness to pay.

Methods

This study conducts a narrative literature review to investigate the factors that shape consumer behavior toward green and environmentally friendly cosmetic products. The review aims to synthesize findings across existing empirical studies in order to identify consistent patterns, differences, and gaps in the literature. A narrative approach was appropriate because the included studies differed substantially in their definitions of sustainability, study populations, product categories, research design and outcome measures. Fifteen studies were carefully selected for analysis based on their relevance to consumer attitudes, purchase preferences, willingness to pay, and sustainability perceptions within the cosmetics industry. Additionally, this review included studies from countries around the world to ensure regional perspectives on eco-friendly products were considered, and differences in views and attitudes could be differentiated across cultures and market contexts. Lastly, the steps followed in this review are summarized in Figure 1 and explained in more detail in the following paragraphs.

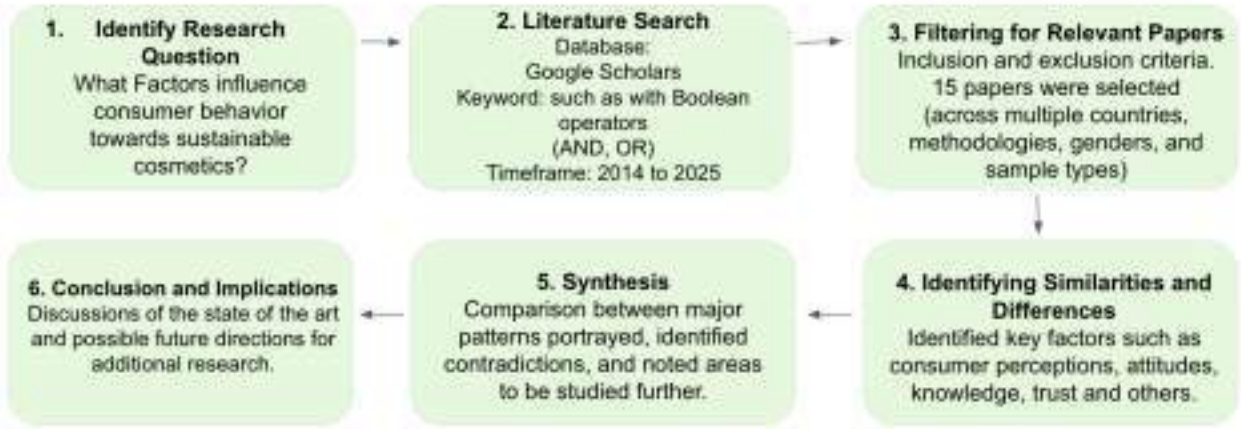

Figure 1: How this literature review was conducted.

(Flowchart illustrating the methodology of the literature review, including the identification of the research question, literature search, paper selection, analysis of similarities and differences, synthesis of findings, and formulation of conclusions and future research implications)

To examine consumer behavior towards sustainable cosmetics across multiple countries, Google Scholar was used as the primary academic database. The search included keywords combinations such as “sustainable cosmetics”, “green cosmetics”, “eco-friendly products”, “consumer demand and sustainable products”, “cosmetics industry”. Boolean Operators such as AND and OR were used to refine research results and ensure relevance. For this review, sources were selected within a time range from 2014 to 2025 with the goal of capturing recent developments and changes in consumer behavior and trends. Studies were included if they examined factors that influence consumer behavior towards green, sustainable, natural, organic or eco-friendly products or if they provided clear results on green consumer behavior that helped explain purchasing decisions in sustainable markets. Studies were excluded if they were opinion or commentary papers, and if their methodology was not clearly stated, or did not focus directly on green or sustainable cosmetic products.

The studies included in this review used a range of research methodologies, including quantitative, qualitative, and mixed-method approaches, which together provide a broader understanding of the research topic. The main methods utilized in the studies are survey-based, commonly using online questionnaires, Likert-scale measurements, and regression analysis to examine relationships between variables. In contrast, the remaining studies apply different approaches, such as focus group discussions or analyses of secondary data. The sample sizes of the studies varied from approximately 142 to more than 500 participants. While some research focuses exclusively on female consumers, others include both male and female participants. However, male consumers were less emphasized than females; only one study focused specifically on male participants, which limits this study’s ability to differentiate between male

and female participants' consumption behaviors.

The literature review covers several geographic regions, including Europe, Asia, North America, and Africa. European studies were conducted in countries such as Finland, Hungary, Croatia, Romania, and the United Kingdom. For North America, one study from the USA was included. In Asia, research includes locations such as Korea, China, Vietnam, and Indonesia (specifically Jakarta and Surabaya). African research is conducted in the Jigjiga town in eastern Ethiopia. In addition, several studies used cross-national datasets, including analyses covering 23 countries, primarily in Asia (Amberg and Fogarassy; Bigovic; Bilal et al.; Chen and Deng; Chin et al.; Joung et al.; Kardos et al.; Limbu et al.; Matić and Puh; Nguyen et al.; Rhee and Lee; Wang et al.). This geographic coverage allows this review to consider global patterns and regional differences, nonetheless direct comparisons across countries are limited since studies used different methods, samples and product categories.

In order to maintain the consistency of the analysis across all studies, each paper was carefully reviewed and organized based on its research goals, research methods, key variables, and main findings. Factors such as environmental awareness, consumer knowledge about products, trust in green claims, sensitivity to price, and perceptions of product effectiveness were specifically emphasized. The results from these studies were then combined and analyzed to identify common patterns, including consumers' willingness to pay for green cosmetics, the impact of sustainability marketing, overall consumer attitudes towards sustainable products, and challenges that limit the adoption of environmentally friendly products.

Results

Consumer Characteristics and External Influences

The existing literature suggests that consumer behavior toward sustainable cosmetics is influenced by multiple factors, including attitudes, price, product knowledge, trust, and perceived product performance. Research such as (Amberg and Fogarassy; Chen and Deng; Wang et al.) has shown that consumers with stronger environmental concerns, positive attitudes toward sustainability, and greater awareness of the personal and environmental benefits of green products are more likely to develop favorable perceptions and higher purchase intentions toward sustainable cosmetics. However, sustainability alone does not determine consumer decisions, as price is often a barrier because high priced sustainable cosmetics may reduce consumers' willingness to purchase these products (Amberg and Fogarassy; Chen and Deng). In addition, knowledge plays a significant role in shaping consumer choices. Greater product knowledge allows consumers to better evaluate environmental claims and build confidence in sustainable brands, thereby increasing trust and purchase intention (Wang et al.). However, higher levels of knowledge may also lead consumers to become more critical of vague, misleading, or exaggerated sustainability claims, increasing skepticism toward green marketing (Nguyen et al.). Furthermore, the distinction between mass-market and luxury cosmetics provides an important perspective when examining sustainable consumption. Mass-market cosmetics are typically

associated with accessibility, affordability, benefit of the product, and frequent purchasing, meaning consumers may prioritize price and functionality when making decisions. In contrast, luxury cosmetics often rely more heavily on brand reputation, exclusivity, quality, and perceived performance. Although existing studies have explored factors affecting sustainable cosmetic consumption, limited research has directly compared how these factors differ between sustainable mass-market and luxury cosmetic products, suggesting a need for further investigation into how product category influences consumer decision-making.

Purchase Intention versus Willingness to Pay

Consumer behavior plays a critical role in determining the success of sustainable cosmetic products because increased environmental awareness does not necessarily result in sustainable purchasing decisions. While many consumers express positive attitudes towards environmentally friendly products, their final purchasing choices are often influenced by factors such as price, perceived product quality, effectiveness, and confidence in sustainability claims. As illustrated in Figure 2, these and other factors interact to shape consumers' perceptions of sustainable cosmetics, ultimately influencing their willingness to pay and purchase intentions. Therefore, understanding consumer behavior is important for identifying the factors that encourage or discourage the adoption of green cosmetics. Two concepts frequently used to measure consumer behavior are purchase intention and willingness to pay (Joung et al.). Purchase intention describes the likelihood that a consumer will purchase a product, whereas willingness to pay refers to the amount of money a consumer is prepared to spend on that product (Matić and Puh). Although these concepts are related, they capture different aspects of consumer decision-making. For example, a consumer may have a strong intention to purchase a sustainable cosmetic product but may be unwilling to pay the higher price along with it, highlighting the difference between purchase intention and willingness to pay (Amberg and Fogarassy; Joung et al.;Wang et al.). Likewise, consumers may express a willingness to pay more for environmentally friendly products, although the amount they are willing to spend depends on factors such as perceived product value, quality, and effectiveness (Amberg and Fogarassy; Rhee and Lee; Joung et al.). Therefore, examining both purchase intention and willingness to pay provides a more comprehensive understanding of the factors influencing consumer behavior toward sustainable cosmetic products.

The findings from this literature review indicate that several factors shape how consumers behave toward green and sustainable cosmetics. Across the reviewed studies, a general pattern appeared: sustainability generally tends to increase consumers' willingness to purchase green products. However, this relationship was not always consistent. Consumer responses were often shaped by factors such as price, purchasing power, perceived product performance, marketing, knowledge and social influence and level of trust they have in the brand and their sustainable claims (Joung et al.; Matić and Puh). These findings suggest that while sustainability is a strong driver of a consumer interest in green products, practical and psychological factors continue to play a significant role in influencing whether consumers intend

to purchase sustainable cosmetic products or pay more for them. Figure 2 summarizes the main relationships identified in the literature, showing how consumer characteristics and external influences shape perceptions of trust, sustainability, value, and quality, which then influence purchase intention and willingness to pay.

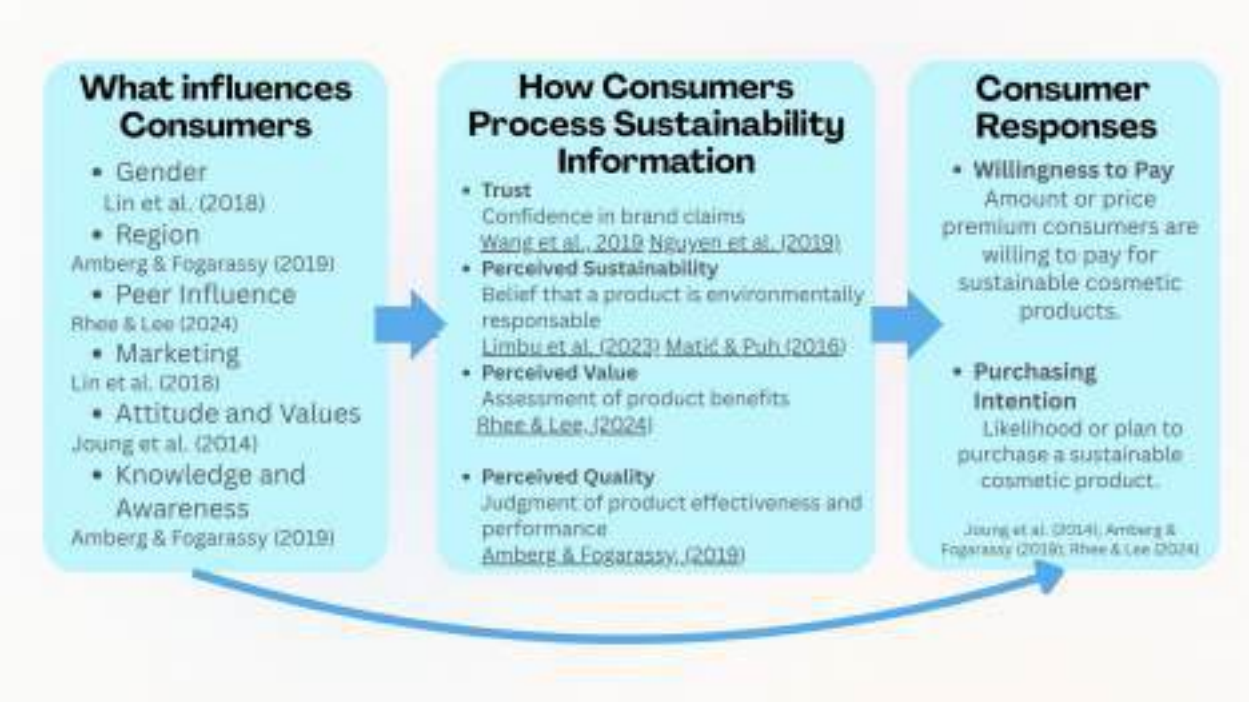

Figure 2: Conceptual Framework of Consumer Behavior Toward Sustainable Cosmetic Products

This figure summarizes the main factors identified in this literature review. influencing consumer behavior toward sustainable cosmetic products. Consumer characteristics and external influences, such as gender, geographic location, peer influence, marketing, attitudes, values, knowledge, and awareness, shape how consumers process sustainability information. These factors influence trust, perceived sustainability, perceived value, and perceived quality, which then affect purchase intention and willingness to pay for sustainable cosmetic products.

Trust, Perceived Sustainability, and Product Quality

Limbu et al. define sustainability as attributing to the concept of green cosmetics and involves two essential issues, which are the ingredients and the packaging of the products. Furthermore, Limbu et al. listed specific examples of sustainability, such as non-chemical organic solutions, renewable sources, and ingredients processed by environmentally supported packaging. Additionally, a study conducted by Nguyen et al. identified key drivers of purchase intention, specifically factors like consumers' cognitive knowledge and greenwashing. Nguyen et al. mentioned that their study in Vietnam discovered that cognitive knowledge is affected by the overuse of greenwashing actions, which caused increased suspicion and skepticism among consumers of the validity of green claims. Hence, it impacts consumers' purchase intention towards green products. Additionally, the study found that greenwash has a negative effect on green purchase intentions due to the negative impact on consumer perceptions, creating an

inverse effect on consumer's willingness to purchase, and possibly creating damaging influence on brands. Granted, the study acknowledges that consumers may be misled because of inability to distinguish truth in some occasions. However, consumers do not always share the same understanding of sustainability. Matić and Puh noted that many consumers struggle to distinguish which products are genuinely natural or organic because numerous brands make similar environmental claims, leading to confusion and distrust in sustainability labels. Furthermore, Matić and Puh mentioned that green products may be interpreted differently by consumers, with some associating sustainability primarily with health benefits, others with environmental protection, and others with superior product quality. Therefore, consumers' perceptions and understanding of sustainability are critical factors influencing purchase intentions, as confusion, skepticism, and differing interpretations of green claims can affect their willingness to buy sustainable cosmetic products.

Consumer confidence in green cosmetic products is strongly influenced by both trust and awareness. Research by Wang et al. suggests that individuals who better understand sustainable products are often more willing to purchase them because they feel their choices can contribute positively to environmental protection. Wang et al. explained that knowledge about green products strengthens consumers' trust in environmentally responsible brands, which in turn increases purchase intention. At the same time, Nguyen et al. found that consumers with more knowledge about sustainability may also become more skeptical of environmental claims, especially when companies use unclear or exaggerated information. Similarly, Matić and Puh noted that there is increasing distrust in "organic" and "natural" labels because many products make similar environmental claims, making it difficult for consumers to determine which products are genuinely sustainable and which are not. These findings emphasize that honesty, transparency, and credible sustainability claims are essential for maintaining consumer trust in the green cosmetics industry.

Marketing strategies also strongly influence consumer behavior toward green cosmetics, especially through social media and peer influence. Bigovic suggests that consumers, particularly Generation Z, are more likely to trust peer reviews and online recommendations than traditional ways of paid advertisements. Bigovic also found that transparency is a key factor shaping consumer's purchasing decisions, as consumers prefer brands that communicate their mission and sustainability efforts more honestly. In addition, Kardos et al. show that green marketing initiatives can increase environmental awareness and encourage more responsible consumer behavior. Furthermore, Wang et al. argued that marketing strategies can influence green purchase intentions by increasing consumers' knowledge of green products through multiple media channels. Their study also found that greater product knowledge strengthens consumer trust and confidence in environmental claims, making consumers more likely to purchase sustainable products. These findings highlight the importance of credibility and transparency in effectively promoting sustainable cosmetic products.

A consistent finding across the literature is that many consumers show willingness to support environmentally friendly products, even if it means purchasing at a higher price level.

For example, Joung et al. found a premium of roughly 15-28% using a dichotomous-choice method, and that the premium shrank as the price rose. Likewise, Amberg and Fogarassy found that many consumers are willing to spend more on cosmetics made with natural ingredients and sustainable packaging, even when the perceived product effectiveness is lower. Furthermore, Rhee and Lee found that consumers who highly prioritize appearance-enhancing attributes such as coverage and pigmentation were less willing to purchase eco-friendly makeup products when sustainability claims led them to perceive the products as less effective. These results indicate that sustainability can increase the perceived value of cosmetic products, although the level of added value differs depending on the product's category, price, and expected performance.

Barriers to Sustainable Consumption

Although many studies suggest that sustainability can encourage consumers to purchase green cosmetic products, there are still existing barriers when it comes to making purchasing decisions. One recurring challenge is the higher cost often associated with sustainable products, as Wang et al. mentioned, which can discourage consumers despite their positive attitude towards environmental protection. Additionally, Rhee and Lee found that another barrier is eco-friendly ingredient claims negatively affect consumers' perceptions of appearance-enhancing qualities in makeup products, reducing consumers' willingness to purchase them. Furthermore, uncertainty surrounding environmental claims presents another obstacle. Matić and Puh reported that consumers often find it difficult to determine which products are truly natural or organic because many brands use similar sustainability claims, leading to confusion and distrust. Together, these findings demonstrate that concerns about price, appearance enhancing performance, and credibility can possibly limit the growth of sustainable cosmetic consumption.

Geographic location can also shape consumers' responses to sustainable cosmetic products because environmental awareness, market development, and trust in sustainability claims differ across regions. For example, Nguyen et al. found that Vietnamese consumers were particularly influenced by concerns about greenwashing, as uncertainty surrounding environmental claims reduced confidence in green products. Furthermore, Chen and Deng found that consumers in China with more positive attitudes toward environmentally friendly products were more likely to intend to purchase them. The study also highlighted the importance of subjective norms, suggesting that recommendations and expectations from friends, family, and society can influence consumers' decisions to purchase sustainable products. Unlike studies conducted in Asian markets that emphasize environmental concern and social influences, Amberg and Fogarassy found that consumers in Hungary generally supported environmentally friendly products and reported willingness to pay more for them, but perceived product effectiveness remained a barrier. Additionally, research conducted by Kardos et al. found that although Romanian consumers showed interest in green products, limited environmental knowledge and awareness remains a challenge in markets where sustainability culture is still developing. Therefore, these findings suggest that the factors influencing sustainable cosmetic purchasing behavior vary across regions, with consumer decisions being shaped by different

combinations of environmental awareness, social influences, trust in sustainability claims, and practical considerations such as price and product quality.

Gender appears to play a significant role in consumers' willingness to purchase sustainable beauty products. Matić and Puh reported an odds ratio of 2.262, indicating that women had higher odds of purchasing natural cosmetics than men. Matić and Puh also mentioned that according to Banerjee and McKeage, female consumers tend to be more ecologically conscious than male consumers, and similarly, Mainieri and Pillai suggested that women are more likely to buy a green product because they believe the product is better for the environment. These findings show consistent patterns that women are more concerned about green issues than men and are more likely to purchase environmentally friendly products because of their belief that these products are more beneficial for the environment. However, men are underrepresented across many of the studies reviewed, highlighting a gap in the literature and the need for further research in this area.

Discussion

Overall, the literature suggests that consumers generally express positive attitudes toward cosmetics (Rhee and Lee; Amberg and Fogarassy; Lin et al.), although these positive attitudes do not always translate into actual purchasing behavior (Joung et al.). Factors such as trust in sustainability claims, perceived product quality, environmental awareness, and product knowledge encourage consumers to consider sustainable cosmetics. In contrast, high prices, concerns about greenwashing, and uncertainty about product performance are barriers when consumers make purchasing decisions.

The relationship between price, trust, and performance is particularly important. Consumers may be willing to support sustainable cosmetics, but this willingness can decrease when environmentally friendly products are perceived as significantly more expensive or less effective than conventional alternatives (Rhee and Lee). Trust can therefore influence how consumers evaluate both price and performance, such that unclear or exaggerated claims may increase skepticism (Amberg and Fogarassy). At the same time, sustainability alone may not be sufficient if consumers doubt a product's effectiveness. Companies should therefore communicate sustainability transparently and honestly while demonstrating that sustainable products can provide comparable quality and performance. Some individual studies emphasize different factors influencing sustainable consumption. For example, Lin et al. highlights the importance of product knowledge, while other studies emphasize trust in sustainability claims or perceived product quality. Rhee and Lee provide evidence concerning price, whereas Joung et al. demonstrates the gap between positive attitudes and actual purchasing behavior. These differences may result from variations in study samples, cosmetic categories, and research settings, as well as in how sustainability, purchase intention, trust, and consumer behavior are defined. These findings suggest several implications for companies. Cosmetic companies can learn that sustainability claims are most effective when it is being transparently presented with credible evidence and clear demonstrations of product quality. However, the available evidence

does not justify assuming that positive consumer attitudes will automatically result in increased sales. Companies should therefore consider sustainability alongside price, product performance, and consumer trust rather than treating environmental sustainability as a selling point. Regulators and certification organizations can also learn from the importance of trust and skepticism toward sustainability claims. More consistent sustainability definitions, transparent information, and credible certification systems could help consumers distinguish between trustworthy environmental claims and misleading ones. However, the literature does not provide sufficient evidence to determine which specific certification systems, regulations, or communication strategies would be most effective across the entire cosmetics industry.

Limitations

Despite these insights, this review has several limitations. First, because this is a narrative review rather than a systematic one, it did not use a pre-registered protocol, formal risk-of-bias checks, or meta-analysis. As a result, the synthesis is interpretive and the weighting of studies reflects the author's judgement. Second, it relied on existing published research rather than primary data, so its conclusions depend on the quality and scope of previous studies. Third, a single database (Google Scholar) was searched, which likely underrepresents non-English work and makes selection bias difficult to rule out. Fourth, the geographic coverage is uneven: Africa and North America are each represented by a single study, Latin America and Australia by none. Fifth, the conceptual model in Figure 2 is an organizing device derived from the literature rather than a tested causal model.

Conclusion and Future Direction

Future research should investigate actual purchasing behavior to better understand the gap between consumer attitudes and actions. Researchers should also examine more diverse populations and compare consumer behavior across different cosmetic categories, such as skincare and makeup, as well as across various price ranges and market segments. Additionally, establishing more consistent definitions of sustainability within the cosmetics industry would also improve the comparability of future research. By addressing these limitations, future studies can provide a more complete understanding of the factors that drive sustainable cosmetic consumption and help businesses develop more effective sustainability strategies. Ultimately, understanding how consumers process sustainable attitudes into purchasing decisions will be important for closing the gap between intention and action and supporting the growth of the sustainable cosmetics industry.

Works Cited

Alviri, Hajar, et al. “Beyond Green Chemistry: A Comprehensive Review of How Sustainability Has Been Integrated into Cosmetic Research.” Global Sustainability, vol. 8, 2025, p. e16, https://doi.org/10.1017/sus.2025.5.

Amberg, Nora, and Csaba Fogarassy. “Green Consumer Behavior in the Cosmetics Market.” Resources, vol. 8, no. 3, 2019, p. 137, https://doi.org/10.3390/resources8030137. Bigovic, Olivia. “The Impact of Sustainable Social Media Marketing on Generation Z’s Buying Behavior in the Cosmetics Industry.” 2025. Haaga-Helia University of Applied Sciences, bachelor’s thesis. Theseus, https://www.theseus.fi/handle/10024/893105.

Bilal, Arebu I., et al. “Cosmetics Utilization Practice in Jigjiga Town, Eastern Ethiopia: A Community Based Cross-Sectional Study.” Cosmetics, vol. 3, no. 4, Dec. 2016, p. 40, https://doi.org/10.3390/cosmetics3040040.

Chen, Kai, and Ting Deng. “Research on the Green Purchase Intentions from the Perspective of Product Knowledge.” Sustainability, vol. 8, no. 9, Sept. 2016, p. 943, https://doi.org/10.3390/su8090943.

Chin, Jacky, et al. “The Investigation of Consumers’ Behavior Intention in Using Green Skincare Products: A Pro-Environmental Behavior Model Approach.” Sustainability, vol. 10, no. 11, Oct. 2018, p. 3922, https://doi.org/10.3390/su10113922.

Joung, Soon Hee, et al. “Willingness to Pay for Eco-Friendly Products.” Asia Marketing Journal, vol. 15, no. 4, Jan. 2014, https://doi.org/10.53728/2765-6500.1565.

Kardos, Mihaela, et al. “Green Marketing’s Roles in Sustainability and Ecopreneurship. Case Study: Green Packaging’s Impact on Romanian Young Consumers’ Environmental Responsibility.” Sustainability, vol. 11, no. 3, Feb. 2019, p. 873, https://doi.org/10.3390/su11030873.

Limbu, Yam B., et al. “What Influences Green Cosmetics Purchase Intention and Behavior? A Systematic Review and Future Research Agenda.” Sustainability, vol. 15, no. 15, Aug. 2023, p.11881, https://doi.org/10.3390/su151511881.

Lin, Yifeng, et al. “An Exploratory Study of Consumer Attitudes toward Green Cosmetics in the UK Market.” Administrative Sciences, vol. 8, no. 4, Dec. 2018, p. 71, https://doi.org/10.3390/admsci8040071.

Matić, Matea, and Barbara Puh. “Consumers’ Purchase Intentions Towards Natural Cosmetics.” Ekonomski Vjesnik/Econviews - Review of Contemporary Business, Entrepreneurship and Economic Issues, vol. 29, no. 1, June 2016, pp. 53–64.

Nguyen, Thi Thu Huong, et al. “Greenwash and Green Purchase Intention: The Mediating Role of Green Skepticism.” Sustainability, vol. 11, no. 9, Jan. 2019, p. 2653, https://doi.org/10.3390/su11092653.

Rhee, Hyunjeong, and Kyu-Hye Lee. “What Going ‘Green’ Costs for Cosmetics: The Effect of Sustainability Claims on Colored Makeup Products.” International Textile and Apparel Association Annual Conference Proceedings, vol. 80, no. 1, 2024.

Turcu, Oana, and Cristian Brancu. “The Impact of the Beauty Industry on the Economic Growth

at the European Union Level." Proceedings of the International Conference on Business Excellence, vol. 17, no. 1, 2023, pp. 468–74, https://doi.org/10.2478/picbe-2023-0045.
Wang, Hong, et al. "How Does Green Product Knowledge Effectively Promote Green Purchase Intention?" Sustainability, vol. 11, no. 4, Jan. 2019, p. 1193, https://doi.org/10.3390/su11041193.

Evaluating The Longitudinal Structural Decay Of Composite Hockey Sticks

By Ethan Zhu

Abstract

Modern composite hockey sticks store and release elastic energy through engineered kickpoint zones, but the rate at which this performance decays with use is largely uncharacterized. This study measured natural frequency (f_n) and static deflection (Δx) across 14 stick combinations from CCM, Bauer, and Warrior. Of those 14, three CCM Ribcor Trigger 10 Pros were tested, new and after 1 and 3 months of on-ice use, with all frequencies measured using Phyphox. A one-way ANOVA found a statistically significant difference in f_n across the 14 groups, and post-hoc comparisons resolved which stick pairs differed. Ten of 14 groups showed an increase in f_n after static loading, and the Ribcor Trigger 10 Pro's f_n rose monotonically across the three usage conditions (11.26 to 12.22 Hz). Kickpoint category was associated with the largest and most consistent difference in f_n, while price-tier comparisons were confounded with kickpoint and could not be isolated. These results show that, at least over the timescales and loading conditions tested here, fixture seating and possible resin post-cure effects outweigh microcrack-driven softening. Flagship pricing is not matched by measurable durability gains within this stick matrix.

KeyWords

Engineering, Materials Science, Composite Materials, Fatigue Testing, Sports Engineering, Structural Health Monitoring

Introduction

Over the past two decades, the ice hockey equipment industry has transitioned from wood and aluminum shafts to advanced carbon composite constructions. Engineered through complex resin matrices and specialized carbon weaves, high-end composite sticks store and release elastic potential energy during a shot, creating what players colloquially term pop. Custom nanotoughened resin systems and precisely tuned fiber architectures now allow manufacturers to place dual flex zones at specific points along the shaft,[1] giving rise to the marketed kickpoint categories that define modern stick lines. However, this mechanical efficiency comes at a steep financial cost, with top-tier flagship models regularly exceeding $370.

A critical issue in modern stick mechanics is functional fatigue: the progressive breakdown of internal resin binders and the propagation of micro-fractures within the carbon fiber matrix. While structural failure is obvious, functional fatigue causes a subtle yet significant loss in flexural rigidity, torsional stiffness, and energy retention long before the stick visibly breaks.[2–3] This mechanical degradation varies depending on key structural design variables: Kickpoint, Flex Rating, and Manufacturing & Cost tier.

Low-kickpoint tapers (e.g., CCM Ribcor Trigger 9 Pro, Warrior Covert QR7 Pro, Bauer Twitch) or Ultra Low Kickpoint tapers (e.g. CCM Ribcor Trigger 10 Pro, CCM Vizion 2025) utilize narrow, highly localized flex zones near the blade for rapid shot release, whereas mid kick (e.g.CCM Vizion 2026, CCM Tacks XF Ghost, Bauer Tracer) and hybrid profiles (e.g. CCM JetSpeed FT8 Pro) distribute bending strain across a broader section of the shaft.

The relative stiffness of the shaft dictates the maximum flexural strain the composite fibers undergo during loading cycles. Flagship models from legacy brands utilize specialized resin layering to minimize overall mass, which claims to improve year after year, and have lower tier sticks priced significantly lower. Current literature focuses primarily on initial puck release velocity rather than long term mechanical degradation rates under repeated stress. The goal of this research is to quantify the rate of material decay across different flex ratings, kickpoints, and material composition to establish an empirical cost-to-effectiveness benchmark for players.

Composite Fatigue Mechanisms in Fiber-Reinforced Polymers

Fatigue in carbon fiber-reinforced polymer (CFRP) composites is a well documented, multi-stage process governed by the interaction of fiber, matrix, and interface behavior under repeated cyclic loading.[2] Damage initiates as matrix microcracking and fiber-matrix debonding, known as inter-fiber fracture, before progressing to delamination between plies as interlaminar shear stresses accumulate at free edges and ply interfaces.[3] Ju and Morgan demonstrated that microcrack density in bismaleimide-carbon fiber composites increases measurably under combined mechanical and thermal cycling, driven by fiber-matrix thermal expansion mismatch,[4] while comparative studies of thermoset versus thermoplastic matrices show that tougher, more ductile matrices can paradoxically exhibit faster fatigue degradation due to more widely distributed damage accumulation.[5] At the fiber-matrix interface itself, microcracks that form under tribological and cyclic loads serve as initiation sites that propagate into macroscopic interlaminar failure.[6]

Residual stiffness and residual strength models are two dominant frameworks used to characterize this degradation quantitatively. Stiffness degradation is typically modeled as a three-stage process: an initial sharp decline associated with early matrix cracking, a prolonged and comparatively linear middle stage of gradual stiffness loss, and a final rapid decline preceding failure.[7–8] This linear middle stage represents the bulk of a composite's operational lifespan, during which micro-flaws slowly and steadily accumulate throughout the resin without causing immediate, catastrophic structural failure. This three-stage curve has been experimentally confirmed in CFRP tendons under cyclic loading, where residual modulus decreases with a characteristic degradation rate tied to the accumulated cycle count,[9] and in matrix-dominated laminates where bidirectional fiber layups show measurably different degradation trajectories than unidirectional layups.[10] Because different fiber architectures dictate how stresses are transferred around these growing microcracks, manufacturers must carefully balance layup geometry to delay the material's transition into the final, rapid-decline stage.

Statistical treatments of stiffness loss further show that the modulus degradation rate

correlates directly with fatigue life across varying microstructures, providing a basis for using stiffness or its dynamic proxy, natural frequency, as a non-destructive proxy for remaining fatigue life.[11] For athletes using composite equipment, this proxy is highly practical; internal structural degradation often manifests as a "dead" feeling in the hockey stick long before any macroscopic surface cracking becomes visible to the naked eye. Digital image correlation studies of glass-fiber-reinforced polymers further confirm that flexural modulus reduction tracks closely with the physical accumulation of matrix damage visible under repeated three-point bend testing.[12] Recent multi-scale reviews emphasize that environmental characterization of long-term composite durability remains time- and cost-intensive, motivating accelerated or indirect testing methodologies such as the one adopted in this study,[13–14] and very high cycle fatigue work on CFRP under ultrasonic loading illustrates that composite fatigue behavior continues to evolve even beyond conventional fatigue life regimes, underscoring the need for sensitive, low-cost detection methods applicable to consumer-grade composites.[15] By establishing a reliable link between vibrational frequency drops and internal fatigue, researchers can better evaluate the true functional lifespan of high-stress sporting goods without relying on expensive, destructive laboratory equipment.

Vibration-Based Methods for Damage Detection

Because internal micro-cracking reduces local stiffness without necessarily producing visible surface damage, vibration based methods have become a preferred non-destructive approach for detecting fatigue damage in composite beams and plates. The underlying principle is that a structure's natural frequencies are a direct function of its mass and stiffness distribution, so a loss of stiffness from microcracking or delamination manifests as a measurable decline in natural frequency.[16–17] Experimental and numerical studies on composite cantilever beams have shown that increasing crack depth and damage severity produce progressively larger reductions in the first several natural frequencies, and that these reductions can be used to estimate not only the presence but also the approximate location and severity of damage.[18–19] Comparable work on hybrid composite cantilever beams combining finite element analysis with artificial neural networks confirms that mode shifting and modal shape mutation accompany the frequency decline, offering additional diagnostic signals beyond the fundamental frequency alone.[20–21]

A recent comprehensive review of vibration-based damage detection methods across beams, plates, and composite structural members confirms that natural-frequency monitoring, while occasionally limited in sensitivity to very small or multiple simultaneous cracks, remains one of the most practical and experimentally accessible damage indicators available outside a specialized laboratory setting,[22–23] and numerical experimental comparisons on cracked composite beams report close agreement (within a few percent) between theoretical and measured frequency shifts, supporting the validity of simplified free vibration cantilever testing as a fatigue monitoring tool.[24] This body of work directly motivates the present study’s use of free vibration cantilever testing, in which a smartphone accelerometer captures the acceleration response and a Fast Fourier Transform isolates the dominant natural frequency; the Phyphox platform used for

data acquisition has been independently validated for exactly this class of accelerometer-based physics measurement.[25]

Hockey Stick Biomechanics and Composite Sports Equipment

Existing biomechanical literature on ice hockey sticks has focused overwhelmingly on the relationship between shaft stiffness, energy storage, and puck velocity at the moment of a new stick's use, rather than on how these properties change with accumulated wear. Early comparative testing across shaft stiffness levels found no statistically significant difference in flex strength between manufactured stiffness classes, attributing shot velocity primarily to player strength and technique rather than to the composite construction itself.[26] However, this conclusion fundamentally treats the hockey stick as a static variable, ignoring the reality that a player's equipment mechanically degrades over a grueling season. If technique heavily outweighs material choice, elite players may unknowingly alter their shooting mechanics to compensate for a softening, "dead" stick, masking the underlying material fatigue from traditional observation. Subsequent biomechanical studies quantified how shaft stiffness and player-applied force jointly determine potential energy storage and puck speed,[27] and a broader literature review concluded that player technique and strength exert substantially more influence over shot puck velocity than stick composition, with stick flexibility, rather than material choice, identified as the key mechanical variable.[28]

To remove the confounding variable of human biomechanics and isolate the material itself, mechanical cantilever bend testing across sticks from multiple manufacturers has since become the standard method for characterizing shaft stiffness independent of player variability.[29] This shift toward isolated mechanical testing is crucial; to truly understand a stick's longevity, researchers must separate the composite's physical decay from the player's physical compensation. Building on this, more recent work using embedded internal measurement units and machine learning has begun to quantify how kickpoint and flex profile interact with individual shooting mechanics in real time.[30] Direct measurement of shaft recoil during a slap shot has further shown that the timing and magnitude of shaft deflection recovery is itself a measurable, repeatable mechanical signature of a given stick's construction.[31] Yet, all of these sophisticated metrics are typically applied only to pristine, factory-new models. The present study leverages these conditions, specifically cantilever deflection and recoil recovery measured with natural frequency, but pivots the focus longitudinally. By tracking these proven mechanical signatures across months of game-worn equipment, this research addresses a glaring blind spot in sports engineering: quantifying exactly how quickly a precisely engineered carbon matrix loses its advertised elasticity.

Because no equivalent fatigue focused literature yet exists for hockey sticks specifically, this study draws on the more mature composite fatigue literature from an analogous sporting goods application: a golf club. Golf shafts share the essential engineering problem of a tapered, carbon fiber composite cantilevered structure engineered for a target flex profile and repeated high strain loading. Quasi-static and dynamic testing of golf shaft composites has shown that

damping and stiffness characteristics are governed primarily by fiber orientation, strain, and strain rate, and that low strain rate cantilever testing is generally sufficient to characterize a shaft's in service dynamic performance.[32–33] Comparable experimental and numerical characterization of golf shaft mechanical performance further demonstrates that composite sporting goods can be rigorously benchmarked using the same flexural and vibrational methods developed for aerospace-grade composites,[34] and broader reviews of composite materials in sporting and exercise equipment confirm that fatigue resistance and vibration damping are now treated as first-order design considerations across the sporting goods industry, not merely as secondary effects of a stiffness-optimized layup.[35–36] Taken together, this literature establishes both the mechanical vocabulary (kickpoint, flex, damping) that differentiates hockey stick marketing categories and the composite fatigue vocabulary (residual stiffness, natural frequency decline, microcrack propagation) needed to evaluate how these categories perform under sustained use, the gap this study addresses.

As the preceding review demonstrates, current literature focuses primarily on initial puck release velocity and single use mechanical characterization rather than on long term mechanical degradation rates under repeated stress. No published study to date has directly compared composite hockey stick fatigue rates across kickpoint and flex categories, nor benchmarked a direct-to-consumer manufacturer against legacy flagship brands on a cost-to-durability basis. The goal of this research is to quantify the rate of material decay across different flex ratings, kickpoints, and material composition to establish an empirical cost-to-effectiveness benchmark for players.

Hypotheses

Based on the composite fatigue and vibration-detection literature reviewed above, three directional hypotheses are proposed and tested in this study:

Research Question 1: How does the engineered kickpoint geometry of a composite hockey stick affect its rate of structural degradation over an accumulated period of active use?

Hypothesis 1: Ultra-low and low-kickpoint sticks, which concentrate bending strain in a narrower zone near the blade, will show a significantly larger decline in stiffness (Δx) and natural frequency (f_n) between the New and 3-month-use conditions than mid- or hybrid-kickpoint sticks, which distribute strain across a broader shaft section.

Research Question 2: Within a given kickpoint, how does a stick's flex rating or shaft stiffness influence its proportional rate of material fatigue under standardized loading?

Hypothesis 2: Within a given kickpoint category, lower-flex sticks will show a smaller proportional decline in f_n over the same usage period than higher-flex (stiffer) sticks, consistent with greater strain energy distribution and reduced localized stress concentration at lower flex ratings.

Research Question 3: Does the premium cost of current-generation flagship composite sticks correlate proportionally with measurable advantages in baseline stiffness and fatigue resistance when compared to their mid-tier or previous-generation counterparts?

Hypothesis 3: Current-generation flagship models (e.g., Ribcor Trigger 10 Pro, Vizion 2026) will show measurably higher baseline stiffness and/or slower initial-condition frequency decline than their manufacturer's own previous-generation or mid-tier counterparts (e.g., Ribcor Trigger 9 Pro, Vizion 2025) tested new. However, the magnitude of that advantage will be disproportionately smaller than the percentage difference in retail price.

Figure 1 visually represents this hypothesis and represents a high frequency retention for sticks that have a higher base price.

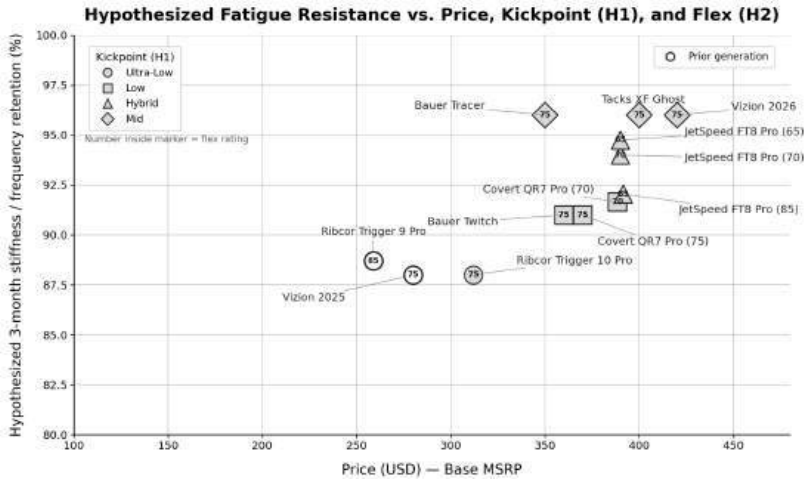

Figure 1: Hypothesized frequency retention by price graph. This graph was created using Python and Matplotlib.

Methods

The experiment utilizes a controlled data set of composite sticks across brand profiles, flex ratings, and wear states (Table 1):

Brand/Manufacturer	Model Profile	Kickpoint	Tested Flex ratings	Condition State
CCM	Ribcor Trigger 10 Pro	Ultra-Low	70	New, 1 and 3 months of heavy use
CCM	Ribcor Trigger 9 Pro	Low	85	New
CCM	JetSpeed FT8 Pro	Hybrid	65, 70, 85	New
CCM	Vizion 2025	Ultra Low	70	New
CCM	Vizion 2026	Mid	70	New
CCM	Tacks XF Ghost	Mid	70	New
Bauer	Tracer	Mid	70	New
Bauer	Twitch	Low	70	New
Warrior	Covert QR7 Pro	Low	70, 75	New

Table 1: Overview of composite stick experimental matrix and structural properties.

Composite Beam Theory and Standardized Flexural Testing

Both static and dynamic testing in this study rest on classical Euler-Bernoulli cantilever beam theory, which relates applied load, beam geometry, and Flexural Rigidity to measure deflection and natural frequency. Standardized single and double cantilever beam configurations are already used industrially to characterize the flexural and fracture properties of sandwich and laminate composites.[37–38] Comparative work in cantilever beam versus mid-point bending test configurations for thin composite panels confirms that a cantilever configuration, following an ASTM-consistent span-to-thickness ratio, produces bending stiffness estimates comparable to standard mid-point flexural testing while requiring simpler fixturing.[39] This precedent supports the use of simplified cantilever and three-point deflection rigs, rather than a full universal testing machine, as an appropriate and sufficient method.

Static Three-Point Bend Testing

To assess static Flexural rigidity (E_I) and permanent structural deformation, each stick is supported horizontally on dual beam supports separated by a fixed span distance (L). A calibrated mass (m) is applied at the geometric midpoint of the shaft. Downward displacement (Δx) under load is measured using precision dial calipers to calculate bending stiffness via elastic beam deflection equations: $\Delta x = \frac{FL^3}{48E_I}$

Following load removal, a permanent set is recorded to evaluate plastic matrix deformation.

For each set of trials: a fixed span of $L = 30$cm centered on the shaft taper region, a calibrated test mass of $m = 5kg$ applied at midspan, and dial caliper displacement readings to $\pm\ 0.01mm$ precision is used. Each stick combination is tested across $n = 5$ trials, with the shaft rotated 90º for one trial to average out any blade-face asymmetry, and the mean and standard deviation of Δx reported per stick.

Vibrational Frequency Analysis

Dynamic softening is quantified using free harmonic oscillation testing. The handle end of the shaft is fixed in a rigid cantilever clamp. A smartphone running the Phyphox sensor suite is mounted securely near the blade heel to measure tri-axial acceleration. The tip of the shaft is displaced by a fixed distance and released. Fast Fourier Transform analysis isolates the primary natural frequency (f_n), where a decline in natural frequency corresponds directly to internal micro-crack propagation and loss of elasticity: $f_n = \frac{1}{2\pi}\sqrt{\frac{k_{eff}}{m}}$

The testing procedure will use a fixed tip placement of 5cm before release, an accelerometer sampling rate of 400 Hz, a minimum 10-second free decay recording window per trial to ensure adequate FFT frequency resolution, and $n = 5$ trials per stick/condition with the

phone remounted between trials to check for mounting-induced variance. The FFT is windowed to reduce spectral leakage before the dominant peak below 100 Hz is extracted as f_n.

Definition of the Game Fatigued Condition

To isolate the effect of accumulated use on the CCM Ribcor Trigger 10 Pro baseline stick, game-fatigued condition states are defined by elapsed calendar time in regular competitive and practice use (1-month and 3-month period) rather than by a cycle count, since the exact shot count is not recoverable retrospectively. However, it is known that the 1-month stick was used in an average of 3 ice sessions a week for 5 weeks, and the 3-month stick had an average ice time of 5 times a week, with every session being an hour and a half long. The conditions spanned separate Trigger 10 Pro sticks.

Statistical Analysis Plan

Deflection and natural frequency data will be summarized as mean $\pm$ standard deviation across trials for each stick, flex rating, and condition state. Prior to inferential testing, each stick/condition group was screened for outliers and non-normal spread using the Grubbs test.

Two comparisons are planned. The first evaluates the CCM Ribcor Trigger 10 Pro across its three condition states to isolate the rate of stiffness and frequency decline. Because the used conditions represent distinct stick specimens rather than a single stick measured longitudinally, this will be evaluated using an independent one-way ANOVA or an independent t-test. The second is a one-way ANOVA across all 14 stick/condition groups in the experimental matrix. To satisfy the statistical assumption of independence, the repeated trial measurements for each stick will be averaged into a single mean value prior to analysis. Normality of the ANOVA residuals was then assessed using the Shapiro-Wilk test. Because the normality assumption was not met, a Kruskal-Wallis test was utilized as a non-parametric robustness check on the same grouping structure. Because the overall tests returned a statistically significant result, post-hoc pairwise comparisons were performed across all stick/condition pairs, at the 95% family-wise confidence level, to identify which specific groups differed. Cost-to-effectiveness is then computed as the measured stiffness or frequency retention per dollar of retail price, allowing lower-tier or prior-generation models (e.g., the Ribcor Trigger 9 Pro, Vizion 2025) to be benchmarked directly against their manufacturer's own current flagship equivalents(e.g., Ribcor Trigger 10 Pro, Vizion 2026) at matched flex ratings.

Results and Discussion

Results are reported separately for cross-sectional differences in baseline stiffness and natural frequency across kickpoint and flex categories in new sticks, the within-session frequency shift observed immediately before versus after the static deflection test, and the aging comparison across the three CCM Ribcor Trigger 10 Pro specimens representing new, 1 month, and 3 month used conditions.

Baseline Differences in Natural Frequency Across Stick, Flex and Condition Groups

An analysis was run on the preload natural frequency across all 14 sticks in the experimental matrix. To ensure statistical independence, the 5 repeated trials for each stick were averaged into a single mean value per stick prior to analysis, following the corrected analysis plan described above. Each group's repeated trials were screened for measurement outliers using Grubbs' test before averaging.

The analysis found a significant overall difference in natural frequency across the 14 stick/condition groups, confirming that kickpoint flex and price tier are jointly associated with substantial differences in natural frequency.

Table 2 provides the full summary of testing, documenting the natural frequencies (f_n) and their standard deviations (SD) before and after static load testing, alongside the recorded maximum deflection (Δx) and resulting plastic deformation (Permanent Set).

Model	# Flex	fn Before (H	fn After (Hz)	Δx (mm)	Perm. set (mm)	Kickpoint
Ribcor Trigger 10 Pro - New	70	11.26±0.04	11.12±0.03	18.29±5.61	−4.41	Ultra-Low
Ribcor Trigger 10 Pro - 1m	70	11.74±0.02	11.72±0.03	27.50±1.54	−0.87	Ultra-Low
Ribcor Trigger 10 Pro - 3m	70	12.22±0.01	12.15±0.01	28.87±3.56	−2.89	Ultra-Low
Ribcor Trigger 9 Pro - New	85	11.94±0.02	11.02±0.02	22.18±5.35	−1.23	Low
JetSpeed FT8 Pro - New	65	10.72±0.02	10.80±0.04	30.05±19.49	−1.24	Hybrid
JetSpeed FT8 Pro - New	70	11.55±0.06	11.66±0.04	14.86±6.96	−0.75	Hybrid
JetSpeed FT8 Pro - New	85	11.85±0.02	12.43±0.03	40.44±19.43	−1.74	Hybrid
Vizion 2025 - New	70	9.98±0.17	11.19±0.07	35.31±16.25	−0.71	Ultra-Low
Vizion 2026 - New	70	8.61±0.02	10.09±0.12	49.24±23.95	−0.95	Mid
Tacks XF Ghost - New	70	11.33±0.03	11.91±0.02	12.28±5.93	−0.72	Mid
Tracer - New	70	10.43±0.02	11.00±0.03	53.32±25.21	−0.56	Mid
Twitch - New	70	12.26±0.04	12.39±0.05	51.06±23.45	−0.85	Low
Covert QR7 Pro - New	70	10.23±0.24	10.73±0.04	19.61±11.55	−2.80	Low
Covert QR7 Pro - New	75	9.34±0.02	10.07±0.02	10.25±2.67	- New	Low

Table 2: Natural Frequency immediately before versus after the static deflection test by stick, kickpoint, flex rating, and condition. This comprehensive data matrix summarizes the mean and standard deviation for natural frequency, static deflection, and permanent deformation across all tested stick configurations.

Normality of the ANOVA residuals was assessed using the Shapiro-Wilk test, which rejected the null hypothesis of normally distributed residuals. The Normal Q-Q plot of the model's residuals (Figure 2) shows a corresponding departure from the reference line at both the lower and upper tails, including heavier-than-normal tails in the residual distribution. Because the normality assumption underlying the ANOVA F-test was violated, the Kruskal-Wallis test confirmed the ANOVA's conclusion of a significant overall difference across groups, indicating that the finding of group differences is robust to the non-normality of the residuals.

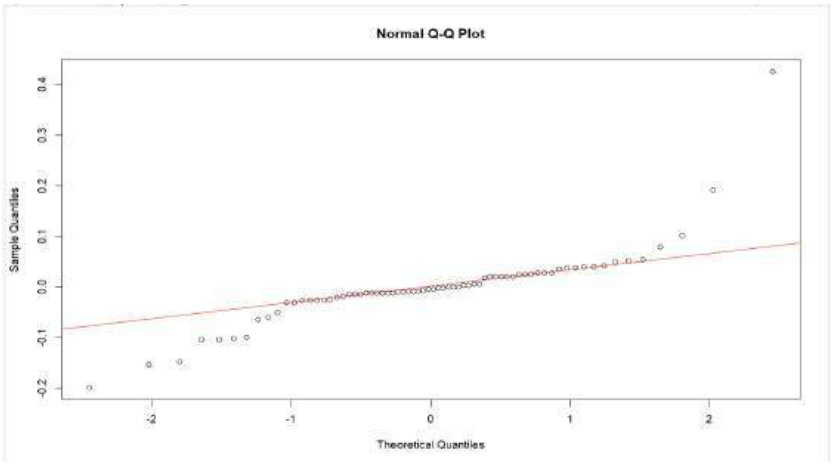

Figure 2: Normal q-q plot of model residuals displaying moderate heavy-tailed deviations and positive outliers relative to the theoretical normal distribution (red line). Figure generated using R from data collected in this study.

Using the post-load frequency retention data, Figure 3 compares their average natural frequencies with the Bauer Twitch having the largest in magnitude.

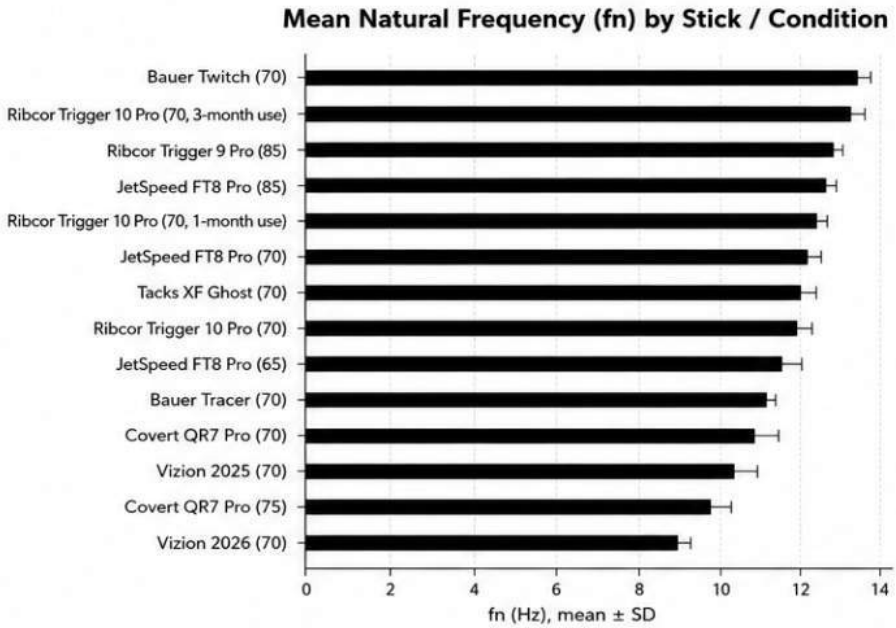

Figure 3: Average natural frequency by stick group. Figure generated using Python (Matplotlib) from data collected in this study.

Over three months of use, the Ribcor Trigger 10 Pro experienced a significant increase in static deflection, indicating a progressive loss of structural stiffness. Figure 4 shows the comparisons of static deflection and permanent set between the three aged conditions. The Static deflection shows an increase in magnitude associated with the increase of time used on ice while the permanent set doesn't follow a consistent trend, dropping from new to 1-month but increasing from 1-month to 3-months.

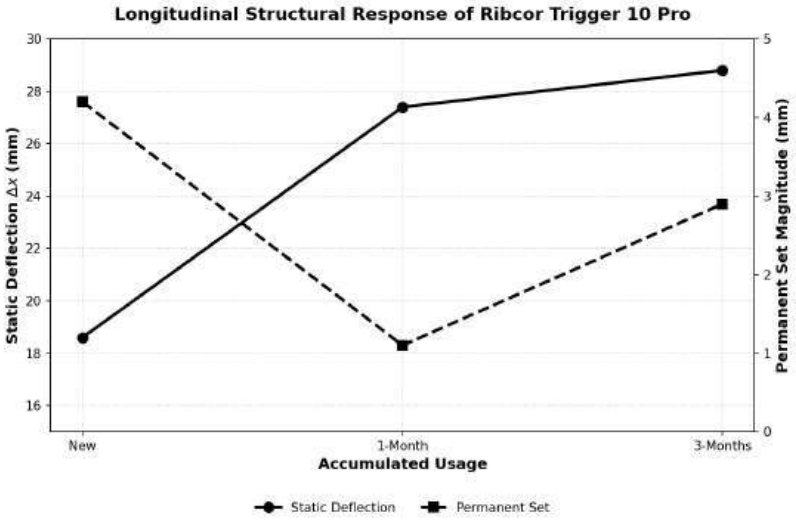

Figure 4: Longitudinal structural decay of the CCM Ribcor Trigger 10 Pro across 0, 1, and 3 months of accumulated usage. Figure generated using Python (Matplotlib) from data collected in this study.

Figure 5 shows the natural frequencies before and after their Static Deflections Tests. Contrary to expected material softening, the vast majority of tested sticks exhibited an increase in natural frequency following the static deflection test, indicating an initial structural break-in phase. This can be seen in all sticks besides the three Ribcor Trigger 10 Pros and 9Pro which all saw a decrease after the Static Deflection Test.

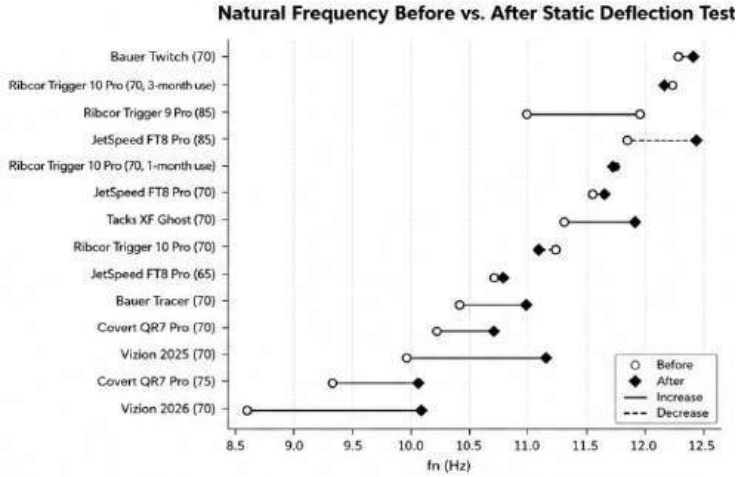

Figure 5: Natural frequency before vs. after the static deflection test, all stick groups. Figure generated using Python (Matplotlib) from data collected in this study.

To highlight the differences in static deflection and permanent set, Figure 6 compares the magnitude of static deflection from stick to stick with the Bauer Tracer having the highest Δx.The Permanent set is then ordered based on their static deflection, with the Tracer also having the lowest magnitude. The Static deflection and permanent deformation vary widely across different

stick models, highlighting significant differences in baseline flexibility and structural resilience among brands.

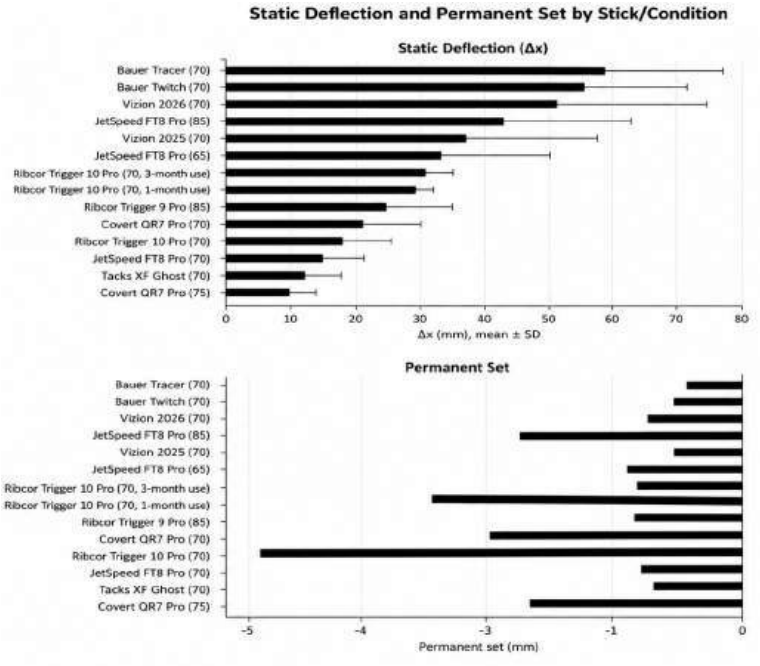

Figure 6: Static deflection and permanent set by stick group. Figure generated using Python (Matplotlib) from data collected in this study.

Table 3 and Figure 7 both the post-hoc comparisons of natural frequency. Post-hoc pairwise comparisons reveal that variations in kickpoint geometry, flex rating, and model generation produce statistically significant differences in a stick's natural frequency.

Comparison	Diff	95% CI	p adj
BauerTwitch (70 Flex) – Vizion2026 (70 Flex)	3.642	[3.454, 3.830]	<0.001
QR7Pro (70 Flex) – Vizion2026 (70 Flex)	1.62	[1.432, 1.808]	<0.001
QR7Pro (75 Flex) – Vizion2026 (70 Flex)	0.722	[0.534, 0.910]	<0.001
BauerTracer (70 Flex) – TacksXFGhost (70 Flex)	-0.898	[-1.086, -0.710]	<0.001
BauerTwitch (70 Flex) – TacksXFGhost (70 Flex)	0.926	[0.738, 1.114]	<0.001
QR7Pro (70 Flex) – TacksXFGhost (70 Flex)	-1.096	[-1.284, -0.908]	<0.001
QR7Pro (70 Flex) – TacksXFGhost (70 Flex)	-1.994	[-2.182, -1.806]	<0.001
BauerTwitch (70 Flex) – BauerTracer (70 Flex)	1.824	[1.636, 2.012]	<0.001
QR7Pro (70 Flex) – BauerTracer (70 Flex)	-0.198	[-0.386, -0.010]	0.03
QR7Pro (70 Flex) – BauerTracer (70 Flex)	-1.096	[-1.284, -0.908]	<0.001
QR7Pro (70 Flex) – BauerTwitch (70 Flex)	-2.022	[-2.210, -1.834]	<0.001
QR7Pro (75 Flex) – BauerTwitch (70 Flex)	-2.92	[-3.108, -2.732]	<0.001
QR7Pro (75 Flex) – QR7Pro (70 Flex)	-0.898	[-1.086, -0.710]	<0.001

Table 3: Selected post-hoc pairwise comparisons of natural frequency (f_n), relevant to H1- H3; full pairwise results (all 91 comparisons) are provided in Appendix A.

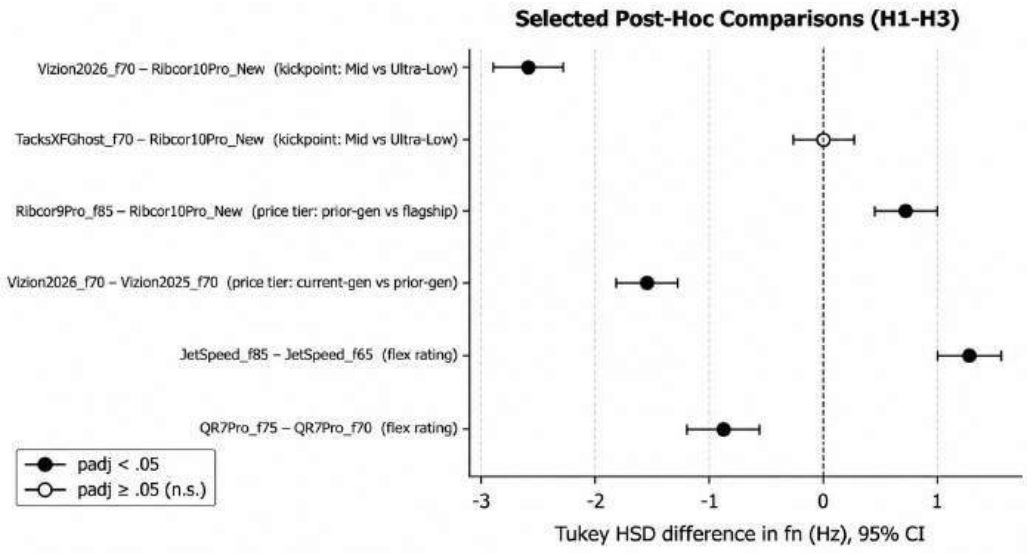

Figure 7: Visual representation of selected post-hoc comparisons of f_n; full comparison results are in Appendix A. Figure generated using Python (Matplotlib) from data collected in this study.

Brand	Model - Condition	Kickpoint	Flex	fn Before (Hz) mean	fn Before SD	fn After (Hz) mean	fn After SD	Δx (mm) mean	Δx (mm) SD	Permanent Set (mm) mean
CCM	Ribcor Trigger 10 Pro - New	Ultra-Low	70	11.26	0.04	11.12	0.03	18.29	5.61	-4.41
CCM	Ribcor Trigger 10 Pro - 1m	Ultra-Low	70	11.74	0.02	11.72	0.03	27.50	1.54	-0.87
CCM	Ribcor Trigger 10 Pro - 3m	Ultra-Low	70	12.22	0.01	12.15	0.01	28.87	3.56	-2.89
CCM	Ribcor Trigger 9 Pro - New	Low	85	11.94	0.02	11.02	0.02	22.18	5.35	-1.23
CCM	JetSpeed FT8 Pro - New	Hybrid	65	65.00	60.68	10.80	0.04	30.05	19.49	-1.24
CCM	JetSpeed FT8 Pro - New	Hybrid	70	11.55	0.06	11.66	0.04	14.86	6.96	-0.75
CCM	JetSpeed FT8 Pro - New	Hybrid	85	11.85	0.02	12.43	0.03	40.44	19.43	-1.74
CCM	Vizion 2025 - New	Ultra Low	70	9.98	0.17	11.19	0.07	35.31	16.25	-0.71
CCM	Vizion 2026 - New	Mid	70	8.61	0.02	10.09	0.12	49.24	23.95	-0.95
CCM	Tacks XF Ghost - New	Mid	70	11.33	0.03	11.91	0.02	12.28	5.93	-0.72
Bauer	Tracer - New	Mid	70	10.43	0.02	11.00	0.03	53.32	25.21	-0.56
Bauer	Twitch - New	Low	70	12.26	0.04	12.39	0.05	51.06	23.45	-0.85
Warrior	Covert QR7 Pro - New	Low	70	10.23	0.24	10.73	0.04	19.61	11.55	-2.80
Warrior	Covert QR7 Pro - New	Low	75	9.34	0.02	10.07	0.02	10.25	2.67	-2.59

Table 4: Final raw testing data of all 14 sticks/conditions (Full data in Appendix B). The comprehensive raw data matrix details the baseline structural metrics, post-load frequency shifts, and permanent deformation across all fourteen tested stick configurations.

Condition	Est. ice times (hrs)	fn mean (Hz)	fn SD	n trials
New	0	11.26	0.04	5
1 -month use	~27	11.74	0.04	5
3 - month use	~91	12.22	0.07	5

Table 5: usage data of the three Ribcor Trigger 10 Pros. Longitudinal tracking of the Ribcor Trigger 10 Pro demonstrates a direct correlation between accumulated on-ice hours and a progressive increase in baseline natural frequency.

Discussion

Evaluating the Kickpoint, Flex Rating, and Price Tier Hypotheses

As outlined in Table 1, the experimental matrix provided a comprehensive cross-section of modern carbon composite sticks, allowing for multi-variable analysis across kickpoints, flex profiles, and longitudinal wear. Despite the non-normal residual distribution illustrated in Figure 2, non-parametric robustness checks verified that the observed shifts in mechanical behavior were highly statistically significant. The ANOVA and post-hoc tests didn't locate any abnormal data, as shown in Table 3 and Figure 7.

Kickpoint was expected to produce the clearest and most statistically robust signal in this dataset. As shown in Table 2 and Figure 5, mid-kickpoint, hybrid and most low/ultra-low sticks actually rose in natural frequency post-load. Most notably, the Vizion 2026, which rose 17% from 8.61 to 10.09 Hz. Only the three Ribcor Trigger 10 and 9 Pro showed a post-load drop in natural frequency. The Composites Part B review of CFRP fatigue performance[2] and the PMC review of Mode I fatigue in fiber-reinforced polymeric composites[3] both establish that localized stress concentration is a primary driver of microcrack initiation and propagation rate, and that kickpoint zones are specifically engineered to localize bending strain near the blade. Together, these sources support the expectation that kickpoint geometry is a first-order determinant of local fatigue rate. Because low-kickpoint zones are specifically engineered to localize bending strain in a narrow taper near the blade, they inherently concentrate material stress, leading to a faster functional fatigue rate. Ultra-low kickpoint models (e.g., the CCM Ribcor Trigger 10 Pro) were therefore expected to show the largest decline in f_n from the new to 3-month-use condition, while mid-kickpoint models are expected to show comparatively smaller declines since they distribute this same strain across a broader section of the upper shaft, allowing the internal carbon matrix to resist delamination much more effectively. However, hypothesis 1 was not supported by the natural frequency data: the Ribcor Trigger 10 Pro instead rose from 11.26 to 11.74 to 12.22 Hz across the new, 1 month, 3 month conditions, and 10 of the 14 groups increased in frequency after static loading.

Inferential statistical testing on baseline natural frequency confirmed significant cross-sectional differences across stick models. Post-hoc Turkey HSD comparisons revealed statistically significant differences (adjusted $p < 0.001$) across almost all cross-sectional comparisons between ultra-low kickpoint models (e.g., CCM Ribcor Trigger 10 Pro) and mid-kickpoint models (e.g., CCM Vizion 2026 and Bauer Tracer). However, these baseline differences reflect initial shaft geometry rather than superior frequency decay resistance in mid and hybrid-kickpoint sticks. Following static loading, mid and hybrid-kickpoint sticks did not simply retain their baseline frequency; instead they showed substantial post-load frequency increases as shown in Table 2 and Figure 5. Most notably, the mid-kickpoint Vizion 2026 rose 17.2% from 8.61 Hz to 10.09 Hz after loading, while the mid-kickpoints XF Ghost (11.33 to 11.91 Hz, +5.1%), Bauer Tracer (9.98 to 11.19 Hz, +12.2%), and ultra-low-kickpoint Vizion 2025 (9.98 to 11.19 Hz, +12.2%) all rose as well.

The Flex rating hypothesis is only partially testable with the current design, since flex-rating comparisons are limited to new sticks and no aged condition exists across multiple flex ratings of the same model. Within the CCM JetSpeed FT8 Pro (Hybrid kickpoint, flex 65/70/85), f_n increased with flex rating ($f_n = 7.84 + 0.0482 \times \text{flex}$, $R^2 = .729$, $p < .0001$), consistent with elementary beam theory, in which a stiffer shaft raises natural frequency at constant mass; this is corroborated by the significant Tukey comparison between JetSpeed f85 and f65 ($\Delta = +1.132$ Hz, padj < .0001). For the Warrior Covert QR7 Pro (Low kickpoint, flex 70/75), f_n instead decreased with flex rating ($f_n = 22.81 - 0.180 \times \text{flex}$, $R^2 = .897$, $p < .0001$). Because only two flex levels were tested for the QR7 Pro, each represented by a single physical stick, this comparison cannot separate a true flex effect from ordinary manufacturing variation between two individual units, and should be treated as a single-pair comparison rather than a validated flex effect. This same caution applies to the three Ribcor Trigger 10 Pro usage conditions, which were also three separate physical sticks measured once each rather than a single stick tracked over time. The total 3-month difference of 0.96 Hz is nearly identical to the 0.90 Hz gap between the two QR7 Pro units and therefore cannot be separated form unit-to-unit manufacturing variance.

However, a novel secondary analysis of the data in Table 2 reveals an unpredicted relationship between flex rating and the break-in frequency shift. By isolating the CCM JetSpeed FT8 Pro line across its 65, 70, and 85 flex ratings, a distinct pattern emerges. The softest stick (65 flex) experienced a minimal post-load frequency shift (+0.08 Hz), the 70 flex shifted slightly more (+0.10 Hz), while the stiffest model (85 flex) exhibited a massive shift (+0.58 Hz). This directly contradicts the assumption that stiffer carbon matrices resist initial breakdown better. Instead, it suggests that under a standardized 5kg cantilever load, stiffer shafts (which contain denser plies or highly rigid epoxy matrices) undergo a more violent internal "break-in" of the resin layer when forced to bend, whereas softer flexes accommodate the load more elastically without immediate structural reorganization.

A confound for the CCM Ribcor Trigger 9 vs 10 Pro and Vizion 2025 vs 2026 was anticipated for price tier, because the price-tier pairs available in the current stick matrix also differ in kickpoint classification (Low vs. Ultra-Low, and Ultra-Low vs. Mid, respectively); any observed difference in stiffness or frequency retention between these pairs cannot be attributed to price tier or model generation alone. As established by the CompositesWorld kickpoint-engineering overview[1] and the CFRP fatigue reviews cited above,[2–3] kickpoint geometry is independently expected to influence the same outcome measures under the kickpoint itself.

Material Composition and Construction

This study did not directly measure fiber orientation, resin chemistry, or ply layup for any stick, as those properties are proprietary to each manufacturer and were not independently characterized here. The discussion below is therefore an interpretation of the measured

mechanical behavior in light of the general CFRP fatigue and vibration literature reviewed in the Introduction, not a claim about a specific stick's measured material composition.

Ultra-low and low-kickpoint sticks (Ribcor Trigger 10 Pro and 9 Pro, Vizion 2025, Covert QR7 Pro, Bauer Twitch) are marketed as using a narrower, more localized taper near the blade, which, by elementary beam theory, concentrates a larger share of bending strain in a shorter section of shaft. Localized strain concentration is the same condition that CFRP fatigue reviews identify as a primary driver of matrix microcracking and fiber-matrix debonding[2–3]. This is consistent with the Ribcor Trigger models being the only two sticks in the matrix to show the microcrack-softening signature (a decline in f_n after static loading) rather than the fixture-seating/stiffening signature seen elsewhere. The same localized-strain construction that gives these sticks their fast-release feel may also be what makes them the only sticks in the matrix where softening was large enough to be visible against the broader stiffening trend.

Mid-kickpoint (Vizion 2026, Tacks XF Ghost, Bauer Tracer) and hybrid sticks (JetSpeed FT8 Pro) distribute bending strain across a longer section of the shaft, which should reduce local strain concentration at any single point and, by the same literature, correspondingly reduce the rate of local microcrack initiation[2 3]. Along with the ultra-low Vizion 2025, these models showed some of the largest positive natural frequency shifts observed after static loading(Table 2, Figure 5), which is directionally consistent with a matrix that represents local microdamage but does not rule out the fixture-seating explanation discussed below as a contributing or dominant mechanism. To understand how these sticks might have stiffened rather than resist degradation, other factors must be considered.

Two candidate material-level mechanisms, both grounded in the composite fatigue literature rather than in a material test performed here, could explain a genuine stiffness increase over repeated use rather than the anticipated decrease[4]. First, continued post-manufacture cross-linking ("post-cure") of the thermoset resin matrix can continue for a period after a part leaves the mold, particularly under the thermal cycling a stick experiences moving between locker-room and rink-ice temperatures; a resin that is not fully cross-linked at the time of purchase would be expected to stiffen, not soften, over its first months of use. Second, moisture content in the as-manufactured composite can decrease with handling and use, and matrix stiffness in carbon fiber composites is known to be sensitive to moisture content,[13]. Both mechanisms would be expected to act on a similar timescale to the three months tested here, and neither can be confirmed or ruled out without additional material-level testing (e.g., differential scanning calorimetry for cure state, or controlled moisture-conditioning trials) that was outside the scope of this study.

The Counterintuitive Increase in Natural Frequency

Contrary to the Kickpoint hypothesis, 10 of the 14 stick groups showed an increase, not a decrease, in f_n after the static deflection test (Figure 6, Figure 3, Figure 5); only the Ribcor Trigger 10 Pro and 9 Pro declined. This pattern also held over the longer term: f_n for the Ribcor

Trigger 10 Pro rose monotonically from 11.26 Hz (new) to 11.74 Hz (1 month) to 12.22 Hz (3 months) of real game use, and its static deflection and permanent set both increased over the same period (Figure 4).

Testing spanned multiple days and ambient temperatures, so temperature drift cannot be fully ruled out as a contributing factor, but it would be expected to affect all sticks similarly rather than singling out the two Ribcor Trigger models specifically. A more likely explanation is that the clamp/support fixture seats more firmly on a stick's first load, modestly raising measured f_n for most sticks independent of any change in the composite itself, while the Ribcor Trigger's more localized kickpoint construction instead softens measurably under the same load, a change large enough to be visible against that fixture-seating effect. Critically, because static deflection and permanent set also rose over the same conditions, the static and dynamic metrics point in opposite directions, providing strong evidence that the rise in natural frequency is tracking a boundary-condition artifact rather than true shaft stiffening.

Another critical factor lies in the material composition and its reaction to the test. Under elementary beam theory, increased flexibility from resin break-in should decrease natural frequency while increasing static deflection; the fact that natural frequency instead rose alongside high static deflection further confirms that the frequency measurement is dominated by fixture seating or damping changes rather than true flexural stiffness. This suggests that the proprietary material composition of the CCM Ribcor Trigger series (utilizing Nanolite resin and Sigmatex weaves) prevents the stick from being overly stiff to begin with, requiring less force to load but suffering from fatigue decay immediately, whereas other constructions require a brief mechanical break-in period before reaching peak structural "pop".

Finally, the realities of long-term fatigue are undeniably captured in Table 5. Comparing the new, 1-month and 3-month sticks shows higher static deflection with greater ice time, although the permanent set dropped from new to 1 month before rising at 3 months, and unit-to-unit variation across the three separate sticks cannot be ruled out. While the observed increases in static deflection and permanent set strongly indicate structural fatigue, macroscopic frequency and deflection data cannot mathematically prove specific internal mechanisms such as interlaminar-shear degradation. However, it can still provide an explanation for why athletes report a "dead" feeling in older sticks despite the absence of visible external fractures.

Conclusion

This study successfully utilized static cantilever deflection and vibrational frequency analysis to track functional fatigue in composite hockey sticks. The empirical findings address the three primary hypotheses of this research: Kickpoint, Flex Rating, and Price & Material.

Contrary to Hypothesis 1, the Low and Ultra-Low Kickpoint sticks did not show the fastest decline in natural frequency. Instead, the Ribcor Trigger 10 rose in baseline natural frequency across the three usage conditions and 10 of the 14 groups increased in frequency after static loading. Only the Ribcor Trigger 9 Pro and 10 Pros decreased in frequency immediately

after static loading, where narrow tapers concentrate bending strain near the blade compared to Hybrid and Mid-Kickpoint sticks.

The influence of flex yielded inverse results. Stiffer sticks exhibited larger post-load deviation spikes and larger frequency shifts compared to their 65-flex equivalents. This indicates that highly rigid resin matrices suffer from rapid, non-uniform internal fracturing during the break-in phase, while softer flexes accommodate the load more elastically.

Similarly, the data indicates that premium flagship pricing does not yield a consistent advantage in fatigue resistance or frequency retention, though price-tier comparisons in this matrix were confounded with changes in kickpoint geometry. Rather than decaying uniformly, nearly all tested sticks exhibited a post load increase in natural frequency, demonstrating that short-term fixture seating or break-in stiffening outweighs immediate microcrack-induced softening across price tiers.

While the observed increases in static deflection and permanent set strongly indicate structural fatigue, macroscopic frequency and deflection data cannot mathematically prove specific internal mechanisms such as interlaminar-shear degradation or microcracking. Confirming these exact material behaviors would require future microscopy imaging or direct material characterization.

If additional time and funding were available, future iterations of the experimental design would benefit from greatly expanding the matrix to include exact price-tier comparisons within fixed kickpoint categories (e.g., directly comparing a $150 Low-Kick to a $350 Low-Kick) to isolate the economic cost-to-effectivness ratio more cleanly. Increased funding would also permit destructive fatigue testing using an automated pneumatic cyclic loading machine. This would allow for the precise counting of exact flexural cycles before failure, effectively removing the estimation required when tracking human ice-time in months. Furthermore, extending the test pool to include professional stock (Pro-Stock) sticks against retail equivalents would challenge differences in manufacturing quality.

AI Acknowledgment

During the preparation of this work, Gemini was used to fit all of the tested data into the table shown in Appendix B. After the use of Gemini, I reviewed and edited the content as necessary and took full responsibility for the content of this publication.

Works Cited

1"Advanced Hockey Stick Design Delivers Optimal Performance." CompositesWorld, www.compositesworld.com/articles/advanced-hockey-stick-design-delivers-optimal-performance.

2"The Fatigue Performances of Carbon Fiber Reinforced Polymer Composites – A Review." Composites Part B: Engineering, ScienceDirect, www.sciencedirect.com/science/article/pii/S2238785422017598.

3"Mode I Fatigue of Fibre Reinforced Polymeric Composites: A Review." National Center for Biotechnology Information, PMC, pmc.ncbi.nlm.nih.gov/articles/PMC9659105/.

4Ju, J., and R. J. Morgan. "Characterization of Microcrack Development in BMI-Carbon Fiber Composite under Stress and Thermal Cycling." Journal of Thermoplastic Composite Materials, 2004, journals.sagepub.com/doi/abs/10.1177/0021998304044773.

5"Fatigue Damage Mechanisms in Unidirectional Carbon-Fibre-Reinforced Plastics." Journal of Materials Science, Springer, link.springer.com/article/10.1023/A:1004684228765.

6"Friction Behaviors and Wear Mechanisms of Carbon Fiber Reinforced Composites for Bridge Cable." National Center for Biotechnology Information, PMC, www.ncbi.nlm.nih.gov/pmc/articles/PMC11644159/.

7"A Stiffness Degradation Model for Composite Laminates under Fatigue Loading." Composites Science and Technology, ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S0263822397001426.

8"A New Coupled Approach of Residual Stiffness and Strength for Fatigue of Fibre-Reinforced Composites." International Journal of Fatigue, ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S0142112301001943.

9"Experimental and Analytical Study on Residual Stiffness/Strength of CFRP Tendons under Cyclic Loading." National Center for Biotechnology Information, PMC, pmc.ncbi.nlm.nih.gov/articles/PMC7763749/.

10"Modulus Reduction and Fatigue Damage of Matrix Dominated Composite Laminates." International Journal of Fatigue, ScienceDirect, www.sciencedirect.com/science/article/abs/pii/026382239290044D.

11Jain, Atul. "A Statistical Treatment of the Loss of Stiffness during Cyclic Loading for Short Fiber Reinforced Injection Molded Composites." International Journal of Fatigue, ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S1359836816314962.

12"Fatigue Damage Analysis of GFRP Composites Using Digital Image Correlation." Journal of Ocean Engineering and Marine Energy, Springer, link.springer.com/article/10.1007/s40722-020-00184-6.

13Ferreira Rocha, P. R., et al. "Mechanisms of Component Degradation and Multi-Scale Strategies for Predicting Composite Durability: Present and Future Perspectives." Journal of Composites Science, vol. 8, no. 6, 2024, p. 204. MDPI, www.mdpi.com/2504-477X/8/6/204.

14"Fatigue Testing Verifies Material and Component Durability for Long-Term Applications."

Element, www.element.com/resources/articles/fatigue-testing-verifies-material-and-component-durability-for-long-term-applications.
15Cui, W., et al. "Very High Cycle Fatigue (VHCF) Characteristics of Carbon Fiber Reinforced Plastics (CFRP) under Ultrasonic Loading." Materials, vol. 13, no. 4, 2020, p. 908. National Center for Biotechnology Information, PMC, pmc.ncbi.nlm.nih.gov/articles/PMC7079593.
16"Detection of Cracks Present in Composite Cantilever Beam by Vibration Analysis Technique." ResearchGate, www.researchgate.net/publication/370902190_Detection_of_Cracks_Present_in_Composite_Cantilever_Beam_by_Vibration_Analysis_Technique.
17"Vibration-Based Delamination Detection in Composite Beams through Frequency Changes." ResearchGate, www.researchgate.net/publication/275437400_Vibration-based_delamination_detection_in_composite_beams_through_frequency_changes.
18"Damage Identification in Composite Beam by Vibration Measurement and Fuzzy Inference System." Academia.edu, www.academia.edu/36207542/Damage_Identification_in_Composite_Beam_by_Vibration_Measurement_and_Fuzzy_Inference_System.
19"Vibration-Based Structural Health Monitoring of Laminated Composite Beams Using Finite Element Modal and Harmonic Analysis." MDPI, www.mdpi.com/2504-477X/10/2/79.
20Ravichandran, A., and P. K. Mohanty. "Vibration Based Damage Detection for Hybrid Composite Cantilever Beam Using Finite Element Analysis and Artificial Neural Network." Journal of Thermoplastic Composite Materials, 2024, journals.sagepub.com/doi/10.1177/09574565231222617.
21"Improved Vibration Based Damage Detection in Laminated Composite Plate Structures under Free and Forced Modal Analysis." Mechanics Research Communications, ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S0997753823001237.
22"A Review on Vibration-Based Damage Detection Methods for Civil Structures." MDPI, www.mdpi.com/2571-631X/6/4/51.
23"Crack Detection in Beam Type Structure Using Frequency Data." ResearchGate, www.researchgate.net/publication/223708003_Crack_Detection_in_Beam_Type_Structure_Using_Frequency_Data.
24"Numerical and Experimental Studies on Vibration Based Methods for Detection of Damage in Composite Beams." Composite Structures, ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S0263822316329543.
25Staacks, S., et al. "Advanced Tools for Smartphone-Based Experiments: Phyphox." Physics Education, vol. 53, no. 4, 2018. IOPscience, iopscience.iop.org/article/10.1088/1361-6552/aac05e.
26"Biomechanical Investigations of Performance Characteristics of Various Types of Ice Hockey

Sticks." Academia.edu, www.academia.edu/49197375/Biomechanical_investigations_of_performance_characteristics_of_various_types_of_ice_hockey_sticks.
27"Ice Hockey Stick and Puck Biomechanical Predictors of Wrist Shot Accuracy." McGill University eScholarship, escholarship.mcgill.ca/downloads/fq977x69h.pdf.
28Laliberte, D. J. "Biomechanics of Ice Hockey Slap Shots: Which Stick Is Best?" The Sport Journal, 2009, thesportjournal.org/article/biomechanics-of-ice-hockey-slap-shots-which-stick-is-best/.
29"Experimental Characterization of Ice Hockey Sticks and Pucks." Washington State University, ssl.wsu.edu/documents/2015/10/experimental-characterization-of-ice-hockey-sticks-and-pucks.pdf.
30"Pilot Study of Embedded IMU Sensors and Machine Learning Algorithms for Automated Ice Hockey Stick Fitting." National Center for Biotechnology Information, PMC, www.ncbi.nlm.nih.gov/pmc/articles/PMC9105185/.
31Villaseñor, A., et al. "Recoil Effect of the Ice Hockey Stick during a Slap Shot." Journal of Applied Biomechanics, vol. 22, no. 3, 2006, pp. 202–211. doi.org/10.1123/jab.22.3.202.
32"Investigation of Flexural Behaviour of Polymer Composite Golf Shaft." ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S2214785317315924.
33"The Quasi-Static and Dynamic Testing of Damping in Golf Clubs Shafts Fabricated from Carbon Fibre Composites." Procedia Engineering, ScienceDirect, www.sciencedirect.com/science/article/pii/S1877705810004121.
34"Evaluation of the Mechanical Performance of Golf Shafts." ScienceDirect, www.sciencedirect.com/science/article/abs/pii/S1350630705000579.
35Fan, Q., et al. "A Review of Composite Materials for Enhancing Support, Flexibility and Strength in Exercise." Alexandria Engineering Journal, ScienceDirect, www.sciencedirect.com/science/article/pii/S1110016824002928.
36"Study on Carbon Fiber Composite Materials in Sports Equipment." ResearchGate, www.researchgate.net/publication/269355387_Study_on_Carbon_Fiber_Composite_Materials_in_Sports_Equipment.
37"Flexural Testing – an Overview." ScienceDirect Topics, www.sciencedirect.com/topics/materials-science/flexural-testing.
38"Flexure Testing of Sandwich Composites." CompositesWorld, www.compositesworld.com/articles/flexure-testing-of-sandwich-composites.
39"Cantilever Beam Static and Dynamic Response Comparison with Mid-Point Bending for Thin MDF Composite Panels." BioResources, bioresources.cnr.ncsu.edu/resources/cantilever-beam-static-and-dynamic-response-comparison-with-mid-point-bending-for-thin-mdf-composite-panels/.

The Effects of Feynman Technique, Blurting, and SQ3R on Immediate Recall of Biological Information Among Middle School Students By Maria Astezheva

ABSTRACT

Effective learning requires not only the acquisition of large amounts of information but also the use of strategies that support long-term retention. Despite the widespread use of various learning methods, the question of which strategies are most effective in supporting memory retention remains relevant.

The aim of this study was to compare the effects of three learning strategies - the Feynman Technique, Blurting, and SQ3R - on the immediate recall of biological information among school students.

The study involved 15 students aged 11-12 years, who were randomly assigned to three groups of five participants. Each group studied the same material on the topic of "the structure and functions of animal cell organelles" for 40 minutes using one of the three learning strategies. After the preparation period, participants completed a test consisting of multiple-choice questions and short-answer tasks.

The highest average score was achieved by the Feynman Technique group (8.0/10), followed by the Blurting group (7.8/10), while the SQ3R group demonstrated the lowest average score (6.8/10). The results indicate that strategies based on active information processing and retrieval may be more effective for immediate recall under the conditions of this experiment.

The findings suggest that active learning strategies may contribute to more effective short-term information retention compared with structured learning approaches. However, the small sample size and limited duration of the experiment require further research with larger groups of participants and evaluation of long-term retention.

Keywords biology education; learning strategies; Feynman Technique; Blurting; SQ3R; immediate recall; retrieval practice

INTRODUCTION

Learning is a complex cognitive process that involves the acquisition, processing, storage, and subsequent retrieval of information. For school students and university learners, the ability to effectively acquire and retain large amounts of knowledge is essential, as modern educational programs require continuous engagement with substantial volumes of material. However, many students experience difficulties with information retention, which may be associated with limitations of working memory, insufficient attention, and the use of ineffective learning strategies.

One of the common challenges in education is the widespread reliance on passive learning methods, such as repeatedly rereading texts or mechanically repeating information. Although these approaches may create a sense of familiarity with the material, research in cognitive psychology suggests that they do not always lead to deep understanding or long-term

retention of knowledge. In contrast, active learning strategies that involve retrieving information from memory, explaining concepts, and establishing connections between ideas may contribute to more effective knowledge acquisition.

One of the most extensively studied mechanisms for improving memory is **retrieval practice**, which suggests that actively attempting to recall information strengthens later memory performance. Additionally, research indicates that deeper processing of information, which requires understanding and explaining the material, contributes to the formation of more stable cognitive connections. However, different learning strategies may influence the learning process through different mechanisms, making the question of which strategy is most effective in specific educational contexts still relevant.

RESEARCH QUESTION

How do the Feynman Technique, Blurting, and SQ3R compare in immediate post-learning performance among middle school students learning biological information?

Currently, various learning methods are widely used among students, including the **Feynman Technique, Blurting, and SQ3R**. The Feynman Technique is based on explaining complex information using simple language, which encourages deeper understanding of the material. The Blurting method involves actively recalling information without access to the original source, which is closely related to the mechanism of retrieval practice. In contrast, SQ3R is a structured reading strategy that includes surveying the material, generating questions, reading, recalling information, and reviewing the content.

Despite existing research focusing on individual learning strategies, comparative analysis of these methods in the context of immediate recall biological information among school students remains limited. In particular, it is important to investigate which strategies may be most effective when students have limited preparation time and need to acquire a large amount of scientific information.

The aim of this study is to compare the effectiveness of three learning strategies, such as **Feynman Technique, Blurting, and SQ3R**, in relation to immediate recall of biological information among school students.

HYPOTHESIS

Learning strategies based on active information retrieval and deeper cognitive processing (Feynman Technique and Blurting) will be more effective in immediate recall information compared to the structured learning strategy SQ3R.

THEORETICAL FRAMEWORK

Cognitive Processes of Learning and Memory Formation

Learning is a complex cognitive process that involves the acquisition, processing, storage, and subsequent retrieval of information. The studying material but also by the way

information is processed during learning.

One of the fundamental models of memory is the **multi-store model of memory** proposed to be stored in long-term memory; it must undergo the process of encoding, during which new knowledge is transformed and connected with existing mental representations.

Traditionally, memory processes are described through three main stages: **encoding, storage, and retrieval**.

Encoding refers to the initial processing of information, during which material is transformed into a meaningful representation. Storage involves maintaining information in memory over time. Retrieval allows previously learned information to be accessed when needed.

Therefore, the effectiveness of different learning methods may depend on how actively they engage the processes of encoding and retrieval.

Working Memory and Cognitive Load in Learning

Working memory plays a key role in learning because it is responsible for temporarily holding and processing new information. However, its capacity is limited: individuals can only process a restricted amount of new information at one time. According to Alan Baddeley's model of working memory (2), working memory consists of the effectiveness of learning is determined not only by the amount of time spent.

Atkinson and Shiffrin (1). According to this model, information passes through several stages of processing: sensory memory, working memory, and long-term memory. In order for information to be of several components responsible for processing different types of information and managing attention. The central executive component regulates the allocation of attention and coordinates cognitive tasks.

The limitations of working memory are particularly important when studying complex subjects such as biology, where students must simultaneously understand numerous new terms, processes, and relationships between concepts.

John Sweller's Cognitive Load Theory (3) suggests that excessive amounts of information or ineffective organization of material can overload working memory and reduce learning efficiency.

Therefore, effective learning strategies should not only provide information but also help students organize knowledge, establish connections between concepts, and actively engage with the material.

Depth of Processing and Meaningful Learning

One of the factors determining the quality of memory retention is the depth of cognitive processing.

According to Craik and Lockhart's Levels of Processing Theory (4), information is retained more effectively when it is processed at a deeper semantic level. Shallow processing is usually associated with mechanical repetition or memorization of individual pieces of information, whereas deeper processing requires understanding meaning, establishing

connections, and explaining the material.
For example, simply memorizing the definition:

"A mitochondrion is a cell organelle."

represents a form of shallow processing. In contrast, understanding the relationship:

"Mitochondria provide energy for the cell; therefore, cells with higher energy demands contain more mitochondria."

represents a deeper level of processing.

Deep processing contributes to the creation of stronger semantic connections between elements of knowledge, making later retrieval of information easier.

This principle is closely connected to the Feynman Technique, as this strategy requires students to explain complex information using simple language, reorganize their knowledge, and identify their own level of understanding.

Retrieval Practice and the Testing Effect

One of the most extensively studied mechanisms for improving memory is retrieval practice.

Unlike passive rereading of material, retrieval practice requires individuals to actively recall information without continuous access to the original source.

Research by Roediger and Karpicke (5) demonstrated that attempting to recall previously studied material can significantly improve later memory performance. This phenomenon is known as the testing effect.

Importantly, testing in this context is not only a method of evaluating knowledge but also an independent learning tool. When students attempt to retrieve information from memory, they strengthen their ability to access this knowledge in the future.

This mechanism is directly connected to the Blurting method. During this strategy, students first attempt to reproduce learned material without using textbooks or notes, then identify gaps in their knowledge and return to insufficiently learned information.

Therefore, Blurting is based on the principle of retrieval practice and may contribute to stronger memory retention.

Generation Effect and the Role of Explanation in Learning

Another important mechanism contributing to effective learning is the generation effect (6).

This effect suggests that information is remembered more effectively when individuals actively create, formulate, or transform information themselves rather than simply receive a prepared explanation.

Creating personal explanations requires more active cognitive processing: students must

identify key ideas, establish connections between concepts, and represent information in a new form.

The Feynman Technique is based on this principle. Instead of mechanically repeating information, students attempt to explain complex concepts using simple language, as if they were teaching another person.

This process requires deep understanding of the material and contributes to the creation of a more organized knowledge structure.

Structured Learning Strategies: The SQ3R Method

Unlike strategies primarily based on active retrieval and explanation, some learning methods focus mainly on organizing and structuring information (7).

One example of such a method is SQ3R, which consists of five stages:

- Survey - reviewing the material before studying;
- Question - generating questions about the content;
- Read - carefully studying the information;
- Recite - recalling and summarizing the main ideas;
- Review - revisiting the material.

The main purpose of SQ3R is to provide a systematic approach to studying texts, improve understanding of material structure, and increase student engagement during reading.

Although SQ3R includes elements of active learning, its primary focus is on information organization and comprehension. Therefore, this strategy may engage different cognitive mechanisms compared with methods primarily based on active retrieval (8).

Theoretical Comparison of the Studied Strategies

The three strategies examined in this study are based on different cognitive mechanisms.

Method	**Main cognitive mechanism**	**Theoretical part**
Blurting	Retrieval practice	Retrieval Practice, Testing Effect
Feynman Technique	Deep processing and self-generated explanations	Levels of Processing Theory, Generation Effect

Table 1. *Learning Methods: Underlying Cognitive Mechanisms and Theoretical Foundation*

Overall, different learning strategies may influence memory through different cognitive processes. Methods that require active retrieval of information and deeper semantic processing may theoretically contribute to more effective knowledge retention compared with strategies primarily focused on organization and repetition of material (9).

Based on existing research in cognitive psychology of learning, it can be hypothesized that the **Feynman Technique and Blurting** will demonstrate higher effectiveness in improving

immediate recall information compared with **SQ3R**.

METHODOLOGY

Research Design

This study used an experimental comparative design to compare immediate post-learning performance following the use of three different learning strategies: the Feynman Technique, Blurting, and SQ3R.

The Feynman Technique involves explaining material in one's own words and identifying gaps in understanding. Blurting involves retrieving information from memory without referring to the original material. SQ3R (Survey, Question, Read, Recite, Review) is a structured reading strategy that combines reading, questioning, recitation, and review.

The independent variable was the learning strategy assigned to each participant. The dependent variable was the score on the post-learning assessment, which was completed immediately after the 40-minute learning intervention.

Participants

The study included 15 students aged 11–12 years: 4 boys and 11 girls. All participants were classmates and members of the researcher-led Science Lab Club, a biology-focused extracurricular program. The participants were recruited from this existing group because the research was conducted as part of the club's educational activities.

The inclusion criteria were membership in the Science Lab Club and willingness to participate in the study. No additional exclusion criteria were applied.

Before the intervention, participants were asked about their previous exposure to the study material and their familiarity with the three learning strategies. None of the participants had previously studied the cell-related material used in the study, and none reported prior experience with the Feynman Technique, Blurting, or SQ3R.

Prior biology knowledge was not formally assessed using a pretest. The study used convenience sampling because the participants were recruited from the existing Science Lab Club group.

Before the experiment, participants were randomly assigned to three groups of five students:

● Feynman Technique group (n = 5); ● Blurting group (n = 5);

● SQ3R group (n = 5).

Random assignment was used to reduce the influence of individual differences between participants and improve the objectivity of comparisons between groups.

Research Materials

The learning material used in the study covered the structure and functions of animal cell organelles. The same material was provided to all participants in the same format and was used

by all three groups during the intervention.

The material included information about the main structures and functions of the animal cell, including the cell membrane, cytoplasm, and nucleus, as well as basic information about cells and their interaction with the environment.

The complete instructional material used in the study is provided in Appendix A.

Experimental Procedure

The experiment was conducted in 2 stages.

Stage 1. Random Assignment

Participants were randomly assigned to one of three intervention groups using a simple draw-lots procedure. Fifteen identical paper slips were prepared, with five slips labeled "Feynman," five labeled "Blurting," and five labeled "SQ3R." The slips were folded so that the labels were not visible. Each participant independently selected one slip without knowing its label in advance. Each participant was allowed one draw, and no slips were replaced or redrawn. This procedure resulted in five participants being assigned to each intervention group.

Stage 2. Learning Interventions

All participants studied the same instructional material on cell structure and organelles for 40 minutes. The material was provided in the same format to all three groups. Participants were also given blank sheets of paper, which they could use freely for notes, explanations, diagrams, or written recall during the intervention. The researcher remained present throughout the intervention and monitored participants' adherence to the assigned strategy as the instructor of the Science Lab Club.

Feynman Technique

Participants assigned to the Feynman group were instructed to first study the material and then close their textbook or learning materials. They were asked to explain the concepts in simple language as if explaining the material to a younger student. When they became unable to explain a concept or used a term they did not fully understand, they were instructed to return to the material, review the relevant information, and attempt the explanation again. The stated goal was conceptual understanding rather than simple memorization.

Blurting

Participants assigned to the Blurting group were instructed to study the material and then close their textbook and notes. They were asked to write down everything they could remember about the topic without referring to the source material. They then reopened their materials, identified information they had forgotten or recalled incorrectly, reviewed those parts, and attempted the recall process again. The stated goal was to strengthen memory through active retrieval.

SQ3R

Participants assigned to the SQ3R group followed the five stages of the SQ3R method. During the survey, they briefly examined headings, images, and keywords. During Question, they converted headings into questions. During Read, they read the material to find answers to those questions. During Recite, they closed the material and explained what they had learned.

During Review, they revisited the material and reviewed the main points. The stated goal was to support structured comprehension and retention.

ASSESSMENT

The assessment consisted of six multiple-choice questions, each worth 1 point, and one short-answer question worth up to 4 points. The maximum possible score was therefore 10 points. The complete assessment, answer key, and scoring rubric are provided in Appendix A. After completing the 40-minute learning intervention, all participants completed the same test without access to the learning materials.

For the multiple-choice section, each correct answer received 1 point, for a maximum of 6 points. The short-answer question was scored using a 4-point rubric based on the accuracy and completeness of the response. The maximum total score was 10 points.

STATISTICAL ANALYSIS

Descriptive statistics were calculated for each learning-strategy group. These included the mean, standard deviation, median, and range of total assessment scores. Individual participant scores were also recorded and examined.

Because the study included only five participants per group and the assessment scores were bounded between 0 and 10, an exact permutation version of the Kruskal–Wallis test was used to explore differences in total assessment scores between the three groups. The test statistic was $H = 1.066$. The exact permutation p-value was .655.

Due to the small sample size, the statistical analysis was considered exploratory, and the results were interpreted together with the descriptive statistics and individual participant scores.

ETHICAL CONSIDERATIONS

This research did not receive formal IRB approval because it was conducted as an independent high school research project and was not affiliated with an institution that provided an ethics review process.

Participants were informed that they were taking part in a research project and were told about its purpose and procedures. Participation was voluntary, and participants were free to stop participating at any time. Parents or legal guardians were aware that their children were participating in the educational program in which the research activities were conducted and had given permission for their participation in the program. However, no separate written parental consent form was collected specifically for the research component. Participants provided verbal agreement to participate; no separate written assent form was administered.

No personally identifying information was collected. Participants were assigned anonymous identifiers, and names, contact information, and other personally identifying data were not recorded. Results were stored and analyzed anonymously and used solely for research purposes.

The study involved educational activities related to learning strategies and biological information retention and did not involve physical procedures, medical interventions, or activities expected to cause harm. The researcher maintained respectful treatment of participants and responsible handling of research data throughout the study.

RESULTS

A total of 15 students aged 11–12 participated in the study. Five students were assigned to each of the three learning-strategy groups: the Feynman Technique, Blurting, and SQ3R.

After 40 minutes of studying the same material, all participants completed the same post-learning assessment. The test was given immediately after the intervention and included six multiple-choice questions and one short-answer question. The maximum score was 10 points: 6 points for the multiple-choice section and 4 points for the short-answer section. Individual participant scores are shown in Table 1.

Participant	Strategy	MC (0–6)	Short answer (0–4)	Total (0–10)
P1	Feynman Technique	6	4	10
P2	Feynman Technique	5	4	9
P3	Feynman Technique	5	4	9
P4	Feynman Technique	3	3	6
P5	Feynman Technique	3	3	6
P6	Blurting	5	4	9
P7	Blurting	5	4	9
P8	Blurting	5	3	8
P9	Blurting	4	2	6
P10	Blurting	4	3	7
P11	SQ3R	6	4	10
P12	SQ3R	5	3	8
P13	SQ3R	4	2	6

P14	SQ3R	4	2	6
P15	SQ3R	3	1	4

Table 1. *Individual participant scores by learning strategy*

The Feynman group had a mean total score of 8.0/10 (SD = 1.87), with a median score of 9 and scores ranging from 6 to 10. The Blurting group had a mean score of 7.8/10 (SD = 1.30), with a median of 8 and scores ranging from 6 to 9. The SQ3R group had a mean score of 6.8/10 (SD = 2.28), with a median of 6 and scores ranging from 4 to 10. The descriptive statistics are presented in Table 2.

Learning strategy	**n**	**Mean**	**SD**	**Median**	**Range**
Feynman	5	8	1.87	9	6–10
Blurting	5	7.8	1.30	8	6–9
SQ3R	5	6.8	2.28	6	4–10

Table 2. *Descriptive statistics of post-learning assessment scores*

The Feynman group had the highest mean total score, followed closely by the Blurting group. The difference between the Feynman and Blurting group means was 0.2 points, while the difference between the Feynman and SQ3R group means was 1.2 points. The SQ3R group showed the widest range of individual scores, from 4 to 10 points.

The results also differed between the two sections of the assessment. For the multiple-choice section, the Blurting group had the highest mean score (4.6/6), while the Feynman and SQ3R groups both had a mean score of 4.4/6. For the short-answer section, the Feynman group had the highest mean score (3.6/4), followed by the Blurting group (3.2/4) and the SQ3R group (2.4/4).

DISCUSSION OF RESULTS

The aim of this study was to compare immediate post-learning performance following the use of the Feynman Technique, Blurting, and SQ3R. The Feynman group had the highest mean total score (8.0/10), followed by Blurting (7.8/10).

The relatively high mean score in the Feynman group may be related to the cognitive activity required by the strategy. According to the Levels of Processing Theory, deeper semantic processing may support memory. The Feynman Technique also requires participants to explain concepts in their own words, identify gaps in understanding, and reorganize information. This may promote active encoding and self-explanation. However, depth of processing and self-explanation quality were not directly measured in this study, so these mechanisms remain

possible explanations rather than demonstrated causes.

The Blurting group also showed a high mean score (7.8/10). Its emphasis on retrieving information without access to the source may have contributed to performance. This is consistent with research on retrieval practice and the testing effect. The Blurting group also had the For total assessment scores, the Feynman group had the highest mean score (8.0/10), followed by the Blurting group (7.8/10) and the SQ3R group (6.8/10). Individual scores ranged from 6 to 10 in the Feynman group, from 6 to 9 in the Blurting group, and from 4 to 10 in the SQ3R group. and SQ3R (6.8/10). However, the exact permutation Kruskal–Wallis analysis did not show a statistically significant difference between the groups ($H = 1.066$, $p = .655$).
Therefore, the observed differences should be considered descriptive and exploratory.
highest mean score on the short-answer section (3.2/4), possibly because the strategy required repeated free recall. However, the small sample does not allow this observation to be interpreted as evidence that Blurting is more effective for open-ended questions.

The SQ3R group had the lowest mean score (6.8/10). This does not indicate that SQ3R is generally ineffective. SQ3R includes surveying, questioning, reading, recitation, and review, and its benefits may become more apparent during longer or repeated study sessions. The lower score observed in this experiment may also be related to the immediate assessment following a single 40-minute session.

Overall, the Feynman and Blurting groups had higher descriptive mean scores than the SQ3R group, but the statistical analysis did not show a significant difference. Therefore, the results do not provide sufficient evidence to conclude that either Feynman or Blurting is generally superior to SQ3R.

Several limitations should be considered. The sample was small ($n = 15$) and was recruited through convenience sampling from a single Science Lab Club, limiting generalizability. No formal pretest was conducted, so possible differences in baseline knowledge or learning ability could not be measured. Although none of the participants had previous experience with the three strategies, differences in how quickly they adapted to each method may have affected performance. The study also examined only one biological topic and one 40-minute learning session, and immediate rather than delayed assessment. The absence of a conventional-study control group also prevents comparison with ordinary studying.

Future research should use larger and more diverse samples, include baseline testing and repeated sessions, and use delayed assessments. Including a conventional-study control group and several biological topics would also help determine whether the observed patterns remain consistent across different learning conditions.

CONCLUSION

The aim of this study was to compare immediate post-learning performance following the use of the Feynman Technique, Blurting, and SQ3R among school students studying biological information.

The Feynman group had the highest mean score (8.0/10), followed by Blurting (7.8/10)

and SQ3R (6.8/10). However, the exact permutation Kruskal–Wallis analysis did not show a statistically significant difference between the three groups ($H = 1.066$, $p = .655$). Therefore, the observed differences should be interpreted as descriptive rather than evidence that one learning strategy was more effective than the others.

The results showed that the Feynman and Blurting groups performed somewhat better on average than the SQ3R group under conditions of this experiment. These differences may be related to the active explanation and retrieval required by the Feynman Technique and Blurting, but the underlying cognitive mechanisms were not directly measured.

The findings are limited by the small sample size, convenience sampling from a single educational group, the use of one biological topic, Future research should include larger and more diverse samples, baseline testing, repeated learning sessions, delayed assessments, and a conventional-study control group. Studying multiple biological topics and other academic subjects could also help determine whether the observed patterns are consistent across different learning conditions.

CONFLICT OF INTEREST

The author declares that there are no conflicts of interest related to this work.

Works Cited

1. Atkinson RC, Shiffrin RM. Human memory: a proposed system and its control processes. In: Spence KW, Spence JT, editors. The psychology of learning and motivation: advances in research and theory. Vol. 2. New York: Academic Press; 1968. p. 89-195.https://escholarship.org/uc/ite m/5kd4s4j3
2. Baddeley A. Working memory: theories, models, and controversies. Annu Rev Psychol. 2012;63:1-29.https://doi.org/10.1146 /annurev-psych-120710-100422
3. Sweller J. Cognitive load during problem solving: effects on learning. Cogn Sci. 1988;12(2):257-285. https://doi.org/10.1207/s15516709co g1202_4
4. Craik FIM, Lockhart RS. Levels of processing: a framework for memory research. J Verbal Learn Verbal Behav. 1972;11(6):671-684. https://doi.org/10.1016/S0022-5371(72)80001-X
5. Roediger HL, Karpicke JD. Test-enhanced learning: taking memory tests improves long-term retention. Psychol Sci. 2006;17(3):249-255. https://doi.org/10.1111/j.1467-9280.2 006.01693.x
6. Slamecka NJ, Graf P. The generation effect: delineation of a phenomenon. J Exp Psychol Hum Learn Mem. 1978;4(6):592-604. https://doi.org/10.1037/0278-7393.4. 6.592
7. Chi MTH, Bassok M, Lewis MW, Reimann P, Glaser R. Self-explanations: how students study and use examples in learning to solve problems. Cogn Sci. 1989;13(2):145-182. https://doi.org/10.1207/s15516709co g1302_1
8. Robinson FP. Effective study. New York: Harper & Brothers; 1946. 9. Dunlosky J, Rawson KA, Marsh EJ, Nathan MJ, Willingham DT. Improving students' learning with effective learning techniques: promising directions from cognitive and educational psychology. Psychol Sci Public Interest. 2013;14(1):4-58. https://doi.org/10.1177/15291006124 53266

Geometric Predictive Control of Wheeled Mobile Robots with an Actuated Directional Sensor on $SE(2) \times SO(2)$ By Winson Zhang

Abstract

This paper presents a geometric model predictive control framework for trajectory tracking of a nonholonomic wheel robot equipped with an independently actuated directional sensor. The configuration is modeled on the product Lie group $SE(2) \times SO(2)$, and group-consistent tracking errors are defined through Lie-group logarithms. A smooth sensor reference is constructed using a continuous angular lift, while the local prediction model in the product Lie algebra yields a convex quadratic program. Numerical results demonstrate simultaneous base tracking and sensor pointing, and show that the proposed Product GMPC outperforms Product NMPC in overall position-tracking accuracy on the tested trajectories. The formulation provides a coordinate-consistent basis for predictive control of mobile robots with steerable directional sensors.

I. INTRODUCTION

Mobile robots increasingly carry steerable cameras, lidars, antennas, or inspection heads whose viewing direction must be controlled while the base follows a path. This creates a coupled geometric tracking problem: the planar pose belongs to $SE(2)$, the relative sensor orientation belongs to $SO(2)$, and both must respect actuator limits. Treating their angles as unconstrained real coordinates obscures the equivalence of representatives separated by 2π and can introduce artificial errors at a coordinate branch cut [2], [6], [7], [9].

The robot–sensor state is

$$Z = (X, S) \in SE(2) \times SO(2),$$

where X is the base pose and S is the sensor orientation relative to the body. If $S = R(\alpha)$ and the base heading is θ, the sensor axis has global bearing $\theta + \alpha$. The task is to follow a prescribed base trajectory while steering this axis toward a stationary observed point q.

Model predictive control (MPC) is attractive because it handles nonholonomic dynamics and input bounds in one finite-horizon problem [15]–[17]. A coordinate NMPC formulation, however, inherits the selected angular chart. Inspired by geometric tracking and manifold MPC [1], [11], [20], we formulate the prediction error on the product group and solve a convex quadratic approximation in its Lie algebra.

The circular variable must be handled carefully. A scalar angle in the interval $[-\pi, \pi)$ is only a local coordinate. It necessarily jumps at the branch cut, although the corresponding rotation matrix changes continuously. Direct subtraction or finite differencing of principal coordinates can therefore create an artificial error or angular velocity of order $2\pi/\Delta t$. Our formulation instead computes errors through $\mathrm{Log}_{SO(2)}$, uses circular losses, and differentiates a continuous Lie-group lift.

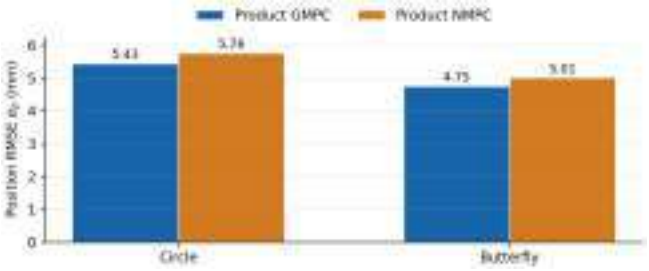

Fig. 1. Overall position-trajectory error for the wheel robot with sensor. Product GMPC reduces the full-trajectory RMSE relative to Product NMPC by 5.6% on the circle and 5.1% on the butterfly. This claim is restricted to e_p: Product NMPC has slightly lower orientation RMSE in these nominal runs.

The contributions are fourfold. First, we formulate the wheel robot with a sensor intrinsically on $SE(2) \times SO(2)$. Second, we construct a smooth sensor reference using a periodic pointing loss and a continuous lift for differentiation. Third, we derive Product GMPC, a constrained receding horizon controller whose online problem remains a convex QP. Fourth, we compare Product GMPC and Product NMPC on multiple references; the sensor-free wheel robot appears only as a reduction check that isolates the effect of appending the sensor subsystem.

II. RELATED WORK AND SCOPE

Trajectory tracking for wheeled mobile robots is complicated by the nonholonomic no-side-slip constraint and periodic orientation. Classical geometric and time-varying feedback laws establish foundational tracking results for unicycle-type vehicles [12]–[14]. NMPC adds predictive constraint handling but requires repeated nonlinear optimization [3], [15], [17]. Real-time iteration methods and sparse solvers make this increasingly practical [4], [5], [18].

Lie-group formulations preserve rigid-motion structure and provide coordinate-consistent local errors [2], [6]–[8]. Riemannian optimization supplies a broader framework for computation on manifolds [10], while geometric controllers exploit the same structure in feedback design [11]. For predictive control, Lu *et al.* formulate tracking on differentiable manifolds [20]; Tang *et al.* specialize Lie-algebra linearization to wheeled robots and obtain a QP-based GMPC on $SE(2)$ [1]. Product groups naturally describe systems with multiple geometric components, but actively steered planar sensor orientation has not been the focus of these works.

The setting covers a pan camera, scanning range sensor, or small turret with an axis normal to the ground plane. It models the sensor as an actuated orientation, not as a perception pipeline: field of view, occlusion, measurement noise, and target uncertainty remain outside scope. The prescribed base path and sensor-pointing objective are separated so that the effect of sensor actuation and its rate limit can be identified.

III. ROBOT–SENSOR MODEL ON $SE(2) \times SO(2)$

A. State Representation

Let $X \in SE(2)$ denote the base pose,

$$X = \begin{bmatrix} R(\theta) & p \\ 0_{1\times 2} & 1 \end{bmatrix}, \quad p = \begin{bmatrix} x \\ y \end{bmatrix}, \quad R(\theta) \in SO(2).$$

Let $S = R(\alpha) \in SO(2)$ represent the sensor orientation relative to the robot body. The matrix Lie group $SO(2)$ is isomorphic to the unit circle S^1, but we use $SO(2)$ throughout so that the group operations, exponential, logarithm, and Lie algebra remain explicit. The full state is

$$Z = (X, S) \in G := SE(2) \times SO(2).$$

If α is a sensor yaw angle relative to the robot body, then the global sensor bearing is

$$\beta = \theta + \alpha.$$

An element of the Lie algebra g = se(2) × so(2) is represented by

$$\eta = \begin{bmatrix} v_x & v_y & \omega & \nu \end{bmatrix}^\top \in \mathbb{R}^4,$$

where (vx, vy, ω) is the base body twist and ν is the sensor angular velocity. The product wedge map is

$$\eta^\wedge = \left(\begin{bmatrix} 0 & -\omega & v_x \\ \omega & 0 & v_y \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & -\nu \\ \nu & 0 \end{bmatrix} \right).$$

B. Circular Coordinates and the Branch Cut

The matrix $S = R(\alpha)$ is unchanged when α is replaced by $\alpha + 2\pi m$ for any integer m. Consequently, the physical sensor state is an equivalence class rather than a unique real number. The principal coordinate $\operatorname{atan2}(S_{21}, S_{11}) \in (-\pi, \pi]$ has a branch cut. For example, the two coordinates π−ε and −π+ε are separated by almost 2π in R, while their geodesic separation on SO(2) is only 2ε.

For this reason, the controller uses the Lie-group error

$$\phi = \operatorname{Log}_{SO(2)}(S_d^{-1} S), \qquad (1)$$

where Sd = R(αd). Likewise, the pointing residual is

$$e_{\text{point}} = \operatorname{Log}_{SO(2)}\left(R(\gamma)^{-1} R(\theta) S\right).$$

The loss $1 - \cos(e_{\text{point}})$ is continuous and 2π-periodic. Its local quadratic behavior near zero is appropriate for accurate pointing, while its periodicity prevents an artificial penalty at the coordinate boundary.

Reference differentiation uses a continuous lift $\alpha e_k \in \mathrm{R}$ satisfying $R(\alpha e_k) = S_k$. It is constructed by accumulating successive group logarithms,

$$\tilde{\alpha}_{k+1} = \tilde{\alpha}_k + \operatorname{Log}_{SO(2)}(S_k^{-1} S_{k+1}).$$

We then compute $\nu_{d,k} \approx \dfrac{\widetilde{\alpha}_{k+1} - \widetilde{\alpha}_k}{\Delta t}$

and obtain the stored coordinate from the rotation matrix only for serialization and plotting. Directly differentiating the principal coordinate would incorrectly report a rate near $2\pi/\Delta t$ whenever the reference crosses the branch cut. The numerical code audit below verifies that continuous lifts, rather than principal coordinates, enter differentiation and the Euler-coordinate baseline.

C. Kinematics

The control input is

$$u = \begin{bmatrix} \mu & \omega & \nu \end{bmatrix}^\top,$$

where μ is the base forward velocity, ω is the base angular velocity, and ν is the relative sensor angular velocity. The nonholonomic constraint removes lateral base velocity, so the admissible product velocity is

$$\eta = Du, \qquad D = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The continuous-time dynamics are

$$\dot{X} = X\zeta^\wedge, \qquad \zeta = \begin{bmatrix} \mu & 0 & \omega \end{bmatrix}^\top, \tag{2a}$$
$$\dot{S} = S\nu^\wedge. \tag{2b}$$

In local coordinates,

$$\dot{x} = \mu\cos\theta, \quad \dot{y} = \mu\sin\theta, \tag{3a}$$
$$\dot{\theta} = \omega, \quad \dot{\alpha} = \nu. \tag{3b}$$

With sampling time Δt, the group update is

$$Z_{k+1} = Z_k \operatorname{Exp}_{\mathcal{G}}(\Delta t\, Du_k),$$

or equivalently

$$X_{k+1} = X_k \operatorname{Exp}_{SE(2)}\left(\Delta t \begin{bmatrix} \mu_k & 0 & \omega_k \end{bmatrix}^\top\right), \tag{4a}$$
$$S_{k+1} = S_k \operatorname{Exp}_{SO(2)}(\Delta t\, \nu_k). \tag{4b}$$

IV. SENSOR-POINTING REFERENCE GENERATION

The base reference is generated by the original GMPC reference generator:

$$\{X_{d,k}, u^b_{d,k}\}_{k=0}^{N}, \qquad u^b_{d,k} = \begin{bmatrix} \mu_{d,k} & \omega_{d,k} \end{bmatrix}^\top.$$

We then optimize only the circular reference {αd,k} and its velocity {vd,k}. This keeps

the numerical experiment close to the original code while still introducing a nontrivial manifold component.

Assume that a target point q ∈ R should be observed by the sensor. The desired global bearing from the robot to the target is

$$\gamma_k = \operatorname{atan2}(q_y - y_{d,k}, q_x - x_{d,k}).$$

The optimized circular reference is obtained from

$$\begin{aligned}
\min_{\{\alpha_{d,k},\nu_{d,k}\}} \; & \sum_{k=0}^{N} w_p \left(1 - \cos(\theta_{d,k} + \alpha_{d,k} - \gamma_k)\right) \\
& + \sum_{k=0}^{N-1} w_\nu \nu_{d,k}^2 + \sum_{k=0}^{N-2} w_{\Delta\nu} (\nu_{d,k+1} - \nu_{d,k})^2 \quad \text{(5a)} \\
\text{s.t.} \; & \alpha_{d,k+1} = \alpha_{d,k} + \Delta t \, \nu_{d,k}, \quad \text{(5b)} \\
& \alpha_{d,0} = \alpha_{\text{init}}. \quad \text{(5c)}
\end{aligned}$$

The first term is a circular loss, not a Euclidean angle subtraction. It is minimized when the global sensor direction θd,k + αd,k points to the target. In the implementation, the optimization variable is the continuous lift αed and νd is obtained by differentiating this lift. The weights are wp = 120, wν = 0.08, and wΔν = 0.002 after the discrete-time scaling used in the code; an additional initial-condition penalty of weight 20 and a circular reference penalty of weight 4 are included. The offline reference is not projected onto the actuator-rate bounds. Instead, the online MPC imposes |ν| ≤ νmax, so infeasible pointing requests appear explicitly as actuator saturation. After solving (5), the product reference is

$$Z_{d,k} = (X_{d,k}, R(\alpha_{d,k})), \qquad u_{d,k} = \begin{bmatrix} \mu_{d,k} & \omega_{d,k} & \nu_{d,k} \end{bmatrix}^\top.$$

The separation between base and sensor reference generation is a modeling choice. It preserves the prescribed base path and assigns the pointing task to the sensor actuator. A fully coupled planner could trade base rotation against sensor motion and might improve feasibility when the sensor has a narrow angular or rate range; the decoupled construction instead isolates the sensor-orientation subproblem.

V. PRODUCT-STATE GMPC

A. Tracking Error

Let Z = (X, S) be the actual state and Zd = (Xd, Sd) be the reference. We use the right-invariant product error

$$E = Z_d^{-1} Z = (X_d^{-1} X, S_d^{-1} S).$$

The Lie algebra error is

$$e = \mathrm{Log}_{\mathcal{G}}(E) = \begin{bmatrix} \psi \\ \phi \end{bmatrix} \in \mathbb{R}^4,$$

where

$$\psi = \mathrm{Log}_{SE(2)}(X_d^{-1} X) \in \mathbb{R}^3,$$
$$\phi = \mathrm{Log}_{SO(2)}(S_d^{-1} S) \in (-\pi, \pi].$$

B. Linearized Error Dynamics

Let η = Du and ηd = Dud. Around the identity error, the product error dynamics are approximated by

$$\dot{e} = A_d(t)e + Du + c_d(t), \qquad c_d(t) = -\eta_d(t),$$

with

$$A_d(t) = \begin{bmatrix} -\mathrm{ad}_{\zeta_d(t)} & 0 \\ 0 & 0 \end{bmatrix}.$$

For ζd = [vx,d, vy,d, ωd]$^{\top}$,

$$\mathrm{ad}_{\zeta_d} = \begin{bmatrix} 0 & \omega_d & -v_{y,d} \\ -\omega_d & 0 & v_{x,d} \\ 0 & 0 & 0 \end{bmatrix}.$$

The last diagonal block is zero because SO(2) is Abelian. Thus the sensor-orientation error obeys $\dot{\varphi} = \nu - \nu_d$.

The implementation optimizes the absolute control u, not the increment δu = u − ud. Consequently, the affine term −ηd must remain in the dynamics. The equivalent incremental form is

$$\dot{e} = Ade + D\delta u,$$

but it is valid only if the QP variable and its bounds are transformed consistently. Mixing the incremental dynamics with absolute-control bounds would introduce a reference- dependent bias. The static implementation audit confirmed that the code uses the absolute-control convention consistently in both dynamics and cost.

Using Euler discretization,

$$e_{k+1} = A_k e_k + B_k u_k + c_k,$$

where

$$A_k = I + \Delta t A_d(t_k), \quad B_k = \Delta t D, \quad c_k = -\Delta t \eta_{d,k}.$$

C. Quadratic Program

At each time step, the product-state GMPC solves

$$\min_{\{u_k\}_{k=0}^{T-1}} \; e_T^\top Q_f e_T + \sum_{k=0}^{T-1} \big[e_k^\top Q e_k + (u_k - u_{d,k})^\top H (u_k - u_{d,k}) \big] \tag{6a}$$

$$\text{s.t.} \quad e_{k+1} = A_k e_k + B_k u_k + c_k, \tag{6b}$$

$$\underline{u} \le u_k \le \overline{u}, \tag{6c}$$

$$E_0 = Z_{d,0}^{-1} Z_0, \tag{6d}$$

$$e_0 = \mathrm{Log}_{\mathcal{G}}(E_0). \tag{6e}$$

The problem remains a convex QP. Compared with the original SE(2) GMPC, the state dimension increases from 3 to 4 and the control dimension increases from 2 to 3.

Convexity follows because the prediction dynamics are affine, the actuator limits are box constraints, and Q, Qf , and H are positive semidefinite. In the experiments we use

$$Q = \mathrm{diag}(20000, 20000, 2500, 4000), \quad Q_f = 100Q,$$

and H = diag(0.35, 0.35, 0.08). The horizon is T = 10. Following the TurtleBot limits reported in [1], base controls satisfy $-0.22 \le \mu \le 0.22$ m/s and $-2.84 \le \omega \le 2.84$ rad/s; the sensor satisfies $|\nu| \le 2.0$ rad/s. These values are fixed for all model, controller, and trajectory comparisons.

At every simulation step, the current product error is computed from $X^{-1}{}_d X$ and Eq. (1). The QP is then transcribed in CasADi and solved with qpOASES [19]. The resulting first control is applied to the exact group update for one sampling period. Repeating this process provides a direct receding horizon test of both the geometric error representation and the constrained sensor actuator.

VI. STRUCTURAL PROPERTIES AND IMPLEMENTATION

Proposition 1 (Branch invariance): Let $\alpha' = \alpha + 2\pi m$ and $\alpha'_d = \alpha_d + 2\pi n$ for integers m, n. The physical sensor states, circular tracking error, and pointing loss are invariant under these changes of representatives.

Argument: The identities $R(\alpha') = R(\alpha)$ and $R(\alpha'_d) = R(\alpha_d)$ follow from the 2π periodicity of sine and cosine. Thus $S'^{-1}{}_d S' = S^{-1}{}_d S$, so their SO(2) logarithms are identical; likewise $\cos(\theta+\alpha'-\gamma) = \cos(\theta+\alpha-\gamma)$. Hence all three quantities depend on elements of SO(2), not on a selected real-valued representative. This statement does not mean that the principal logarithm is globally continuous; it means that the controller output is unaffected by an equivalent 2π shift.

***Proposition 2 (Convex online control problem)*:** Suppose $Q \succeq 0$, $Q_f \succeq 0$, and $H > 0$. For a fixed reference sequence and linearized product-error dynamics, the finite-horizon GMPC problem is a convex QP; if the feasible set is nonempty, its control sequence is unique.

Argument: Stacking the predicted errors and controls converts the equality constraints into an affine system. The box actuator limits form a convex polytope. The objective Hessian is positive semidefinite, and the control block contributed by $H > 0$ is positive definite. Therefore the problem is convex and strictly convex in the control sequence. The statement applies to each local QP and does not establish global nonlinear closed-loop stability, since the Lie-algebra dynamics are a first-order approximation about the tracked reference.

Proposition 3 (Base–sensor separation): Under the tested block-diagonal model and diagonal weights, changing only the sensor reference or rate bound leaves the ideal base subproblem unchanged.

Argument: The SO(2) adjoint block is zero and the fourth row and third input column are decoupled from the first three error coordinates. Diagonal Q, Qf , and H introduce no cross term. The finite-horizon objective and constraints therefore separate into base and sensor components. This explains why the measured base RMSE is constant across the sensor experiments. The property would no longer hold if the sensor mass affected base dynamics, the cost included cross terms, the field of view constrained base orientation, or planning jointly modified the base reference.

A. Receding-Horizon Procedure

The complete method has an offline reference stage and an online tracking stage. Offline, the original GMPC generator produces ($X_{d,k}$, $\mu_{d,k}$, $\omega_{d,k}$). A continuous relative-bearing lift is constructed intrinsically by accumulating successive increments $\mathrm{Log}_{SO(2)}(S^{-1}{}_{k-1}\,Sk)$. The smooth circular optimization produces $\widetilde{\alpha}_{d,k};\ S_{d,k} = R(\widetilde{\alpha}_{d,k})$ is stored as the group state and $\nu_{d,k}$ is differentiated from the lift.

Online, the following operations are repeated at each sample:

1) Select the current reference index and construct $Z_{d,k}$ and $u_{d,k}$.
2) Compute $\psi k = \mathrm{Log}_{SE(2)}(X^{-1}{}_{d,k}\,X_k)$ and $\varphi k = \mathrm{Log}_{SO(2)}(S^{-1}{}_{d,k}\,S_k)$.
3) Form the horizon-dependent affine error dynamics using the reference twists.
4) Solve the constrained QP for absolute controls (μ, ω, ν).
5) Apply the first control through the product-group update and repeat at the next sample.

This procedure preserves a continuous physical evolution even when a plotted principal coordinate changes sign at the branch cut. The branch diagnostic used below additionally compares the principal-coordinate step, continuous-lift step, reference rate, and commanded-rate step so that a hidden $2\pi/\Delta t$ artifact cannot be mistaken for valid motion.

VII. NUMERICAL EXPERIMENTS

A. Protocol and Error Definitions

The primary numerical object is the *wheel robot with a sensor*, with state $(X, S) \in SE(2) \times SO(2)$. The *wheel robot,* with state $X \in SE(2)$, is not a competing application setting; it is the sensor-free reduction used to verify that the appended, decoupled sensor subsystem does not alter base tracking. Both are controlled by Product GMPC and Product NMPC on two references, a circle and a butterfly shaped lemniscate. The sample period is $\Delta t = 0.05$ s, the horizon is $T = 10$, and the initial base perturbation is $(-0.03, 0.02, -0.08)$ in (x, y, θ). For the primary sensor model, a fixed landmark at (1.15, 0.85) m defines the desired global pointing direction.

To align the vertical axes with Fig. 6 of [1], we report

$$e_p(k) = \|p_k - p_{d,k}\|_2, \qquad e_R(k) = \left\|\mathrm{Log}_{SO(2)}(R_{d,k}^{-1} R_k)\right\|_2, \tag{7}$$

together with linear velocity μ and angular velocity ω. For the sensor model we additionally report

$$e_s(k) = \left|\mathrm{Log}_{SO(2)}\left(R(\gamma_{d,k})^{-1} R(\theta_k + \alpha_k)\right)\right|. \tag{8}$$

All angular metrics are therefore group logarithms even though Product NMPC optimizes Euler coordinates.

Product NMPC directly augments the reference paper's Euler-angle NMPC state and input with sensor yaw and rate. Its state is $z = [x, y, \theta, \alpha]^\top$ with

$$\dot{x} = \mu \cos\theta, \quad \dot{y} = \mu \sin\theta, \quad \dot{\theta} = \omega, \quad \dot{\alpha} = \nu.$$

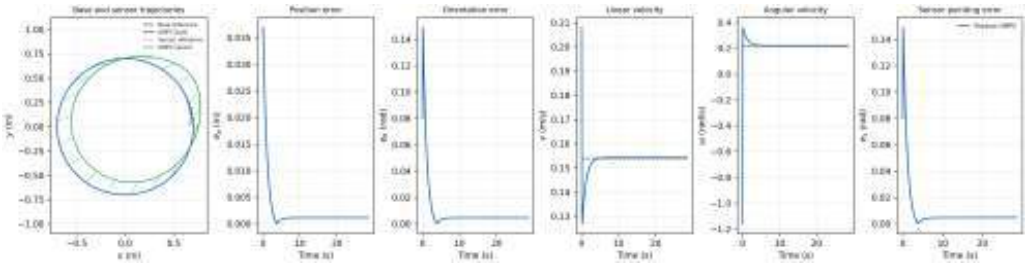

Fig. 2. Product GMPC tracking of the circular reference. The first panel overlays the reference and closed-loop base and sensor-axis endpoint paths; arrows indicate the realized sensor direction. The remaining panels show tracking errors and control inputs.

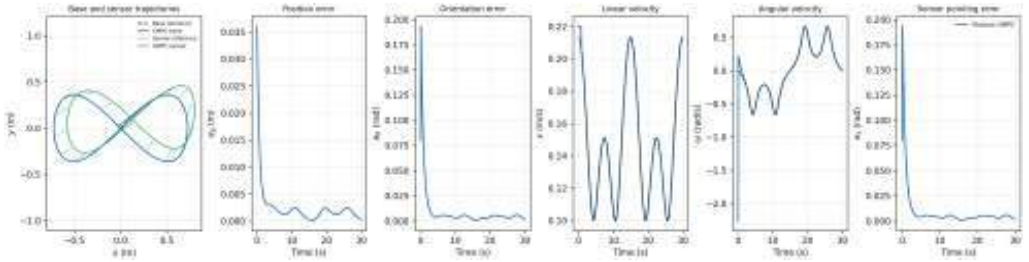

Fig. 3. Product GMPC tracking of the butterfly reference, using the same trajectory, error, and input quantities as Fig. 2.

TABLE I
OVERALL POSITION ERROR ON TWO REFERENCES

Reference	Controller	Overall E_p (m)
Circle	Product GMPC	0.00543
Circle	Product NMPC	0.00576
Butterfly	Product GMPC	0.00475
Butterfly	Product NMPC	0.00501

It retains the reference NMPC's direct transcription and IPOPT solver [5], and uses

the same horizon, reference, sampling period, and input bounds as Product GMPC. Its objective subtracts θ and α as real coordinates. Consequently, the simulation retains continuous lifts of these two coordinates; wrapping the plant state after each integration step would be an implementation error, because it would mix a principal chart with an Euclidean objective.

B. Multiple-Trajectory Comparison

We record one aggregate scalar per run,

$$E_p = \sqrt{\frac{1}{N}\sum_{k=0}^{N-1} e_p(k)^2}, \qquad (9)$$

the position RMSE over the complete trajectory. Figure 1 and Table I show that Product GMPC has slightly lower Ep on both references, while the detailed plots report e_R and e_s.

Figures 2 and 3 present the detailed Product GMPC results for the circular and butterfly references, respectively, including the base and sensor-axis endpoint trajectories, tracking errors, and control inputs. The corresponding Product NMPC results were also obtained but are omitted to save space; they exhibit larger overall position-tracking errors than Product GMPC on both trajectories.

VIII. CONCLUSION

This paper formulated geometric predictive control for a wheeled mobile robot with an actuated directional sensor on SE(2) × SO(2). The model combines nonholonomic base motion with relative sensor orientation, and the local product-algebra prediction produces a convex tracking QP.

A continuous lift is used only for reference differentiation, while group logarithms and periodic losses prevent the principal-angle branch cut from becoming a physical control Discontinuity.

Across the circular and butterfly references, Product GMPC achieves lower overall position RMSE than Product NMPC. The detailed Product GMPC results show close base-trajectory tracking, sensor-axis endpoint tracking, and decreasing position, orientation, and sensor-pointing errors.

These findings are numerical observations, not a proof of global closed-loop stability or hardware robustness. Future work should incorporate sensor acceleration and travel limits, estimator noise, occlusion and field-of-view constraints, and a coupled base–sensor reference planner. Physical experiments will be necessary to determine whether the computational and tracking behavior persists under actuator dynamics and model mismatch.

Works Cited

[1] J. Tang, S. Wu, B. Lan, Y. Dong, Y. Jin, G. Tian, W.-A. Zhang, and L. Shi, “GMPC: Geometric model predictive control for wheeled mobile robot trajectory tracking,” IEEE Robotics and Automation Letters, vol. 9, no. 5, pp. 4822–4829, 2024.

[2] J. Sola, J. Deray, and D. Atchuthan, “A micro Lie theory for state ` estimation in robotics,” arXiv:1812.01537, 2018.

[3] J. T. Betts, Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, 2nd ed. Philadelphia, PA, USA: SIAM, 2010.

[4] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “CasADi: A software framework for nonlinear optimization and optimal control,” Mathematical Programming Computation, vol. 11, pp. 1–36, 2019.

[5] A. Wachter and L. T. Biegler, “On the implementation of an interior point filter line-search algorithm for large-scale nonlinear programming,” Mathematical Programming, vol. 106, pp. 25–57, 2006.

[6] R. M. Murray, Z. Li, and S. S. Sastry, A Mathematical Introduction to Robotic Manipulation. Boca Raton, FL, USA: CRC Press, 1994.

[7] F. Bullo and A. D. Lewis, Geometric Control of Mechanical Systems. New York, NY, USA: Springer, 2005.

[8] G. S. Chirikjian, Stochastic Models, Information Theory, and Lie Groups, vol. 2. Boston, MA, USA: Birkhauser, 2012. ¨

[9] T. D. Barfoot, State Estimation for Robotics. Cambridge, U.K.: Cambridge University Press, 2017.

[10] P.-A. Absil, R. Mahony, and R. Sepulchre, Optimization Algorithms on Matrix Manifolds. Princeton, NJ, USA: Princeton University Press, 2008.

[11] T. Lee, M. Leok, and N. H. McClamroch, “Geometric tracking control of a quadrotor UAV on SE(3),” in Proc. IEEE Conf. Decision and Control, 2010, pp. 5420–5425.

[12] Y. Kanayama, Y. Kimura, F. Miyazaki, and T. Noguchi, “A stable tracking control method for an autonomous mobile robot,” in Proc. IEEE Int. Conf. Robotics and Automation, 1990, pp. 384–389.

[13] C. Samson, “Control of chained systems: Application to path following and time-varying point-stabilization of mobile robots,” IEEE Trans. Autom. Control, vol. 40, no. 1, pp. 64–77, 1995.

[14] P. Morin and C. Samson, “Motion control of wheeled mobile robots,” in Springer Handbook of Robotics, B. Siciliano and O. Khatib, Eds. Berlin, Germany: Springer, 2008, pp. 799–826.

[15] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, “Constrained model predictive control: Stability and optimality,” Automatica, vol. 36, no. 6, pp. 789–814, 2000.

[16] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, Model Predictive Control: Theory, Computation, and Design, 2nd ed. Madison, WI, USA: Nob Hill Publishing, 2017.

[17] L. Grune and J. Pannek, Nonlinear Model Predictive Control: Theory and Algorithms, 2nd ed. Cham, Switzerland: Springer, 2017.
[18] M. Diehl, H. G. Bock, and J. P. Schloder, “A real-time iteration ¨scheme for nonlinear optimization in optimal feedback control,” SIAM J. Control Optim., vol. 43, no. 5, pp. 1714–1736, 2005.
[19] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, “qpOASES: A parametric active-set algorithm for quadratic programming,” Math. Program. Comput., vol. 6, pp. 327–363, 2014.
[20] G. Lu, W. Xu, and F. Zhang, “On-manifold model predictive control for trajectory tracking on robotic systems,” IEEE Trans. Ind. Electron., vol. 70, no. 9, pp. 9192–9202, 2023.

Exploring CRISPR/Cas9 as a Potential Gene Therapy for Tay-Sachs Disease

By Lydia Xia

Abstract

Tay-Sachs disease is a rare genetic disorder that progressively destroys nerve cells in the brain and spinal cord. Sandhoff disease shares similar characteristics with those of Tay-Sachs, but differs in the mutations that distinguish them. Both disorders are neurodegenerative, meaning they involve progressively deteriorating cells in the nervous system. Tay-Sachs and Sandhoff are diseases caused by mutations in the HEXA and HEXB genes, respectively, leading to a buildup of GM2 gangliosides in neurons. Tay-Sachs disease is fatal; those with infantile Tay-Sachs typically do not survive past the age of 5. Currently, treatment focuses on managing symptoms as there is no cure. Using CRISPR, we can explore the possibility of correcting these mutations directly at the source.

Symptoms

Infantile Tay-Sachs disease is the most severe disorder of the nervous system (Kornfeld) and can present differently between individuals (Fernandes Filho and Shapiro). The infant typically does not exhibit symptoms at birth, but symptoms may develop within 6 months, with some cases presenting as early as 1 week (Lui et al.). Common symptoms include seizures, general weakness, and, most notably, a yellow spot in the eyes with large grey/white patches and a brown/red center, most distinct in the fovea centralis (Fernandes Filho and Shapiro). Although an initial examination reveals normal optic nerves, upon further examination, these inconsistencies in the eyes are seen (Fernandes Filho and Shapiro). These ocular inconsistencies initially lead to blindness, but are not lethal, as revealed by a pathological examination of the optic nerves (Fernandes Filho and Shapiro). Hearing and touch appear sensitive to the smallest of sounds and touch startling patients (Fernandes Filho and Shapiro). The muscles of the head, neck, and back are incredibly weak and therefore cannot support movements such as holding the head straight or sitting upright (Fernandes Filho and Shapiro). Juvenile Tay-Sachs manifests in early childhood between ages 2-10, while adult Tay-Sachs is incredibly rare, with diagnosis at 18 and older (Lui et al.).

Diagnosis

Diagnosis of Tay-Sachs disease starts with clinical findings of progressive weakness with developmental delay, seizures, red spots in the eyes, or hyperreflexia (Lui et al.). Patients presenting with these symptoms should be further evaluated for gangliosidosis involving hexosaminidase A (HEXA) and hexosaminidase B (HEXB) by measuring serum levels (Lui et al.). Individuals with infantile Tay-Sachs disease exhibit either no or extremely low enzyme activity (0%-5%) accompanied by normal or elevated HEXB levels (Lui et al.).

Additionally, if this testing reveals reduced enzyme activity and the patient has an Ashkenazi or French Canadian background, targeted gene testing should be conducted (Lui et al.). Tay-Sachs disease is more frequent in the Ashkenazi Jewish heritage (Lui et al.).

Treatment

Currently, treatment for Tay-Sachs disease focuses on symptom relief (Kornfeld). In the developmental stage, gene therapies using adeno-viruses deliver cDNA encoding **α** and **β** HEXA subunit genes (Kornfeld). The effectiveness of this therapy is evaluated in **α** or **β** HEXA subunit-defective model mice or Jacob sheep, which echibit the same pathological features as those in humans (Kornfeld). However, the limitation of adeno-based vectors is their ability to carry constructs (Lui et al.).

Genetics

Lysosomes are acidified cytoplasmic organelles containing hydrolytic enzymes that degrade most biological macromolecules (Solovyeva et al.). Lysosomes are involved in the intracellular components of bone tissue. Lysosomal storage diseases are heterogeneous diseases that involve the deficiency of specific lysosomal enzymes, which leads to an accumulation of the missing enzyme's substrate (Solovyeva et al.; Xiao et al.). Tay-Sachs disease is a lysosomal storage disorder characterized by the acculumation of GM2 gangliosides (Lui et al.).

Molecular Basis

There are over 75 mutations of the **α**-subunit of HEXA deficiency and GM2 gangliosidosis that are responsible for Tay-Sachs disease (Fernandes Filho and Shapiro). These mutations result in significant variations in enzyme activity data, the extent and distribution of glycolipid accumulation in the brain and spinal cord, and a wide range of clinical phenotypes (Fernandes Filho and Shapiro). These mutations include single-gene deletions, substitutions, insertions, splicing alterations, duplications, and complex gene rearrangements (Lui et al.).

CRISPR-Based Gene Editing

Editing Mechanism

Clustered regularly interspaced short palindromic repeats (CRISPR/Cas9) is a gene-editing tool that can be used to repair defective genomes by turning specific genes on or off (Redman et al.). CRISPR/Cas9 has been used in homozygous autosomal recessive disorders, such as cystic fibrosis, and in X-linked recessive disorders, including Duchenne muscular dystrophy (Wojtal et al.). CRISPR/Cas9 is a remarkably simple and versatile system used in various biological applications (Lino et al.). CRISPR/Cas9 utilizes a guide RNA to target a specific gene, and then the Cas9 protein causes a double-stranded DNA break, allowing the genome to be modified (Redman et al.). This is all directed by a synthetic single-guide RNA (sgRNA) designed to bind to your target gene (Redman et al.). This sgRNA is typically intended to be about 18–20 nucleotides long to enhance the specificity of targeting said gene (Redman et al.). However, this occurs in the absence of a homologous DNA template, which can lead to mutations and deletions in the repaired strand (Lino et al.). CRISPR/Cas9 can be classified into two categories: class one uses multi-subunit protein complexes to cleave nucleic acids, while class two uses one large protein for recognition and

cleavage (Tao et al.).

Ethical Dilemmas

There are some ethical dilemmas associated with CRISPR/Cas9 that need to be addressed, including its overall safety (Lorenzo et al.). In the case of interventions on germline cells, the offspring will inherit the modified DNA (Lorenzo et al.). Because future generations directly inherit CRISPR edits, ethical concerns must be considered relating to who is responsible if there is an issue with the manipulated gene (Brokowski and Adli). Some believe that the use of CRISPR will cause an individual or group to become genetically superior to others (Shinwari et al.). Because of these concerns of eugenics and human safety, a moratorium was called in Napa, California, to discuss the ethical and social aspects of CRISPR (Shinwari et al.; Sheridan).

Challenges and Limitations

There is a challenge concerning the delivery of gene editing to the right cells in vivo (Redman et al.). In order to safely deliver Cas9-nuclease-encoding genes and the correct guide RNAs while avoiding harmful effects, a vector is needed (Redman et al.). Many non-viral delivery systems are currently being researched to identify a suitable vector (Redman et al.). Additional concerns include the possibility of limited on-target editing efficiency, incomplete editing, and inaccuracies with target editing (Brokowski and Adli; Piergentili et al.). Off-target mutations in the genome can be difficult to detect because the number and position of the mutation are unknown (Zischewski et al.). These concerns will eventually be diminished over time as CRISPR technology continues to improve (Piergentili et al.).

CRISPR-Based Therapeutic Approaches for Tay-Sachs Disease

Current Research

Adeno-associated viruses (AAVs) can be used to deliver CRISPR/Cas9 to replace defective HEXA and HEXB genes with HEXM (Josefano et al.). HEXM is a modified gene that has the gene codes for the HEX subunit of the same enzyme missing in Tay-Sachs disease (Josefano et al.). To treat mice with Sandhoff disease, a HEXM subunit was applied to a PS813 gene-editing system (Ou et al.). Using AAV delivery, GM2 ganglioside levels were reduced to normal levels in these mice, excluding the brain (Josefano et al.). This provides hope for the future, with the potential to develop in vivo gene editing for those with Sandhoff disease and Tay-Sachs disease (Ou et al.).

Methods

Data Collection

Healthy human HEXA gene and HEXA protein sequences were retrieved from the NCBI Gene database. The known disease-causing variants of the HEXA gene were retrieved from NCBI ClinVar, a trusted, publicly available database that contains clinically observed mutations.

Sequence Alignments

The HEXA variants were input into Vector Builder and translated into their corresponding amino acid sequence. Translation was carried out across three open reading frames to account for potential frameshift mutations that could completely alter the reading frames. This step is crucial to ensure that in-frame and shifted translations are correctly captured. The resulting open reading frames and the healthy reference protein sequences were then aligned using Clustal Omega, a multiple sequence alignment tool from EMBL-EBI. Clustal Omega alignments were then exported and input into ESPRipt where we visualized them to identify both conserved and divergent regions. Sequences were then directly compared to determine where specific mutations alter the primary protein structure.

3D Protein Modeling

The sequences were submitted to AlphaFold to create its projected 3D protein model. AlphaFold can distinguish model confidence of the per-residue local confidence measure.

Results

First, we aligned the healthy HEXA gene with six variants from live patients to examine differences in the gene sequence. Alignments showed that the healthy copy is 317 amino acids longer than D, 319 amino acids longer than E, 318 amino acids longer than F, 326 amino acids longer than B, 336 amino acids longer than A, and 334 amino acids longer than C. When translated into protein, you can see that the proteins are significantly affected, with the healthy HEXA copy being 995 longer than F, 997 longer than E, 998 longer than D, 1044 longer than C, 1031 longer than B, and 1050 longer than A. This effect is as expected, given the shortened gene sequence already observed. In F, we see a similar effect to the other mutations, which may seem incorrect due to the nature of frameshift mutations. However, three reading-frame output predictions where only one aligned this way, with the other two having unsuccessful alignments. This tracks because these are likely frameshift mutations that don't align at all. To compare protein misfolding, we used AlphaFold and Swiss Model software to analyze the amino acid sequences, and only the healthy HEXA sequence could be predicted. None of the diseased mutated variants could have been predicted. This is not surprising, given the large portions missing. We attempted two different programs, and neither could predict a protein for either diseased sample, even at low confidence.

A

B

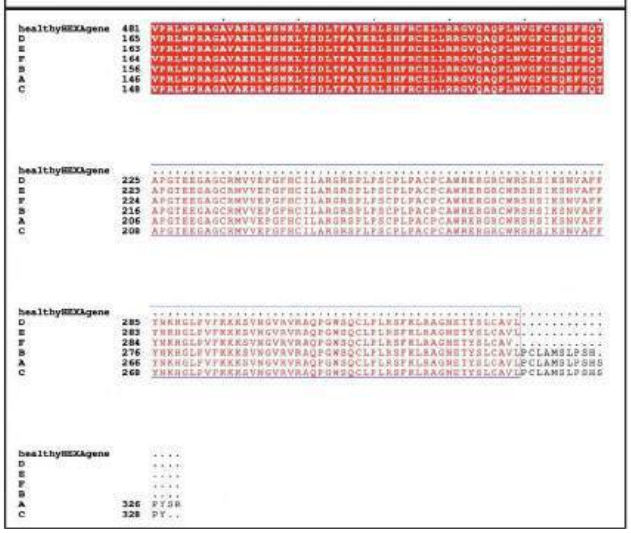

C

Sequence ID	Classification	Mutation
HealthyHEXAgene	Healthy	N/A
A	Likely benign	Synonymous
B	Pathogenic	Nonsense
C	Uncertain	Missense
D	Pathogenic	Splice acceptor
E	Benign	Intron Variant
F	Pathogenic	Frameshift

Figure 1.

Chart A. 3’ end of FASTA format alignment from ESPript derived from Clustal Omega of HEXA DNA sequences. Highlighted red represents 100% alignment of all sequences. Red letters surrounded by a blue box represent areas where it aligns, but not completely. The dots represent areas where the sequence has not started.

Chart B. 5’ end of FASTA format alignment from ESPript derived from Clustal Omega of HEXA DNA sequences. Highlighted red represents 100% alignment with all sequences. Red letters surrounded by a blue box represent areas where it aligns, but not completely. The dots represent areas where the sequence has already ended.

Chart C. Table detailing which mutations and classifications the respective letters correspond to that were used in the alignments.

A

B

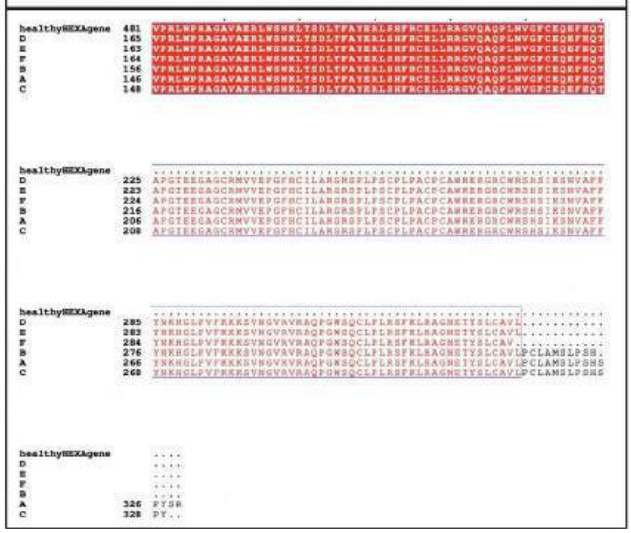

C

Sequence ID	Classification	Mutation
HealthyHEXA gene	Healthy	N/A
A	Likely benign	Synonymous
B	Pathogenic	Nonsense
C	Uncertain	Missense
D	Pathogenic	Splice acceptor
E	Benign	Intron Variant
F	Pathogenic	Frameshift

Figure 2.

Chart A. 3’ end of FASTA format alignment from ESPript derived from Clustal Omega of HEXA protein sequences. Highlighted red represents 100% alignment of all sequences. Red letters surrounded by a blue box represent areas where it aligns, but not completely. The dots represent areas where the sequence has not started.

Chart B. 5’ end of FASTA format alignment from ESPript derived from Clustal Omega of HEXA protein sequences. Highlighted red represents 100% alignment with all sequences. Red letters surrounded by a blue box represent areas where it aligns, but not completely. The dots represent areas where the sequence has already ended.

Chart C. Table detailing which mutations and classifications the respective letters correspond to that were used in the alignments.

```
>protein sequence, frame 1
EVSSVFPDFYLHLGGDEVDFTCWKSNPEIQDFMRKKGFGEDFKQLESFYIQTLLDIVSSYGKGYVVWQEVFDNKVKIQPDTIIQVWREDIPVNY

>protein sequence, frame 2
GKSALSSQIFIFILEEMRLISPAGSPTQRSRTL*GRKASVRTSSSWSPSTSRRCWTSSLLMARAMWCGRRCLIIK*RFSQTQSYRCGERIFQ*T

>protein sequence, frame 3
GASQLCLPRFLSSSWRR*G*FHLLEVQPRDPGLYEEERLR*GLQAAGVLLHPDAAGHRLFLWQGLCGVAGGV***SKDSARHNHTGVARGYSSE
```

Figure 3. Three frames of protein sequences derived from VectorBuilder of pathogenic

frameshift mutation (F). Red asterisks indicate deletions; yellow-highlighted parts indicate potential errors.

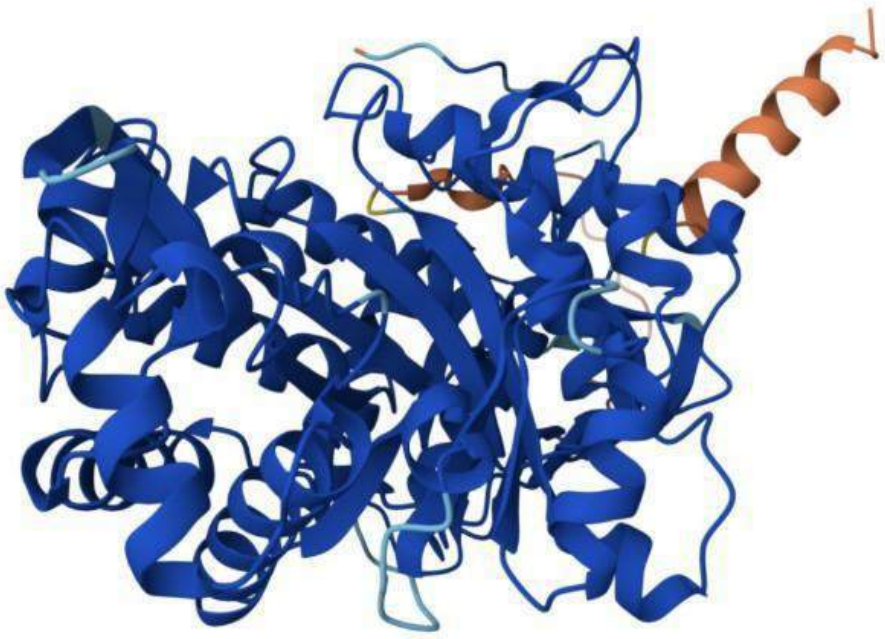

Figure 4. 3D model of a healthy HEXA protein. Derived using protein sequence from NCBI and AlphaFold technology.

Discussion

The healthy HEXA sequence represents the wild-type and serves as the baseline for comparison. The available disease variants include a range of mutation types, such as frameshift mutations, which alter the reading frame of the gene; nonsense mutations, which introduce a premature stop codon; missense mutations, which substitute one amino acid for another; and synonymous mutations, where the amino acid sequence remains unchanged but splicing or expression may be affected. Additional variants include intron mutations and splice acceptor mutations that disrupt mRNA processing.

In ClinVar, each disease variant is classified by clinical significance into likely pathogenic, likely benign, pathogenic, benign, or uncertain significance. These classifications are based on experimental evidence and clinical case reports. By retrieving healthy reference sequences and a large range of documented disease variants, a strong foundation is created to compare sequence alterations and further analyze their potential protein structure and functional consequences.

In VectorBuilder, the DNA sequences that were derived from ClinVar were translated into protein sequences. There were three reading frames given; out of the three, only one matched for each sequence. Those aligned sequences are the sequences used in the subsequent steps. This could explain why the frameshift mutation still aligns with both the

healthy copy and other mutated variants.

In Clustal Omega, the open reading frames and the healthy reference protein sequences are aligned to identify both conserved and divergent regions. This provides a visual alignment that highlights insertions, deletions, and substitutions. By directly comparing healthy and diseased variants, this analysis reveals how and where specific mutations alter the primary protein structure. This alignment data will be further used to study how and why the functional loss of the HEXA protein can manifest as Tay-Sachs disease.

Due to nonsense and intron variant mutations, the start codon was likely uncoded for until much later than it typically is. As shown in Figure 2, the three sequences begin to align when the healthy HEXA sequence reaches amino acid 1021. However, the two diseased sequences have just started at amino acid 26 and amino acid 1, indicating a large chunk missing in the 3 prime end. This also occurs on the 5 prime end, with the two diseased sequences having an early stop codon while the healthy sequence continues. In Figure 1, the two diseased sequences have a large chunk missing from the 3' end and an excess at the 5' end, whereas the healthy sequence has already terminated. We see that in the two diseased sequences, nearly all of the sequence aligns with the healthy sequence, even though these sequences are much shorter in length in comparison to the healthy one.

AlphaFold and Swiss Model programs were unable to generate a predicted 3D model for the diseased sequences. This indicates that the mutations were too impactful and prevented these proteins from folding correctly. Protein misfolding can happen when errors in the amino acid sequence disrupt the folding process, often resulting in nonfunctional or toxic proteins that are targeted for degradation by cellular quality control systems (Gregersen). Missing or altered amino acids can prevent proper folding and can also eliminate essential binding domains, leading to loss of function (Cheung and Deber). Mutations in HEXA can cause misfolding of the α subunit in the endoplasmic reticulum, leading to its degradation and resulting in Tay–Sachs disease due to reduced enzyme activity (Dersh et al.).

Works Cited

Brokowski, Carolyn, and Mazhar Adli. “CRISPR Ethics: Moral Considerations for Applications of a Powerful Tool.” Journal of Molecular Biology, CRISPR: from the basic biology to its technological applications, vol. 431, no. 1, Jan. 2019, pp. 88–101, https://doi.org/10.1016/j.jmb.2018.05.044.

Cheung, Joanne C., and Charles M. Deber. “Misfolding of the Cystic Fibrosis Transmembrane Conductance Regulator and Disease.” Biochemistry, vol. 47, no. 6, Feb. 2008, pp. 1465–73. ACS Publications, https://doi.org/10.1021/bi702209s.

Dersh, Devin, et al. “Tay–Sachs Disease Mutations in HEXA Target the α Chain of Hexosaminidase A to Endoplasmic Reticulum–Associated Degradation.” Molecular Biology of the Cell, vol. 27, no. 24, Dec. 2016, pp. 3813–27. https://doi.org/10.1091/mbc.E16-01-0012.

Fernandes Filho, Jose Americo, and Barbara E. Shapiro. “Tay-Sachs Disease.” Archives of Neurology, vol. 61, no. 9, Sept. 2004, pp. 1466–68. https://doi.org/10.1001/archneur.61.9.1466.

Gregersen, N. “Protein Misfolding Disorders: Pathogenesis and Intervention.” Journal of Inherited Metabolic Disease, vol. 29, nos. 2–3, 2006, pp. 456–70. https://doi.org/10.1007/s10545-006-0301-4.

Josefano, Richelle Bertly, et al. “CRISPR-Cas9 Mediated Gene Therapy: Current Advancements and Applications Towards Tay-Sachs and Sandhoff Disease.” Indonesian Journal of Life Sciences, Sept. 2024, pp. 69–89. https://doi.org/10.54250/ijls.v6i2.195.

Kornfeld, Stuart. “Trafficking of Lysosomal Enzymes.” The FASEB Journal, vol. 1, no. 6, 1987, pp. 462–68. Wiley Online Library, https://doi.org/10.1096/fasebj.1.6.3315809.

Lino, Christopher A., et al. “Delivering CRISPR: A Review of the Challenges and Approaches.” Drug Delivery, vol. 25, no. 1, Jan. 2018, pp. 1234–57. Taylor and Francis. https://doi.org/10.1080/10717544.2018.1474964.

Lorenzo, David, et al. “Ethics and Genomic Editing Using the Crispr-Cas9 Technique: Challenges and Conflicts.” NanoEthics, vol. 16, no. 3, Dec. 2022, pp. 313–21. https://doi.org/10.1007/s11569-022-00425-y.

Lui, Forshing, et al. “Tay-Sachs Disease.” StatPearls, StatPearls Publishing, 2026, http://www.ncbi.nlm.nih.gov/books/NBK564432/.

Ou, Li, et al. “A Novel Gene Editing System to Treat Both Tay–Sachs and Sandhoff Diseases.” Gene Therapy, vol. 27, no. 5, May 2020, pp. 226–36. https://doi.org/10.1038/s41434-019-0120-5.

Piergentili, Roberto, et al. “CRISPR-Cas and Its Wide-Ranging Applications: From Human Genome Editing to Environmental Implications, Technical Limitations, Hazards and Bioethical Issues.” Cells, vol. 10, no. 5, May 2021, p. 969. https://doi.org/10.3390/cells10050969.

Redman, Melody, et al. “What Is CRISPR/Cas9?” Archives of Disease in Childhood - Education and Practice, vol. 101, no. 4, Aug. 2016, pp. 213–15. Research in Practice.

https://doi.org/10.1136/archdischild-2016-310459.
Sheridan, Cormac. “CRISPR Germline Editing Reverberates through the Biotech Industry.” Nature Biotechnology, vol. 33, no. 5, May 2015, pp. 431–32. https://doi.org/10.1038/nbt0515-431.
Shinwari, Zabta Khan, et al. “Ethical Issues Regarding CRISPR Mediated Genome Editing.” Current Issues in Molecular Biology, 2018, pp. 103–10. https://doi.org/10.21775/cimb.026.103.
Solovyeva, Valeriya V., et al. “New Approaches to Tay-Sachs Disease Therapy.” Frontiers in Physiology, vol. 9, Nov. 2018. https://doi.org/10.3389/fphys.2018.01663.
Tao, Shuan, et al. “The Application of the CRISPR-Cas System in Antibiotic Resistance.” Infection and Drug Resistance, vol. 15, Aug. 2022, pp. 4155–68. https://doi.org/10.2147/IDR.S370869.
Wojtal, Daria, et al. “Spell Checking Nature: Versatility of CRISPR/Cas9 for Developing Treatments for Inherited Disorders.” The American Journal of Human Genetics, vol. 98, no. 1, Jan. 2016, pp. 90–101. https://doi.org/10.1016/j.ajhg.2015.11.012.
Xiao, Changrui, et al. “Adults with Lysosomal Storage Diseases in the Undiagnosed Diseases Network.” Molecular Genetics & Genomic Medicine, vol. 10, no. 9, 2022, p. e2013. https://doi.org/10.1002/mgg3.2013.
Zischewski, Julia, et al. “Detection of On-Target and off-Target Mutations Generated by CRISPR/Cas9 and Other Sequence-Specific Nucleases.” Biotechnology Advances, vol. 35, no. 1, Jan. 2017, pp. 95–104, https://doi.org/10.1016/j.biotechadv.2016.12.003.

Progress and limitations of iPSC-derived Brain Organoids in Neurological Disease Research By Sydney Kroker McCormack

Abstract

Brain diseases have historically been difficult to study, as researchers lacked the tools necessary to study human brain cells and their malfunctions in disease conditions. Through stem cell-derived brain organoids, researchers were able to closely investigate many neurological disorders, as these 3D models mimic key features of the brain. What is the current state of stem-cell-derived organoids as models for brain disease? Here, we report on the methodological approaches currently used in brain organoid research, commonly involving reprogramming human cells into induced pluripotent stem cells (iPSCs) and turning them into 3D brain models using targeted growth factors. Many of these studies used CRISPR gene editing as well to create control lines, isolating the effect of the disease-causing mutation by comparing it to the healthy, otherwise identical version of the cells. This review argues that stem cell-derived brain organoids have transformed the study of neurological conditions, allowing researchers to model diseases that were once nearly impossible to study directly in human tissue. However, this review will also highlight that the field must address limitations, such as reproducibility, vascularization, and sample size, to ensure reliable results. Brain organoids are quickly becoming one of the key tools scientists use to study diseases ranging from rare genetic disorders to neurodegenerative conditions, so understanding how well they work and where they fall short is critical. This review brings that research together to help identify where the field should focus to improve these models and what progress has been made.

Introduction

The field of neurological disease research faces the challenge of understanding brain conditions, which requires studying human brain tissue, as the brain is largely inaccessible for direct study during life. By using stem cell-derived organoids, researchers can now study brain diseases in new ways. Before 3D brain organoids, researchers relied on animal models and older cell cultures that did not fully represent human biology, making it difficult to accurately predict how treatments would work in human brains. Animal models remain an important tool in neuroscience available to researchers and continue to be used with newer tools as well.

Stem cell-derived organoids closely resemble human brain tissue, opening up a new way to study brain development and disease.[1] They allow diseases affecting the human brain to be studied through models that replicate real brain tissue. Organoids develop multiple cell types and form bits of the cellular organization in brain tissue, unlike traditional two-dimensional cell cultures. This technology advanced further when stem cells shifted from hESCs, which required destroying embryos, to iPSCs, which are reprogrammed from a patient's cells, making the research personalized with fewer ethical concerns.[2]iPSCs are specialized adult cells (usually skin or blood) that are reprogrammed into stem cells that can then become any type of cell, reprogrammed through resetting the genetic memory.[2] This review investigates the current state

of stem-cell derived organoids as models for brain disease. Although brain organoids have transformed the way neurodegenerative diseases are studied, recent research shows they struggle to accurately model complex diseases due to unreliable study design, small sample sizes, and variable results.[3,4] Recent studies have utilized brain organoids to model specific conditions such as microcephaly, Timothy syndrome, and autism spectrum disorder.[5] This paper reviews how stem cell-derived brain organoids have enabled deeper, more direct study of neurological diseases that was not possible before. However, numerous constraints still need to be addressed in the use of brain organoids to successfully study complex diseases with accuracy and reliability.

This review first walks through the methods used to create and study diseases in these stem cell-derived brain organoids, including CRISPR gene editing, iPSC reprogramming, and organoid growth techniques, revealing how these organoids have evolved into what they are now. Next, this review examines how these brain organoids have been used to research genetic, neurodegenerative, and psychiatric diseases, demonstrating the impact they've made on disease research and revealing their applicable diversity. Finally, this review focuses on the current limitations of this brain organoid research, including small sample size and reliability, discussing possible solutions researchers have proposed that the field should focus on moving forward.

Technical Background

The technical background section explains how stem cell-derived organoids are studied and created through key techniques. These include iPSC reprogramming, 3D differentiation, and CRISPR/Cas 9 gene editing, which provide the groundwork for understanding their later use in disease research.

iPSC reprogramming & hESC→iPSC shift

Researchers must use specific stem cell technologies in the process of selecting the right stem cell source, guiding those stem cells into 3D organoid growth, and gene editing them to model specific diseases. Each of these steps matters individually, starting with choosing the stem cell source.

Stem cells were originally sourced as hESCs, but due to ethical concerns, research shifted to iPSCs, which solved those issues by using a patient's own cells instead of destroying an embryo. hESCs are pluripotent, unspecialized cells derived from early human embryos. They were first isolated many years ago, before reaching their first clinical milestone in 2010. This breakthrough was the first patient treatment using hESC therapy at the Shepherd Center in Atlanta.[2] Using hESCs for research is a controversial topic, as they are derived through a process that destroys live embryos. Many argue that hESCs are needed to help save the lives of those suffering from diseases, while others argue to protect the lives of early embryos. Due to this tension, many places using them are under tight regulation, while some countries even ban hESCs in research entirely.[2]iPSCs emerged from this ethical controversy as they became

appealing because they don't require embryos at all. Much of the field began shifting toward this new stem cell method for this reason, as it allows researchers to continue disease research under less controversial circumstances.[2] iPSCs have many upsides, such as the fact that they are reprogrammed from an adult human cell, usually skin or blood, rather than an embryo, meaning there is no embryonic destruction involved; instead, the stem cells better match the patient genetically, as they are derived from the patient themselves. This allows more personalized disease modeling as well as less of a chance of immune rejection when used therapeutically, making them a less controversial alternative.[2] Reprogramming into iPSCs works through a process of introducing genes (transcription factors) that control which other genes are turned on and off in adult cells. These original four factors include Oct4, Sox2, Klf4, and c-Myc. By using these transcription factors, the cell's gene expression program is reset and pushed back into pluripotency, where they become embryonic-like.[6] Although iPSCs involve no ethical concerns, there were still some areas where researchers were worried. This included the fact that reprogramming involved major epigenetic changes that led to doubt about whether iPSC-derived neurons would be reliable disease models. Researchers compared the iPSC-derived neurons to hESC-derived neurons that were isogenic, meaning genetically identical. Using gene expression data and genome-wide DNA methylation analysis, as well as a narrowed look at the genes linked to neurodegenerative disease, researchers compared these stem cell neurons. The result of this comparison was that the gene expression and DNA methylation patterns were very similar to each other in both neurons. This included both the genome-wide and individual gene levels. While this research demonstrates that iPSC-derived neurons closely resemble hESC ones, continued studies are needed to confirm whether these stem cell-derived neurons match those of a real brain.[7] With a more ethically sourced and reliable stem cell, researchers needed a way to grow these cells into brain-resembling 3D structures for studies.

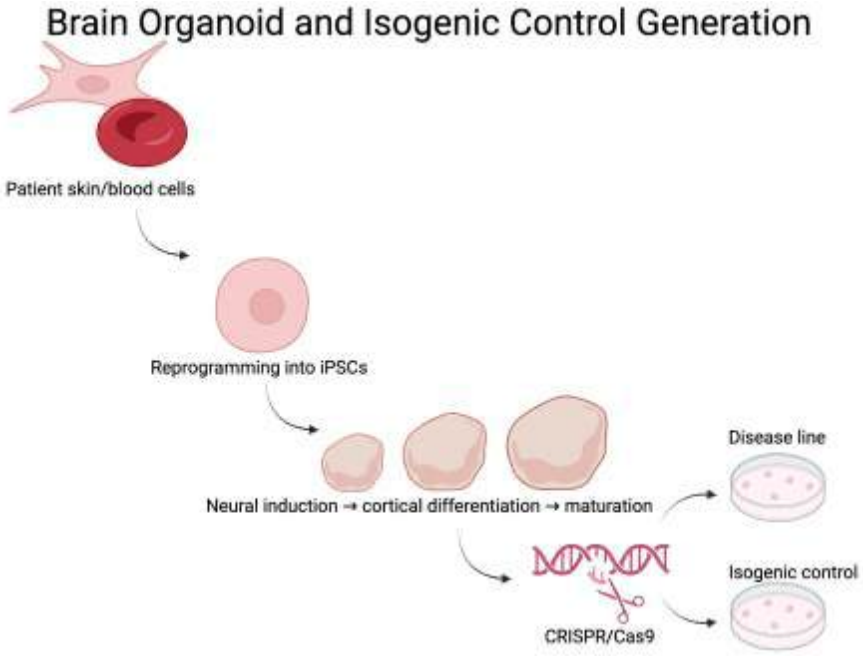

Fig 1. Brain Organoid and Isogenic Control Generation

Patient-derived skin or blood cells are reprogrammed into iPSCs, first guided through neural induction, then cortical differentiation, and finally maturation, where the brain organoids are generated. CRISPR/Cas 9 genome editing is then used to create an isogenic control line, allowing the disease-mutation line to be compared to the genetically matched control.

Growing organoids / 3D differentiation

After iPSCs are established, they can be grown into brain-like structures called organoids through a process called 3D differentiation. To grow the iPSCs into organoids, the protocol first includes growing them into small 3D clumps, embryoid bodies, as the starting point for formation. About nine thousand cells per well are seeded into small V-bottom wells, where they join into an embryoid body. The process then continues through neural induction, cortical differentiation, and maturation. For neural induction, once the embryoid bodies have formed, they are treated with chemical signals that stop them from becoming other cell types and push them toward becoming brain and neural cells instead. Then, during cortical differentiation, the cells are pushed further toward a cortex identity; they start forming radial glia, a type of neural stem cell, along with early progenitor cells. Finally, over 70 or more days, the cells mature into distinct layers of neurons and supporting astrocytes, ending up looking much like a real developing cortex.[8] These cells begin to spontaneously organize themselves into layered structures without needing outside help, similar to the natural layering in cortical tissue during brain development.[9] When researchers compared the genes switched on in the brain organoids to those active in a real fetal brain, they found a close match. This meant that at a genetic level, the organoids were not just resembling a developing brain, but actually behaving like one.[10] Their gene expression pattern directs the brain to form the cerebral cortex rather than any one cortical area.[11] Their stage of development matched that of a real human fetal brain at 8–10 weeks of pregnancy.[9] Not only do these organoids accurately model a developing frontal cortex, but they can also be tools to study specific disease-causing mutations through CRISPR.

CRISPR & isogenic controls

Subsequently, once organoids are grown, researchers must edit their genes to model a specific disease. CRISPR/Cas9 is a tool used to edit DNA. It works like a pair of molecular scissors, cutting DNA at a specific spot to disable or fix a gene.[12]In the Allende study, researchers used skin cells from a patient with infantile Sandhoff disease, a rare genetic disorder that destroys brain and spinal cord nerve cells. These skin cells were reprogrammed into iPSCs, which were found to carry a mutation in a gene called HEXB, the gene responsible for the disease. Using CRISPR, researchers were able to correct the mutated gene in some of the cells.[12] Isogenic controls are made by using CRISPR/Cas9 to edit iPSCs, creating two versions of the same cell line that are identical except for the disease-causing mutation. As they share the same genetic background, any differences researchers see between the two versions can be traced back to that single mutation, rather than other unrelated genetic differences.[13]Isogenic controls

are key because when comparing a sick patient's cells to those from a completely different healthy person, any differences observed could be caused by the disease mutation, or just by normal genetic variation between two different people. This uncertainty creates a research problem that isogenic controls are able to solve. Since both versions come from the exact same person's cells, any difference is likely to be caused by the mutation, rather than genetic noise from two completely different people. This leads to much more reliable and disease-specific results.[13] Researchers now have a precise way to study how a gene mutation affects brain development and function due to isogenic controls.

Overall, iPSC reprogramming, 3D organoid growth, and CRISPR/Cas9 gene editing build up what is needed to construct brain models from a single patient's cells. Together, these core tools create personalized, powerful disease models for patients. With these models built, researchers can now use them to study a wide range of neurological diseases.

Disease-Specific Applications

This section reviews specific neurological diseases that have been modeled using brain organoids, showing the range and impact of this technology.

Having covered how brain organoids are grown, we can now understand how they help reveal insights that would otherwise be impossible to study in the live human brain. This section covers specific diseases modeled using this technology, including genetic and storage diseases, neurodegenerative diseases, and psychiatric disorders. In each category, a specific disease is examined in depth, starting with Sandhoff disease, then Alzheimer's disease, and finally Schizophrenia. Together, these examples showcase the wide range of capabilities that organoid technology offers, as well as its impact on many types of brain disorders.

Genetic/storage diseases

Brain organoids have proven capable of modeling genetic and storage diseases, which are caused by inherited gene mutations that disrupt normal cell function. One disease that has been closely studied using these organoids is Sandhoff disease. Brain cells are heavily affected by this because they are packed with gangliosides, which are crucial for healthy nerve cell membranes and signaling. Because of this, GM2 accumulates there at a much faster rate than in other cells, explaining why this disease destroys the nervous system. The most common form of Sandhoff disease is the infantile form, which is often fatal by early childhood. Its symptoms include loss of motor skills, muscle weakness, and seizures. Currently, there is no approved therapy for GM2 gangliosidoses.[14] Due to there not being a cure for Sandhoff disease, organoids became a way to study how the disorder damages the brain as it progresses quickly and severely. Allende et al. used this approach, growing stem cell-derived organoids from a patient with Sandhoff disease. Researchers compared the Sandhoff disease brain organoids to isogenic controls that were edited by CRISPR/Cas9 to repair the HEXB mutation. This allowed them to isolate the effects of the mutation. The Sandhoff organoids were found to be significantly larger than isogenic HEXB-corrected organoids, measured through a cross-sectional perimeter. This

enlargement was caused by increased cell proliferation. This meant the cells were dividing much more than necessary. The organoids also showed impaired neurodifferentiation alongside this overgrowth, meaning the iPSCs struggled to mature into functioning neurons. As early as week four of organoid growth, researchers used immunostaining to detect GM2 ganglioside buildup. This demonstrated that cell damage begins early in development. Before a real Sandhoff patient would display symptoms, this damage was already detectable in the brain organoids, suggesting that neurological damage could be happening undetected, even before a diagnosis is made. Organoids with Sandhoff disease at week 4 were injected with AAVrh8 vectors encoding HEXA and HEXB, then 11 days later harvested and measured for both size and β-hexosaminidase activity.[12] This silent damage reflects the idea that significant damage already has occurred by the time of the diagnosis. Using organoids to detect this harm highlights that they can reveal disease processes that would otherwise be hidden until toxic damage. Not only can organoids model diseases such as Sandhoff disease, present from birth, but also ones that develop later in life. This includes neurodegenerative diseases such as Alzheimer's.

Neurodegenerative disease (Alzheimer's)

Brain organoids are an impactful resource for modeling neurodegenerative diseases as well. These are conditions where nerve cells break down over time, often appearing later in life. Alzheimer's disease, one of these conditions, has been heavily researched through this technology. Alzheimer's is a progressive neurodegenerative disease. Over time, the brain's function gradually declines. Alzheimer's has two hallmark features, including amyloid beta plaques and neurofibrillary tangles. Amyloid beta plaques are sticky protein clusters that accumulate between neurons. Neurofibrillary tangles are made of twisted tau protein fibers, which build up inside the neurons. These features cause damage by disrupting brain cell communication over time. This disruption begins in the entorhinal cortex and later spreads to the hippocampus, which are responsible for memory and cognitive function, leading to widespread neuronal death. This neurodegeneration in these crucial brain areas leads to patients experiencing confusion, worsening memory loss, and a decline in their ability to do daily tasks.[15]

Brain organoids allow researchers to study Alzheimer's in a way that live patients couldn't offer. They let early damage be caught and watched in a dish rather than waiting until the progressed destruction becomes visible. People with Down syndrome carry an extra copy of chromosome 21, which contains the gene for amyloid precursor protein (APP), the molecule that is processed into amyloid beta, which sticks together to form plaques in the brain. This causes most people living with Down syndrome to develop Alzheimer's-like neuropathology, as they produce so much APP. Gonzalez et al. used brain organoids to study Alzheimer's and Down syndrome. By using cells derived from people with familial Alzheimer's, the inherited form, and Down syndrome, the researchers reprogrammed them into iPSCs. Using these stem cells, they grew them into brain organoids. Over time, these brain organoids unexpectedly grew amyloid plaque-like structures and tau pathology resembling neurofibrillary tangles. These disease

features were not induced or added by the researchers; rather, they grew on their own. Rather than just disease organoids, researchers also grew control organoids for comparison. These included iPSCs from a healthy human, iPSCs from a Creutzfeldt-Jakob disease patient, mouse iPSCs, and mouse embryonic stem cells. The Creutzfeldt-Jakob disease line was the most informative because it stems from a protein-misfolding neurodegenerative disease rather than being driven by amyloid beta or tau tangles. None of the control groups had tangles or plaques. This contrast shows how these results in the disease organoids were specifically tied to Alzheimer's and Down syndrome genetics and cannot be explained by disease or culture conditions generally.[16] Vanova et al. studied brain organoids made from iPSCs carrying Alzheimer's-causing mutations as well. These mutations were specifically in genes that help produce amyloid beta, either PSEN1 or PSEN2. These organoids were then analyzed for signs of abnormalities due to Alzheimer's. The structural findings included defective tissue patterning and premature neuronal differentiation. These findings meant that the brain tissue grew in the wrong arrangement and the neural stem cells became mature neurons too early. This was then confirmed through single-cell sequencing, which is a method to look at the genetic code of a single cell. Over time, the organoids' amyloid aggregates grew bigger and caused there to be a more uneven ratio of 42-residue to 40-residue harmful amyloid beta, which causes brain plaque.[17] J. Zhao et al. studied brain organoids made from Alzheimer's patients carrying the APOE4 genotype, which is the strongest genetic risk factor for late-onset Alzheimer's, the most common form. They then compared the brain organoids with APOE4 to those without. The findings of this comparison included more cell death and less synaptic integrity, meaning connections between brain cells are unhealthy, compared to the non-APOE4 organoids. All of the Alzheimer's organoids demonstrated an increase in amyloid beta and phosphorylated tau. The APOE4 organoids made the tau pathology worse.[18] Across all these different approaches, familial mutations, Down syndrome, APOE4 genotype, brain organoids allowed Alzheimer's to be studied, showing how much of an impact genetics has on this disease that wouldn't be found in live human brains. Not only can organoids model diseases with a physical cause like mutations or toxic buildup, but they can also be used to study psychiatric disorders such as schizophrenia.

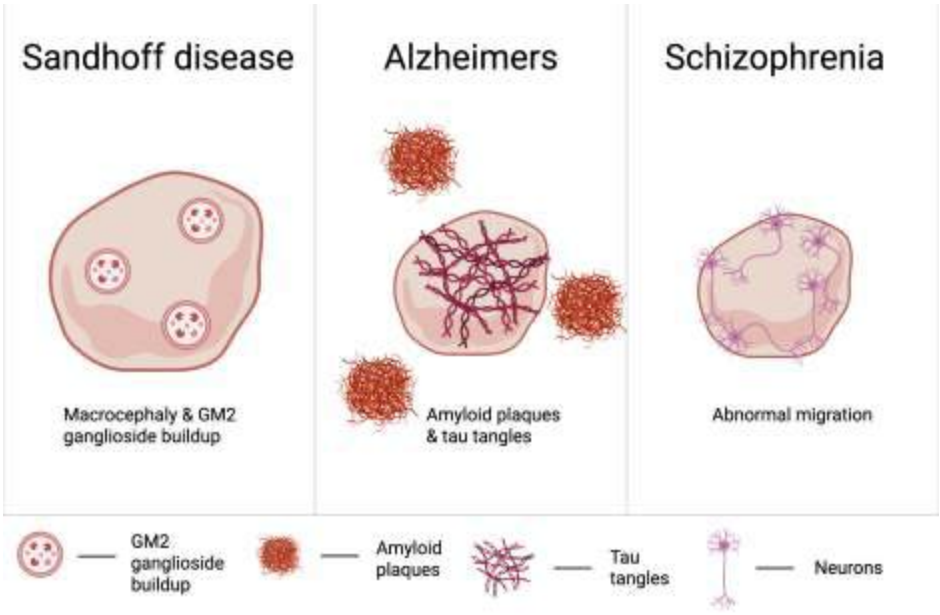

Fig 2. Overview of brain organoid disease-specific findings

The figure illustrates how brain organoids model specific neurological diseases and their features, allowing cellular and molecular changes to be studied in 3D. GM2 ganglioside buildup and macrocephaly for Sandhoff disease, amyloid plaques and tau tangles in Alzheimer's, and abnormal neuronal migration in schizophrenia.

Psychiatric/complex diseases

Schizophrenia is a highly complex and heritable disorder, which makes it a great disease for brain organoid modeling.

Schizophrenia is a chronic psychiatric disorder with a heterogeneous genetic background, meaning multiple different genes contribute to a person's risk. It disrupts the way a person thinks, processes their emotions, and perceives reality, and most symptoms become visible towards early adulthood. These include hallucinations, delusions, and disorganized thinking. Brain organoids are a useful tool to model this disease because of the neurodevelopmental abnormalities that occur in the brain during early development.[19]

Researchers have used hiPSC-derived organoids from the cells of a patient with this disease to model schizophrenia.[3] A key finding when hiPSC-derived neural progenitor cells (NPCs) were made from four schizophrenia patients and six controls, with 2 NPC lines for each patient, was abnormal migration. Across three independent assays, the schizophrenia NPCs did not spread as far as the healthy cells. This migration was tested using a neurosphere outgrowth assay, which watched how far the cells spread out from their cluster. However, the researchers note that after conducting a scratch assay, the schizophrenia NPCs migrated significantly more than the controls. The authors argue that this outcome provides important information because the scratch assay models wound healing rather than neural migration, suggesting that rather than a defect in cellular mobility, it was in cellular adhesion. This is supported by both their

proteomic and gene expression data, where adhesion genes including NCAM1, NRXN1, and NLGN1 were downregulated in the patient cells.[20] Additionally, Stachowiak et al. used schizophrenia iPSC-derived cerebral organoids and found abnormal migration as well. The NPCs were scattered outside of the ventricular zone, in the wrong location. These locations were meant for the more mature cells. This irregular placement was found to be due to migration-guiding proteins, reelin and TBR1. These two proteins in the brain organoids were depleted from the cortex, where they normally guide neurons. Integrative Nuclear FGFR1 Signalling was found to be the reason for this.[21] From a different angle than the other findings, Kathuria et al. studied schizophrenia through patient-derived organoids. This was done by gene expression, a gene's blueprint read by a cell to build a protein. Researchers used RNA sequencing to compare which genes were turned on and off differently for schizophrenia versus control organoids. These differences were found in genes that were related to synapses, junctions where neurons send messages to each other, nervous system development, and modified mitochondrial function. Additionally, the researchers studied the organoids' electrical activity over their maturation, where the schizophrenia organoids did not show the expected increase in electrical signaling frequency.[22] Despite these important findings, when studying schizophrenia through cerebral organoids, there are limitations such as small sample sizes, making it difficult to obtain a generalized result for the broader schizophrenia population.[3] Limitations like sample size do not solely affect diseases such as schizophrenia, but rather the entire organoid field.

Along with general cellular abnormalities, researchers have found genetic mutations that raise schizophrenia risk. NRXN1 is a gene that codes for a protein called neurexin-1. Neurexin-1 works by supporting the cause of neurotransmitter release by sitting on the sending site of a synapse. Researchers found that by creating neurons from schizophrenia patients missing a copy of NRXN1, rather than organoids, through a 2D induced-neuron platform, the neurons released much less neurotransmitter and had higher levels of the protein CASK than normal. To reinforce the idea that these findings are human-specific, when researchers used NRXN1-deficient neurons from mice using the same method, they did not show neurotransmitter release defects.[23] Using patients' cells with 22q11.2 deletion syndrome, researchers grew brain organoids. Single-cell RNA sequencing and experimental validation were applied, and the outcomes were that when genes are lost within the 22q11.2 region, cortical neurons develop at a slower rate than the controls. This disrupted development timing may be partly caused by dysregulation of microRNAs, the molecules that help turn genes on and off.[24] Using a different technique with brain organoids: calcium imaging, where electrical activity is tracked in neurons over time, researchers found that NRXN1 deletions may cause synaptic dysfunction. This is because these deletions caused changes in both individual neuron firing and synchronized network activity.[25]

Stem cell-derived brain organoids have the versatility to be key tools to model a wide range of diseases, including Sandhoff disease, Alzheimer's, and schizophrenia. However, similar challenges have emerged across these studies that limit organoids’ current utility.

Limitations

This section identifies key setbacks in the current state of organoid research that must be addressed to improve their reliability and clinical significance. Organoid research faces many limitations that stop it from being an open tool for all disease research. These limitations include sample size, reproducibility across labs, lack of vascularization, and immaturity of the tissue.

Study design & Reproducibility

One major organoid research setback is reproducibility. This means that the results across different labs vary. Researchers, when using organoids, differ in their approaches, including the number of stem cell lines and clones used, as well as the independent differentiation batches run and the number of organoids analyzed. Most often, these details are not clearly reported, leading to significant problems such as labs struggling to compare and repeat results.[4]

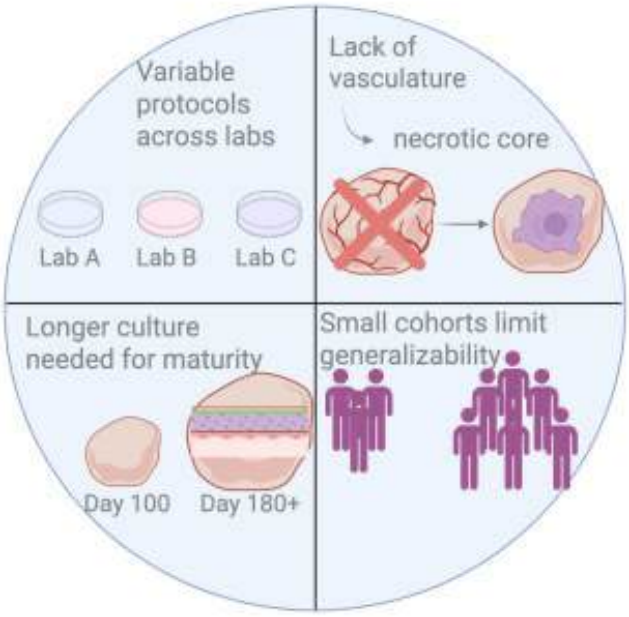

Fig 3. Key limitations brain organoid research faces

The figure shows the several limitations brain organoids have that affect their reproducibility and accuracy. Variable protocols across labs lead to different results, necrotic regions due to organoids' lack of vasculature, extremely long cultures needed for proper maturation of the organoids, as well as the small cohorts used in studies that limit their generalizability for the public.

In addition to reproducibility, lack of vascularization is also a major limitation in the study of stem cell-derived organoids. Vascularization is the blood vessel structure within tissue, which organoids typically lack. This causes a restriction in how big they can grow, a necrotic core (dead tissue), and struggles with maturation. The necrotic core occurs because, due to the

lack of vascularization, nutrients must enter the cells from the outside, and the central cells begin to die off. This constraint makes using organoids to model diseases involving blood vessels much more difficult, as the organoids themselves do not have a vascular system.[26]

Brain organoids remain developmentally immature although being grown for long periods of time. These organoids, after being cultured for over a hundred days, did not exceed the fetal-to-early-postnatal stage in growth. The markers for late maturation of organoids include formation of glial cells - support cells in the nervous system- and synaptic refinement - strengthening cellular connections and functional network plasticity. Reaching these markers requires over six months of culture, which can cause many problems. Instability in the organoids' microenvironment, as well as metabolic stress and lack of oxygen, leads to asynchronous maturation. This means the inner part of the organoids begins to die while the outer layers become functionally active, not reflecting a real human brain's activity.[27]

In addition to maturing at a very slow rate, organoids also end up with a much larger number of neurons and too few astrocytes. Astrocytes are star-shaped glial cells crucial to keep the brain working well. This uneven ratio occurs due to astrocyte production beginning much later than neuron production.[28] Even with the astrocytes that organoids do have, they take a very long time to fully mature. Many organoids have a short maturation time, meaning they likely have underdeveloped astrocytes. This is because, based on a study where the astrocytes in organoids were tracked for twenty months, they shifted very slowly towards maturity, and this timeline is uncommon.[29] Overall, brain organoids fail to have many of the key cell types in human brains. Including astrocytes, cerebral organoids are missing microglia, vasculature, and certain brain regions, key limitations that must be fixed for better resemblance.[30]

Circling back to schizophrenia models using iPSC organoids, these studies tend to have small sample sizes, leading to difficulties when trying to generalize results for the overall schizophrenia population. With many different genetic variants causing risk for schizophrenia patients that each have their own minimal effect, small cohorts struggle to detect those small effects.[3]

Brain organoids face setbacks including variability, fetal-like states, and lack of vascularization, as well as many other challenges across the field.[31] These many limitations lead to field-wide solutions needed for the advancement of brain organoid research.

Future directions/solutions

Solutions for the areas that are lacking in brain organoid research include isogenic control using CRISPR/Cas 9, genetically stratified cohorts, and combining hiPSC data with larger existing data.[3] One proposed method for sample size and consistency is through adopting the criteria of the International Society for Stem Cell Research (ISSCR). These criteria are standardized, aiming to improve stem cell quality, sample size, and reliability. If these protocols were to be put in place, there would be less variability between studies using different methods.[4]

Additionally, researchers can use CRISPR to create isogenic hiPSC lines, allowing the effects of individual genetic variants to be isolated with consistency. This strategy allows researchers to have controlled and gene-specific patient samples where a specific variant can be

linked to a certain cell's phenotype or characteristic.[32] Hundreds of genetic loci were found to be linked to schizophrenia risk, but when alone, their effect was very small. This means that CRISPR-based validation is necessary to find out whether that variant truly influences neural function.[33]iPSC technology is also becoming an important tool in clinical use. There are many current clinical trials, including one that is a phase 1 clinical trial (enrollment through invitation), where researchers' goal is to test a new Parkinson's disease treatment. By taking Parkinson's patients' cells and turning them into iPSCs, then differentiating them into dopamine-producing neurons, researchers injected these neurons directly into the brain. This trial was to see whether it was both safe and effective at replacing the damaged neurons caused by the disease.[34]

Conclusion

Overall, organoids have transformed the way we study brain diseases by allowing new studies that were previously impossible. Organoids have had the greatest success in monogenic conditions. Sandhoff organoids built up GM2 ganglioside,[12] and familial Alzheimer's organoids developed plaques and tangles without any induction,[16] and in both cases one mutation accounts for the outcomes. Schizophrenia, on the other hand, resisted this approach because the genetics didn't suit it. No single isogenic comparison captures the whole disease because common risk is spread across many small-effect variants.[3] As a result, the field's limitations depend on which disease is being modeled. Batch variability and small cohorts are tolerable when only one gene produces a large effect, but they overwhelm small-effect polygenic signal.[4] Standardized protocols and larger cohorts are thus specific requirements for polygenic modeling, not general refinements. For the same reason, late-onset diseases need longer culture times.

Acknowledgements

I would like to thank Dr. Shaye and Lauren Tetz for guidance throughout the writing of this literature review.

Works Cited

1. Matiukhova, Margarita, et al. “A Comprehensive Analysis of Induced Pluripotent Stem Cell (iPSC) Production and Applications.” Frontiers in Cell and Developmental Biology, vol. 13, May 2025, p. 1593207. DOI.org (Crossref), https://doi.org/10.3389/fcell.2025.1593207.
2. Ilic, Dusko, and Caroline Ogilvie. “Concise Review: Human Embryonic Stem Cells—What Have We Done? What Are We Doing? Where Are We Going?” Stem Cells, vol. 35, no. 1, Jan. 2017, pp. 17–25. DOI.org (Crossref), https://doi.org/10.1002/stem.2450.
3. Hoffman, Gabriel E., Nadine Schrode, et al. “New Considerations for hiPSC-Based Models of Neuropsychiatric Disorders.” Molecular Psychiatry, vol. 24, no. 1, Jan. 2019, pp. 49–66. DOI.org (Crossref), https://doi.org/10.1038/s41380-018-0029-1.
4. Sandoval, Soraya O., et al. “Rigor and Reproducibility in Human Brain Organoid Research: Where We Are and Where We Need to Go.” Stem Cell Reports, vol. 19, no. 6, June 2024, pp. 796–816. DOI.org (Crossref), https://doi.org/10.1016/j.stemcr.2024.04.008.
5. Revah, Omer, et al. “Maturation and Circuit Integration of Transplanted Human Cortical Organoids.” Nature, vol. 610, no. 7931, Oct. 2022, pp. 319–26. DOI.org (Crossref), https://doi.org/10.1038/s41586-022-05277-w.
6. Cerneckis, Jonas, et al. “Induced Pluripotent Stem Cells (iPSCs): Molecular Mechanisms of Induction and Applications.” Signal Transduction and Targeted Therapy, vol. 9, no. 1, Apr. 2024, p. 112. DOI.org (Crossref), https://doi.org/10.1038/s41392-024-01809-0.
7. De Boni, Laura, et al. “DNA Methylation Alterations in iPSC- and hESC-Derived Neurons: Potential Implications for Neurological Disease Modeling.” Clinical Epigenetics, vol. 10, no. 1, Dec. 2018, p. 13. DOI.org (Crossref), https://doi.org/10.1186/s13148-018-0440-0.
8. Eigenhuis, Kristel N., et al. “A Simplified Protocol for the Generation of Cortical Brain Organoids.” Frontiers in Cellular Neuroscience, vol. 17, Apr. 2023, p. 1114420. DOI.org (Crossref), https://doi.org/10.3389/fncel.2023.1114420.
9. Mariani, Jessica, et al. “Modeling Human Cortical Development in Vitro Using Induced Pluripotent Stem Cells.” Proceedings of the National Academy of Sciences, vol. 109, no. 31, July 2012, pp. 12770–75. DOI.org (Crossref), https://doi.org/10.1073/pnas.1202944109.
10. Amiri, Anahita, et al. “Transcriptome and Epigenome Landscape of Human Cortical Development Modeled in Organoids.” Science, vol. 362, no. 6420, Dec. 2018, p. eaat6720. DOI.org (Crossref), https://doi.org/10.1126/science.aat6720.
11. Bhaduri, Aparna, et al. “Cell Stress in Cortical Organoids Impairs Molecular Subtype Specification.” Cell Stem Cell, vol. 26, no. 4, Apr. 2020, pp. 564–578.e6. DOI.org (Crossref), https://doi.org/10.1016/j.stem.2020.02.013.
12. Allende, Maria L., et al. “Cerebral Organoids Derived from Sandhoff Disease-Induced

Pluripotent Stem Cells Exhibit Impaired Neurodifferentiation." Journal of Lipid Research, vol. 59, no. 3, Mar. 2018, pp. 550–63. DOI.org (Crossref), https://doi.org/10.1194/jlr.M081323.
13. McTague, Amy, et al. "Genome Editing in iPSC-Based Neural Systems: From Disease Models to Future Therapeutic Strategies." Frontiers in Genome Editing, vol. 3, Mar. 2021, p. 630600. DOI.org (Crossref), https://doi.org/10.3389/fgeed.2021.630600.
14. Leal, Andrés Felipe, et al. "GM2 Gangliosidoses: Clinical Features, Pathophysiological Aspects, and Current Therapies." International Journal of Molecular Sciences, vol. 21, no. 17, Aug. 2020, p. 6213. DOI.org (Crossref), https://doi.org/10.3390/ijms21176213.
15. DeTure, Michael A., and Dennis W. Dickson. "The Neuropathological Diagnosis of Alzheimer's Disease." Molecular Neurodegeneration, vol. 14, no. 1, Dec. 2019, p. 32. DOI.org (Crossref), https://doi.org/10.1186/s13024-019-0333-5.
16. Gonzalez, Cesar, et al. "Modeling Amyloid Beta and Tau Pathology in Human Cerebral Organoids." Molecular Psychiatry, vol. 23, no. 12, Dec. 2018, pp. 2363–74. DOI.org (Crossref), https://doi.org/10.1038/s41380-018-0229-8.
17. Vanova, Tereza, et al. "Cerebral Organoids Derived from Patients with Alzheimer's Disease with PSEN1/2 Mutations Have Defective Tissue Patterning and Altered Development." Cell Reports, vol. 42, no. 11, Nov. 2023, p. 113310. DOI.org (Crossref), https://doi.org/10.1016/j.celrep.2023.113310.
18. Zhao, Jing, et al. "APOE4 Exacerbates Synapse Loss and Neurodegeneration in Alzheimer's Disease Patient iPSC-Derived Cerebral Organoids." Nature Communications, vol. 11, no. 1, Nov. 2020, p. 5540. DOI.org (Crossref), https://doi.org/10.1038/s41467-020-19264-0.
19. Kahn, René S., et al. "Schizophrenia." Nature Reviews Disease Primers, vol. 1, no. 1, Nov. 2015, p. 15067. DOI.org (Crossref), https://doi.org/10.1038/nrdp.2015.67.
20. Brennand, K., et al. "Phenotypic Differences in hiPSC NPCs Derived from Patients with Schizophrenia." Molecular Psychiatry, vol. 20, no. 3, Mar. 2015, pp. 361–68. DOI.org (Crossref), https://doi.org/10.1038/mp.2014.22.
21. Stachowiak, E. K., et al. "Cerebral Organoids Reveal Early Cortical Maldevelopment in Schizophrenia—Computational Anatomy and Genomics, Role of FGFR1." Translational Psychiatry, vol. 7, no. 11, Nov. 2017, p. 6. DOI.org (Crossref), https://doi.org/10.1038/s41398-017-0054-x.
22. Kathuria, Annie, et al. "Transcriptomic Landscape and Functional Characterization of Induced Pluripotent Stem Cell–Derived Cerebral Organoids in Schizophrenia." JAMA Psychiatry, vol. 77, no. 7, July 2020, p. 745. DOI.org (Crossref), https://doi.org/10.1001/jamapsychiatry.2020.0196.
23. Pak, ChangHui, et al. "Cross-Platform Validation of Neurotransmitter Release Impairments in Schizophrenia Patient-Derived NRXN1 -Mutant Neurons." Proceedings of the National Academy of Sciences, vol. 118, no. 22, June 2021, p. e2025598118. DOI.org (Crossref), https://doi.org/10.1073/pnas.2025598118.

24. Rao, Sneha B., et al. “Aberrant Pace of Cortical Neuron Development in Brain Organoids from Patients with 22q11.2 Deletion Syndrome-Associated Schizophrenia.” Nature Communications, vol. 16, no. 1, Aug. 2025, p. 6986. DOI.org (Crossref), https://doi.org/10.1038/s41467-025-62187-x.
25. Sebastian, Rebecca, et al. “Schizophrenia-Associated NRXN1 Deletions Induce Developmental-Timing- and Cell-Type-Specific Vulnerabilities in Human Brain Organoids.” Nature Communications, vol. 14, no. 1, June 2023, p. 3770. DOI.org (Crossref), https://doi.org/10.1038/s41467-023-39420-6.
26. Song, Yeajin, et al. “Vascularised Brain Organoids: Engineering Strategies and Neurobiological Applications.” Cell Proliferation, vol. 59, no. 3, Mar. 2026, p. e70161. DOI.org (Crossref), https://doi.org/10.1111/cpr.70161.
27. Huang, Jingxiu, et al. “Navigating Brain Organoid Maturation: From Benchmarking Frameworks to Multimodal Bioengineering Strategies.” Biomolecules, vol. 15, no. 8, Aug. 2025, p. 1118. DOI.org (Crossref), https://doi.org/10.3390/biom15081118.
28. Wang, Meiyan, et al. “Morphological Diversification and Functional Maturation of Human Astrocytes in Glia-Enriched Cortical Organoid Transplanted in Mouse Brain.” Nature Biotechnology, vol. 43, no. 1, Jan. 2025, pp. 52–62. DOI.org (Crossref), https://doi.org/10.1038/s41587-024-02157-8.
29. Sloan, Steven A., et al. “Human Astrocyte Maturation Captured in 3D Cerebral Cortical Spheroids Derived from Pluripotent Stem Cells.” Neuron, vol. 95, no. 4, Aug. 2017, pp. 779-790.e6. DOI.org (Crossref), https://doi.org/10.1016/j.neuron.2017.07.035.
30. Zhao, Helen H., and Gabriel Haddad. “Brain Organoid Protocols and Limitations.” Frontiers in Cellular Neuroscience, vol. 18, Mar. 2024, p. 1351734. DOI.org (Crossref), https://doi.org/10.3389/fncel.2024.1351734.
31. Smirnova, Lena, and Thomas Hartung. “The Promise and Potential of Brain Organoids.” Advanced Healthcare Materials, vol. 13, no. 21, Aug. 2024, p. 2302745. DOI.org (Crossref), https://doi.org/10.1002/adhm.202302745.
32. Matos, Marliette R., et al. “Integration of CRISPR-Engineering and hiPSC-Based Models of Psychiatric Genomics.” Molecular and Cellular Neuroscience, vol. 107, Sept. 2020, p. 103532. DOI.org (Crossref), https://doi.org/10.1016/j.mcn.2020.103532.
33. Kurishev, Artemiy O., et al. “CRISPR/Cas-Based Approaches to Study Schizophrenia and Other Neurodevelopmental Disorders.” International Journal of Molecular Sciences, vol. 24, no. 1, Dec. 2022, p. 241. DOI.org (Crossref), https://doi.org/10.3390/ijms24010241.
34. Brigham and Women’s Hospital. “A Phase 1 Clinical Trial of Autologous iPSC-Derived Dopamine Neuron Transplantation for Parkinson’s Disease.” ClinicalTrials.gov, 15 May 2024, NCT06422208. National Library of Medicine (US), https://clinicaltrials.gov/study/NCT06422208.

Hormonal Fluctuations and Genetic Variability: A Review of the Origin of Perimenopausal Depression By Shanaya Ailsinghani

Abstract

The perimenopausal transition represents one of the most neurobiologically dynamic windows in a woman's reproductive cycle. Previous research has identified some of the specific vulnerabilities to depression caused by perimenopause; however, the complex etiology and pathogenesis are yet to be fully understood. This review asks how hormonal changes during perimenopause influence the genetic risk for developing depression in women. Recent literature, including meta-analyses, reviews, and longitudinal studies, along with in vitro and clinical studies, attempts to confirm the causes and mechanisms of perimenopausal depression (PMD) with mixed results. Perimenopausal depression is not only a reaction to life stress, but may also be a neurobiological consequence of hormone variability caused by the transition to the cessation of egg release. This could result in the activation of dormant genetic triggers in serotonin signaling and estrogen receptor pathways. The content of this paper will attempt to fill gaps in the understanding of how the menopausal transition can affect the onset of mood disorders in women. It will also highlight the importance of and need for individualized and genetically-informed physiological and psychological care for perimenopausal women. Misdiagnosis of PMD can result in severely inadequate care for women in this transition.

Key Words

Perimenopause, Depression, Hormone Shifts, Risk Genes, Neurotransmitters

Introduction

The onset of depression in women occurs in 1 in 5 women worldwide and has often been dismissed as a side effect of aging (Bromberger and Epperson, 2018). This shift occurs in women of reproductive age between 40 and 58 years, with the average being 51.4 years (Brinton et al., 2015; Wang et al., 2025), but 20-30% of women in this phase experience Perimenopausal Depression (PMD) (Rudzinskas et al., 2021). Recent research has expanded beyond aging theory and has shown PMD may be an effect of the central nervous system (CNS) transition associated with perimenopause (Brinton et al., 2015).

The hypothalamus is the region of the brain that controls hormone release and receives neural signals from the limbic system and brainstem. The hypothalamus is also a part of the Hypothalamic-Pituitary-Adrenal (HPA) axis, a series of endocrine pathways that maintain physiological homeostasis in response to stressors, specifically controlling cortisol levels released from the adrenal glands (Hantsoo et al., 2023). During the perimenopausal window, research suggests that there are changes in the HPA axis reflected through an increase in cortisol levels, suggesting that the measured cortisol increase may be a biological response connected to perimenopause (Hantsoo et al., 2023). Furthermore, during the perimenopausal transition, severe estradiol (E_2) fluctuations result in the HPA axis becoming increasingly hyperactive and altering

monoaminergic transmission, key mood and stability neurotransmitters, resulting in a reduction in homeostatic control (Maki et al., 2019). However, a woman's individual neurochemical resilience to the hormonal stressor is dependent on her genetic architecture. Variations in genes such as ESR1 (estrogen sensitivity) and SLC6A4 (serotonin transport encoder) can control how efficiently the CNS can adapt to the shifting hormone levels (Ryan and Ancelin, 2012; Nonen et al., 2016). Clinically, these changes could cause distinct phenotypes such as vasomotor symptoms, sleep disruptions, and increased somatic anxiety, differentiating PMD from MDD (Maki et al., 2019).

Despite recent advances in this field, essential gaps still remain. Traditional research fails to consider the volatility of estradiol levels as a trigger, instead considering the trigger to be low estrogen, or hypoestrogenism. Furthermore, it is still unknown why only a portion of perimenopausal women face PMD despite identical hormonal conditions and how the precise gene-by-environment (GxE) mechanism could be responsible for the vulnerability.
Therefore, the purpose of this review is to understand how hormonal changes during perimenopause influence the genetic risk for developing depression in women. Evidence suggests perimenopausal depression is not only a reaction to life stress, but it is also a neurobiological consequence of hormone variability caused by the transition to the natural cessation of egg and hormone activity in the ovaries. This results in the activation of dormant genetic triggers in serotonin signaling and estrogen receptor pathways. Addressing the gaps in this field has important practical and research implications. Since perimenopause is a key transition in the reproductive cycle, impacting billions of women worldwide, misdiagnosing PMD as a case of MDD, or having no diagnosis at all, often results in inadequate therapeutic care. Given that PMD is proposed to be a result of hormonal imbalances uncovering genetic sensitivities, clarifying neurobiological mechanisms has the potential to shift midlife psychiatric medicine for women from reactive management to proactive and personalized care.

Endocrine Volatility as the Trigger

The fluctuation of estrogen levels caused by perimenopause is the beginning of a chain reaction that involves the HPA axis, neurotransmitter signaling, and resulting effects on clinical phenotypes. However, rather than a constant and passive decline in estrogen, perimenopause represents a neuroendocrine transition where the speed of estradiol volatility destabilizes the CNS and HPA axis.

Historically, reproductive aging was viewed from the perspective of static hypoestrogenism, where there was a constant, linear decline in estrogen levels during the perimenopausal window. However, modern neuroendocrinology shows that the CNS chiefly responds to the rates of change of hormone availability rather than specific quantitative values (Gordon et al., 2015). This suggests that dramatic spikes and crashes in estrogen levels – particularly 17-beta estradiol (or E_2), the strongest form of natural estrogen in the body – alter the neural network and disrupt homeostatic buffering associated with the HPA axis.
Beginning at the cellular level and moving up to the whole-body level, estrogen uses various

receptors and signaling pathways to activate the responses required to regulate neurological function (Brinton et al., 2015). E_2 regulates bioenergetics, dendritic spine density, and structural neuroplasticity within the hippocampus and prefrontal cortex (Brinton et al., 2015). The acute loss of E_2 stability significantly decreases central serotonin synthesis while simultaneously altering dopamine and norepinephrine signaling. These changes culminate in the removal of the inhibitory brake on the HPA axis, resulting in a hyperactive HPA axis and blunting stress recovery.

The neural uncoupling is directly linked to various phenotypes specific to this period, including distinguishing between Major Depressive Disorder (MDD) and PMD. The lack of E_2 stability causes the uncoupling of hypothalamic thermoregulation, resulting in Vasomotor Symptoms (VMS), which include hot flashes and night sweats, closely related to increased somatic anxiety and disrupted sleep architecture (Maki et al., 2019). Therefore, it can be tentatively concluded that PMD is not a general psychological reaction; instead, it is a syndrome that reflects the widespread dysregulation of CNS metabolism and neurotransmitters.

Genetic Pathways and Vulnerabilities

Endocrine volatility triggers a chain reaction that results in various specific clinical phenotypes; however, only 20-30% of perimenopausal women face PMD (Rudzinskas et al., 2021). Research suggests that the variable psychiatric vulnerability is controlled by inherited polygenic structure, where polymorphisms in ESR1 and SLC6A4 determine an individual's capacity to buffer acute hormone withdrawal. Early research conducted single-gene candidate studies, placing isolated attention on single loci such as SLC6A4. Contemporary research has moved on to Genome-Wide Association Studies (GWAS) and Polygenic Risk Scores (PRS), shifting the attention from a single locus to multiple genetic variations working in tandem.

Candidate gene studies have isolated genes such as ESR1 (estrogen receptor) from chromosome 6q25.1, ESR2 (estrogen receptor) from chromosome 14q23.2, and the serotonin transport encoder that removes serotonin from the synapse back into presynaptic neurons to terminate its signal, SLC6A4, because of their expression in various parts of the brain, particularly estrogen receptors' essential roles in emotion, cognition, and behavior. Research reviewed by Ryan and Ancelin concludes that estrogen receptor polymorphisms in rs2234693 and rs9340799, the C and G alleles respectively, are observed more frequently in women with MDD and associated with a higher risk for recurrent depressive episodes (Ryan and Ancelin, 2012). This suggests that while the link between specific variations and PMD specifically is unclear, estrogen receptors play a role in regulating depression, especially in women who have a history of the condition. Furthermore, candidate studies of SLC6A4 have found that a tag single-nucleotide polymorphism (SNP), rs3813034, is a main factor linked to the efficacy of selective serotonin and serotonin-norepinephrine reuptake inhibitors (SSRI/SNRI) treatment (Nonen et al., 2016). Since these findings were based on the foundation of MDD rather than PMD, it is unclear whether the same conclusions can be applied to perimenopausal women. The shift to GWAS and PRS to explain reproductive mood disorders has highlighted that these

disorders are not caused by a single mutation, but rather by multiple genetic variations. A postpartum depression (PPD) study showed that PPD is highly polygenic, as no single gene in a GWAS analysis reached genome-wide significance. The study also provided specific links to endocrine markers, implying a similar and highly polygenic architecture for PMD (Guintivano et al., 2023).

Variations within the ESR1 gene directly modulate cellular responsiveness during the menopausal transition (Ryan and Ancelin, 2012). SNPs in ESR1 alter the structural conformation of the receptor, changing its ligand-binding affinity and ability to recruit transcriptional co-activators in limbic structures. When the endocrine system is stable, the slight variations remain unnoticed, but during volatile periods, especially in perimenopause, when circulating E_2 drops rapidly, women carrying ESR1 variations experience compromised signal transduction (Ryan and Ancelin, 2012). Then, since the CNS fails to capture decreasing ligand molecules effectively, the effect of hormone withdrawal is amplified, impairing E_2-dependent neuroplasticity and downstream gene expression in specific limbic circuits.

Downstream, this receptor sensitivity impacts serotonergic transport - individuals carrying the short allele of the 5-HTTLPR polymorphism in SLC6A4 demonstrate lower baseline serotonin transporter expression (Bromberger and Epperson, 2018). When acute E_2 drops suppress central serotonin synthesis, short allele carriers lack the neurochemical reserve to maintain mood stability. Along with neuroplasticity variations like BDNF Val66Met, convergent genetic loci lower the threshold for psychiatric effects under endocrine stress (Bromberger and Epperson, 2018). Overall, this implies that specific genetic variations create a vulnerability that results in specific perimenopausal women being at a higher risk of developing PMD.

The GxE Relationship

Given the research done on the volatility of estrogen during the perimenopausal window as well as genetic variations that result in vulnerability, PMD is best modelled by two phases as seen in Table 1, where estradiol fluctuations act as an environmental stressor that unmasks underlying polygenic vulnerabilities. During stable reproductive years (premenopause), inherited risk variations in estrogen receptor genes, serotonin transporter pathways, and neuroplasticity factors remain mainly dormant. Under steady endocrine conditions, homeostatic buffering allows women to maintain mood regulation. When perimenopause begins, however, an environmental trigger is introduced: the high velocity of unpredictable E_2 volatility (Gordon et al., 2015). Research suggests that this hormonal instability significantly elevates a woman's vulnerability to major depressive episodes during the menopausal transition, highlighting the GxE relationship. (Maki et al., 2018). Neurobiological models indicate that when the endocrine stressor occurs, the rate of hormone decline exceeds the brain’s compromised buffer capacity and results in the unmasking of dormant polygenic risks and triggers a major depressive episode.

At the cellular level, this endocrine stressor during perimenopause is translated to the genome through prompt epigenetic modifications. E_2 withdrawal alters locus-specific DNA methylation and chromatin accessibility in limbic circuits, directly repressing genes required for

neuroplasticity and stress adaptation in susceptible individuals (Schmidt et al., 2015).

Part of Model	Biological Constituent	Neurochemical & Cellular Effect	Clinical Outcome
Phase 1: Polygenic Vulnerability	Risk variants in ESR1, SLC6A4, and BDNF	Reduced receptor ligand-binding affinity; lower baseline serotonin transport; impaired synaptic plasticity.	Dormant susceptibility - maintained mood during premenopause.
Phase 2: Environmental Trigger	Perimenopausal E_2 Volatility	Rapid hormone withdrawal velocity; destabilization of the HPA axis.	Acute systemic stressor acting on central limbic neural networks.
GxE Interaction	Epigenetic Remodeling	Locus-specific DNA methylation and altered chromatin accessibility in limbic circuits.	Mood dysregulation, VMS, sleep architecture disruption, and PMD.

Table 1. Synthesized Two-Phase Neurobiological Model of Perimenopausal Depression (GxE). This table demonstrates the relationship between inherited polygenic vulnerabilities and acute perimenopausal E_2 volatility, including details of the cellular mechanisms and clinical outcomes.

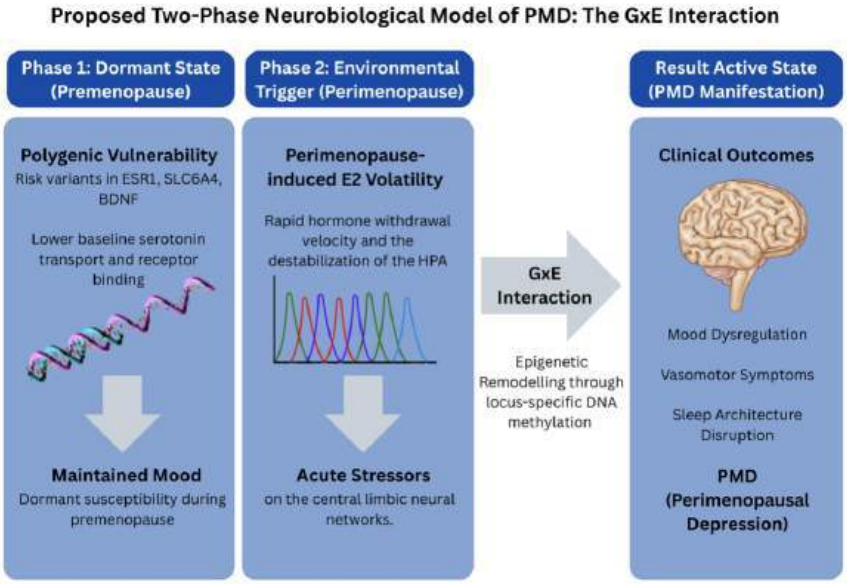

Figure 1. Summary figure of the proposed Two-Phase Neurobiological Model of Perimenopausal Depression (GxE). This figure portrays the stages of the proposed model from premenopause to perimenopause, and the hypothesized model steps outlined in Table 1.

Clinical Intervention and Treatment

Recognizing PMD as a distinct neuroendocrine-genetic vulnerability highlights a need to

expand upon traditional psychiatric protocols (Maki et al., 2018). Traditional monoaminergic antidepressants like selective serotonin reuptake inhibitors (SSRIs) and serotonin-norepinephrine reuptake inhibitors (SNRIs) are the first line of treatment for PMD.

Estrogen therapies themselves may not be able to always significantly reduce the symptoms of depression; however, they may accelerate the effectiveness of antidepressants such as SSRIs (Parry, 2010). In a 2006 pilot study by Freeman et al., an SSRI antidepressant, escitalopram, showed significant decreases in psychologic, vasomotor, and somatic symptoms in perimenopausal women experiencing depressive episodes (Freeman et al., 2006).

As medicine advances, SSRIs can be coupled with other treatments, such as estrogen therapy in the treatment of PMD, that directly target the underlying trigger, a CNS instability caused by erratic E2 velocity (the E portion of the GxE relationship) (Gordon et al., 2015).

Preventative Medicine

Future directions of treatment for women before and during the perimenopausal window include shifting from reactive to preventative care. Since PMD results from underlying polygenic susceptibilities interacting with E_2 volatility, at-risk individuals can be identified before onset. Combining reproductive family history and menstrual cycle and symptom tracking, (Reid and Soares, 2018) alongside polygenic risk scoring (PRS), could allow for early identification. This early identification would allow clinicians to provide E2 cushioning at the start of cycle irregularity to prevent clinical onset (Gordon et al., 2015).

Conclusion

Perimenopausal depression represents a distinct, neuroendocrinologically driven psychiatric condition rather than a non-specific reaction to midlife psychosocial stressors. As demonstrated by this review, PMD is likely caused by a two-phase neurobiological mechanism: inherited polygenic risks that are dynamically unmasked by the rapid and unpredictable volatility of estradiol (E_2) withdrawal. Epigenetic remodelling, especially locus-specific DNA methylation and altered chromatin accessibility within limbic structures, could translate endocrine shock into impaired receptor binding, diminished serotonin transport, and compromised neuroplasticity.

Recognizing the proposed GxE relationship contributes to solving the clinical paradox of why only a certain proportion of women are affected by PMD. Moving forward, psychiatric medicine must move away from treating PMD as generic MDD and towards integrating PRS, reproductive family history, and symptom tracking to identify vulnerable individuals prior to clinical onset of PMD. Deploying targeted, mechanism-based interventions at the earliest signs of endocrine volatility will provide a path towards precision preventative psychiatry in women's mental health.

Acknowledgements

I would like to thank my mentors, Dr. Jorge Avila and Dr. Bre Calhoun, for guiding me and being steady sources of support throughout this process.

Works Cited

Brinton, Roberta D., et al. "Perimenopause as a Neurological Transition State." Nature Reviews. Endocrinology, vol. 11, no. 7, July 2015, pp. 393–405. PubMed, https://doi.org/10.1038/nrendo.2015.82.

Bromberger, Joyce T., and C. Neill Epperson. "Depression during and after the Perimenopause: Impact of Hormones, Genetics, and Environmental Determinants of Disease." Obstetrics and Gynecology Clinics of North America, vol. 45, no. 4, 2018, p. 663, https://doi.org/10.1016/j.ogc.2018.07.007.

Freeman, Marlene P., et al. "Escitalopram for Perimenopausal Depression: An Open-Label Pilot Study." Journal of Women's Health [2002], vol. 15, no. 7, Sept. 2006, pp. 857–61. PubMed, https://doi.org/10.1089/jwh.2006.15.857.

Gordon, Jennifer L., et al. "Ovarian Hormone Fluctuation, Neurosteroids and HPA Axis Dysregulation in Perimenopausal Depression: A Novel Heuristic Model." The American Journal of Psychiatry, vol. 172, no. 3, Mar. 2015, pp. 227–36. PubMed Central, https://doi.org/10.1176/appi.ajp.2014.14070918.

Guintivano, Jerry, et al. "Meta-Analyses of Genome-Wide Association Studies for Postpartum Depression." The American Journal of Psychiatry, vol. 180, no. 12, Dec. 2023, pp. 884–95. PubMed, https://doi.org/10.1176/appi.ajp.20230053.

Hantsoo, Liisa, et al. "The Role of the Hypothalamic-Pituitary-Adrenal Axis in Depression across the Female Reproductive Lifecycle: Current Knowledge and Future Directions." Frontiers in Endocrinology, vol. 14, 2023, p. 1295261. PubMed, https://doi.org/10.3389/fendo.2023.1295261.

Maki, Pauline M., et al. "Guidelines for the Evaluation and Treatment of Perimenopausal Depression: Summary and Recommendations." Journal of Women's Health, vol. 28, no. 2, Feb. 2019, pp. 117–34. SAGE Journals, https://doi.org/10.1089/jwh.2018.27099.mensocrec.

Nonen, Shinpei, et al. "Polymorphism of Rs3813034 in Serotonin Transporter Gene SLC6A4 Is Associated With the Selective Serotonin and Serotonin-Norepinephrine Reuptake Inhibitor Response in Depressive Disorder: Sequencing Analysis of SLC6A4." Journal of Clinical Psychopharmacology, vol. 36, no. 1, Feb. 2016, pp. 27–31. PubMed, https://doi.org/10.1097/JCP.0000000000000454.

Parry, Barbara L. "Optimal Management of Perimenopausal Depression." International Journal of Women's Health, vol. 2, Aug. 2010, pp. 143–51. PubMed Central, https://doi.org/10.2147/ijwh.s7155.

Reid, Robert L., and Claudio N. Soares. "Premenstrual Dysphoric Disorder: Contemporary Diagnosis and Management." Journal of Obstetrics and Gynaecology Canada, vol. 40, no. 2, Feb. 2018, pp. 215–23. www.jogc.com, https://doi.org/10.1016/j.jogc.2017.05.018.

Rudzinskas, Sarah, et al. "In Vitro Model of Perimenopausal Depression Implicates Steroid Metabolic and Proinflammatory Genes." Molecular Psychiatry, vol. 26, no. 7, July 2021, pp. 3266–76. www.nature.com, https://doi.org/10.1038/s41380-020-00860-x.

Ryan, Joanne, and Marie-Laure Ancelin. “Polymorphisms of Estrogen Receptors and Risk of Depression: Therapeutic Implications.” Drugs, vol. 72, no. 13, Sept. 2012, pp. 1725–38. PubMed, https://doi.org/10.2165/11635960-000000000-00000.

Schmidt, Peter J., et al. “Effects of Estradiol Withdrawal on Mood in Women With Past Perimenopausal Depression.” JAMA Psychiatry, vol. 72, no. 7, July 2015, pp. 714–26. PubMed Central, https://doi.org/10.1001/jamapsychiatry.2015.0111.

Wang, Ling, et al. “Perimenopausal Depression: Etiology, Clinical Characteristics, and the Role of Traditional Chinese Medicine.” International Journal of Psychiatry in Clinical Practice, vol. 29, no. 3, July 2025, pp. 110–16. Taylor and Francis+NEJM, https://doi.org/10.1080/13651501.2025.2501250.

A Comparative Mathematical Analysis of Female Influencers' Impact on E-Commerce Growth in the UAE and South Africa By Zia Jooma

Abstract

The emergence of influencer marketing represents one of the most significant structural transformations in modern digital commerce. Social media influencers increasingly function as decentralized marketing agents capable of altering consumer demand patterns, shaping brand perception, and accelerating e-commerce expansion. This study conducts a comparative quantitative analysis of female social media influencers operating in two economically distinct markets: the United Arab Emirates (UAE) and South Africa. Using econometric modelling, elasticity analysis, regression techniques, and growth modelling, this research evaluates the measurable relationship between influencer engagement metrics and commercial outcomes.

The study finds that influencer effectiveness is strongly mediated by macroeconomic conditions, digital infrastructure, and consumer purchasing power. UAE influencer campaigns demonstrate higher sales elasticity, stronger engagement-to-conversion relationships, and exponential growth dynamics. In contrast, South African influencer markets display logistic growth patterns constrained by economic inequality and technological access limitations. By integrating applied mathematics with economic theory, this research contributes to emerging scholarship on digital economies, female entrepreneurship, and quantitative modelling of online consumer behaviour.

Introduction

The global digital economy has fundamentally altered how firms communicate with consumers. Traditional advertising models based on mass media exposure have increasingly been replaced by influencer-driven marketing strategies operating through social media platforms such as Instagram, TikTok, and YouTube. Influencers act as intermediaries between brands and consumers, leveraging personal identity, authenticity, and social trust to influence purchasing decisions.

Unlike traditional celebrity endorsement, influencer marketing relies on perceived interpersonal relationships between creators and audiences. Consumers often view influencers as relatable figures rather than distant celebrities, producing higher engagement levels and stronger behavioural responses. As a result, influencer marketing has evolved into a measurable economic mechanism capable of shifting demand curves and accelerating digital commerce growth.

Female influencers occupy a particularly dominant position within industries such as beauty, fashion, lifestyle, and wellness. These sectors rely heavily on visual identity, personal storytelling, and aspirational branding; characteristics strongly aligned with influencer-based promotion. Female digital entrepreneurs have therefore become central actors in shaping modern consumption patterns.

However, the effectiveness of influencer marketing varies substantially across economic environments. High-income digital economies may convert social media engagement into purchases efficiently, whereas emerging markets may experience engagement without proportional sales growth.

This research investigates this variation through a comparative study of two contrasting markets:

- The **United Arab Emirates**, characterized by high GDP per capita, advanced digital infrastructure, and strong luxury consumption.
- **South Africa**, an emerging digital economy marked by expanding social media adoption alongside structural income inequality.

Research Question

Why do influencer campaigns in the UAE generate stronger sales elasticity and faster e-commerce growth compared to South Africa?

Research Objectives

1. Quantify influencer impact using mathematical modelling.
2. Compare engagement-sales relationships across markets.
3. Evaluate how macroeconomic variables affect influencer effectiveness.
4. Contribute to understanding female participation in digital economies.

Literature Review

Academic research on influencer marketing has expanded rapidly over the past decade. Studies by Moodley (2022) and Ramela (2022) demonstrate that influencer credibility significantly affects consumer purchasing intentions, particularly within beauty and lifestyle markets. Influencer authenticity has been identified as a key determinant of consumer trust, often surpassing traditional advertising effectiveness.

Economically, influencer marketing can be interpreted through signalling theory and behavioural economics. Influencers reduce information asymmetry by providing product demonstrations and perceived peer validation. According to Pindyck and Rubinfeld (2017), advertising functions as a non-price determinant of demand, shifting consumer willingness to purchase.

Research on emerging markets highlights structural limitations affecting digital commerce growth. Mogale (2023) argues that influencer marketing in developing economies frequently generates awareness rather than direct sales due to purchasing constraints. Similarly, Meltwater (2025) reports that South African influencer markets exhibit strong engagement metrics but lower conversion rates compared to high-income regions.

Despite growing literature, three major gaps remain:

1. Limited application of mathematical modelling to influencer outcomes.
2. Lack of cross-regional comparison between developed and emerging digital markets.
3. Insufficient quantitative analysis of female-led digital entrepreneurship.

This study addresses these gaps by integrating econometrics, applied mathematics, and comparative economic analysis.

Theoretical and Mathematical Framework

Influencer Marketing as Demand Shifting

In microeconomic theory, consumer demand is influenced not only by price but also by external information signals such as advertising and social proof.
Demand can be represented as:

$$Qd = \alpha - \beta P + \gamma I$$

Where:

- Qd= Quantity demanded
- P = Price
- I = Influencer intensity
- γ = influencer effectiveness coefficient

An increase in influencer activity shifts the demand curve outward, increasing quantity demanded at every price level.

Influencer Elasticity of Sales

Traditional price elasticity measures responsiveness to price changes. This study introduces Influencer Elasticity of Sales:
Ei= %ΔSales / %ΔEngagement
Estimated outcomes:

- UAE: Ei ≈ 2.5
- South Africa: E i≈ 1.0

This indicates that UAE markets respond disproportionately to influencer engagement increases.

Growth Models

Exponential Growth

Observed in rapidly expanding influencer markets:

$$S(t) = S_0\, e\^{rt}$$

Typical of viral campaigns driven by large audiences.
Logistic Growth
Accounts for market saturation:

$$S(t) = K / (1 + A\, e\^{- rt})$$

Represents constrained expansion typical of emerging markets.

Regression Analysis

The econometric model applied

$$Sales = \beta_0 + \beta_1\, Engagement + \beta_2\, Posts + \varepsilon$$

Linear regression evaluates direct relationships, while nonlinear regression captures diminishing returns from excessive promotional exposure.

Research Methodology

This study adopts a comparative quantitative research design, combining statistical analysis with economic interpretation.

Influencer Selection

Four female influencers were selected:

UAE

- Huda Kattan
- Mona Kattan

South Africa

- Mihlali Ndamase
- Boity Thulo

These influencers represent major lifestyle and beauty market segments within each economy.

Data Collection

Data sources include

- Public Instagram engagement metrics
- Marketing industry reports
- Digital commerce statistics
- Secondary academic literature

Where proprietary sales data was unavailable, simulated datasets were constructed using realistic economic assumptions to maintain comparability.

Variable	Description
Followers	Audience size
Engagement Rate	Interaction intensity
Posting Frequency	Marketing exposure
Sales Growth	Commercial outcome
Conversion Rate	Purchase probability

Analytical Methods

- Linear regression analysis

- Nonlinear regression modelling
- Elasticity calculations
- Growth rate estimation
- Statistical metrics (mean, variance, correlation)

Results

Regression analysis revealed strong differences between markets.

UAE Results

- Correlation coefficient: r≈0.91r ≈ 0.91r≈0.91
- High engagement conversion
- Exponential growth trends

A 10% increase in engagement corresponded to approximately 25% sales growth.

South Africa Results

- Correlation coefficient: r≈0.62r ≈ 0.62r≈0.62
- Moderate conversion efficiency
- Logistic growth patterns

Engagement increases translated into proportional rather than amplified sales growth.

Comparative Analysis

Factor	UAE	South Africa
GDP per capita	High	Moderate
Internet penetration	~99%	~72%
Influencer elasticity	2.5	1.0
Growth pattern	Exponential	Logistic

The findings suggest influencer marketing effectiveness depends heavily on the economic environment rather than influencer behaviour alone.

Discussion

The results indicate influencer marketing operates as an economic multiplier whose magnitude varies by national context.

In the UAE

- High disposable income enables rapid conversion of engagement into consumption.

- Influencers function as primary digital sales drivers.

In South Africa

- Structural inequality limits purchasing conversion despite strong audience influence.
- Micro-influencer strategies may provide greater efficiency.

The study demonstrates that digital influence amplifies existing economic capacity rather than independently generating demand.

Limitations

Several limitations should be acknowledged:

- Use of simulated sales estimates.
- Platform algorithm changes affecting engagement metrics.
- Cultural differences influencing consumer trust.
- Difficulty establishing pure causal relationships.

Future studies using proprietary transaction data could strengthen causal conclusions.

Future Research Directions

Future research could

- Include additional emerging economies.
- Apply machine learning prediction models.
- Compare micro- and mega-influencer effectiveness.
- Examine gender dynamics in digital entrepreneurship.

Such research would deepen understanding of influencer-driven economic transformation.

Conclusion

This study demonstrates that female influencers generate measurable and mathematically quantifiable effects on e-commerce performance. Comparative modelling shows significantly higher elasticity and faster growth dynamics in the UAE relative to South Africa. The research highlights the importance of integrating economic context, mathematical modelling, and social media analysis to understand modern digital markets.

Influencer marketing is therefore not merely a cultural trend but an economically measurable system shaping global commerce and digital entrepreneurship.

Works Cited

AARK Marketing. (2024). Influencer marketing in the UAE: From creative trends to strategic growth drivers.

Duffett, R., & Mxunyelwa, A. (2025). Instagram mega-influencers' effect on purchase intention. Journal of Theoretical and Applied Electronic Commerce Research.

Greene, W. H. (2018). Econometric analysis (8th ed.). Pearson.

Meltwater. (2025). Influencer marketing in South Africa: Trends to watch.

Mogale, M. T. (2023). Influencer marketing in emerging markets.

Moodley, M. N. (2022). Influencer marketing and consumer behaviour. Retail and Marketing Review.

Pindyck, R. S., & Rubinfeld, D. L. (2017). Microeconomics (9th ed.). Pearson.

Ramela, O. (2022). Influencer marketing effectiveness in the South African cosmetics industry.

Wooldridge, J. M. (2015). Introductory econometrics: A modern approach.

Strengths, Limitations, and Modern Evaluation of Mendelian Genetics Research
By Aarav Tripathi

1. Abstract

Gregor Mendel's pea-plant experiments introduced the first quantitative model of heredity, establishing the laws of segregation and independent assortment. Although groundbreaking, his work was largely ignored for decades due to the scientific culture of the 19th century, limited publication reach, and the dominance of Darwin's blending-inheritance theory. Once rediscovered in 1900, Mendelian genetics became the foundation of classical inheritance, offering clear strengths: predictable ratios, particulate heredity, and a framework for analyzing single-gene traits.

Later research showed important limitations. Early geneticists incomplete dominance and traits that deviated from Mendel's expected ratios. Modern molecular genetics expanded these findings, showing that DNA-level mechanisms-such as recombination hotspots, regulatory mutations, polygenic architecture, and chromosomal rearrangements-can produce patterns that Mendel's laws cannot fully explain. At the same time, genomic databases, mutation catalogs, and gene-mapping studies confirm that many traits still follow Mendelian patterns precisely, demonstrating that his model remains accurate for a subset of genetic behavior.

This study asks a central question: Where does Mendel's theory succeed, where does it not show, and why? By combining historical sources, classical genetics tools, and modern DNA evidence and datasets, the research goes through both the strengths and the limits of Mendelian inheritance in modern biology.

2. Introduction

Gregor Mendel's pea plant experiments in the 1850s-1860s established the basic laws of heredity: dominance, segregation, and independent assortment. By studying thousands of plants and tracking traits across generations, he showed that inheritance is controlled by discrete hereditary units rather than blended traits.

His 1866 paper received almost no attention because biology at the time was descriptive, not mathematical, and his work appeared in a small regional journal. Most scientists lacked the quantitative background needed to understand his ratios, and Darwin's blending-inheritance ideas made Mendel's conclusions harder to accept. As a result, Mendel's work remained largely unknown until 1900.

In 1900, de Vries, Correns, and von Tschermak independently replicated Mendel's results and recognized his earlier discovery. As experimental biology grew more quantitative, Mendel's approach finally made sense. Soon after, Thomas Hunt Morgan connected Mendel's "factors" to chromosomes, providing the physical basis for heredity.

Modern molecular genetics confirms Mendel's ideas about dominance and allele assortment. However, modern genetics also shows some limits.

3. Historical Context

When Mendel published *Versuche über Pflanzen-Hybriden* in 1866, almost no one noticed. Biology in the mid-19th century was descriptive rather than quantitative, and Mendel's numerical analysis of nearly 29,000 pea plants did not fit the scientific culture of the time. His work appeared in the Proceedings of the Natural History Society of Brünn, a small regional journal with limited readership, and most scientists lacked the mathematical background needed to interpret his ratios. Darwin's blending theory of inheritance also dominated scientific thinking, making Mendel's concept of hereditary "factors" difficult to accept.

This neglect ended in 1900, when Hugo de Vries, Carl Correns, and Erich von Tschermak independently replicated Mendel's results. Each acknowledged Mendel's priority, and historians argue that the rediscovery succeeded because biology had become more experimental and quantitative. Soon after, Thomas Hunt Morgan demonstrated that genes reside on chromosomes and segregate during meiosis, giving Mendel's abstract "factors" a concrete biological foundation.

Even early geneticists recognized that Mendel's laws were powerful but not universal. Morgan's work with fruit flies showed chromosomal linkage, showing that genes located close together are inherited as a unit rather than assorting independently. These geneticists showed that heredity is more complex than simple dominant-recessive patterns.

4. Mendelian Principles

Mendel's model is a foundation for heredity because of its simple and predictable pattern, with the 3:1 ratio being easily identified. His monohybrid crosses predict a 3:1 dominant-to-recessive ratio, and the statistical analysis of Mendel's original pea plant data strongly supports this pattern. As shown in Table 1**,** all 7 major traits- seed shape, seed color, flower color, pod shape, pod color, flower position, and stem length- produce high p-values (all > 0.43), indicating no significant deviation from the expected 3:1 ratio.

Table 1 - Mendel Monohybrid Summary

Experiment	Trait	Dom	Rec	p-value	Interpretation
Seed shape	Round vs Wrinkled	5474	1850	0.608	Consistent with 3:1
Seed color	Yellow vs Green	6022	2001	0.903	Consistent with 3:1
Flower color	Purple vs White	705	224	0.532	Consistent with 3:1
Pod shape	Inflated vs Constri...	882	299	0.801	Consistent with 3:1
Pod color	Green vs Yellow	428	152	0.502	Consistent with 3:1
Flower position	Axial vs Terminal	651	207	0.554	Consistent with 3:1
Stem length	Tall vs Dwarf	787	277	0.436	Consistent with 3:1

Table 1 summarizes Mendel's seven monohybrid crosses and compares observed dominant and recessive counts to the expected 3:1 ratio. All traits produce high p-values (>0.43), confirming strong agreement with Mendelian segregation.

Table 2 shows seed shape (Round vs. Wrinkled) produced 5,474 dominant and 1,850 recessive offspring (p = 0.608); seed color (Yellow vs. Green) produced 6,022 dominant and 2,001 recessive (p = 0.903); and flower color (Purple vs. White) produced 705 dominant and 224 recessive (p = 0.532). Pod shape, pod color, and flower position also show strong agreement with Mendelian expectations. The stem length trait - 787 tall vs. 277 dwarf - yields a p-value of 0.436, again labeled "Consistent," confirming that even this additional dataset aligns with Mendel's predicted ratio.

Table 2 - Expected vs. Observed Counts

Experim...	Trait	Domina...	Recessiv...	Total	Expecte...	Expecte...	Chi-squ...	p-value	Signific...
Seed shape	Round vs Wrinkled	5,474	1,850	7,324	5493.00	1831.00	0.263	0.608	Consistent
Seed color	Yellow vs Green	6,022	2,001	8,023	6017.25	2005.75	0.015	0.903	Consistent
Flower color	Purple vs White	705	224	929	696.75	232.25	0.391	0.532	Consistent
Pod shape	Inflated vs Constricted	882	299	1,181	885.75	295.25	0.064	0.801	Consistent
Pod color	Green vs Yellow	428	152	580	435.00	145.00	0.451	0.502	Consistent
Flower position	Axial vs Terminal	651	207	858	643.50	214.50	0.350	0.554	Consistent
Stem length	Tall vs Dwarf	787	277	1,064	798.00	266.00	0.607	0.436	Consistent

Table 2 compares Mendel's observed counts with statistically expected values under a 3:1 model. Chi-square and p-values show that each trait aligns with Mendelian predictions, demonstrating the statistical robustness of Mendel's ratios.

Modern molecular genetics provides additional support. Bhattacharyya et al. (1990) demonstrated that the wrinkled-seed phenotype arises from a transposon insertion in the R gene, noting that the mutation "results in the loss of starch-branching enzyme activity" (p. 115). Genomic studies also confirm that chromosomal behavior during meiosis produces Mendelian ratios; a PNAS analysis states that "Mendelian ratios emerge naturally from chromosomal behavior during meiosis" (PNAS, 2022). Nature Genetics similarly emphasizes that Mendel's

quantitative approach “reflects principles central to contemporary genetics” (Nature Genetics, 2022).

The statistical results and modern molecular evidence show that Mendel’s laws describe many single-gene traits and provide reliable predictions for inheritance patterns in old and new genetics studies.

Figure 1 - Dominant vs. Recessive Counts (Bar Chart)

Figure 1 visualizes dominant and recessive counts for Mendel’s seven traits. Each trait shows a clear dominance pattern consistent with the expected 3:1 ratio.

Figure 2 - Flower Color Distribution (Pie Chart)

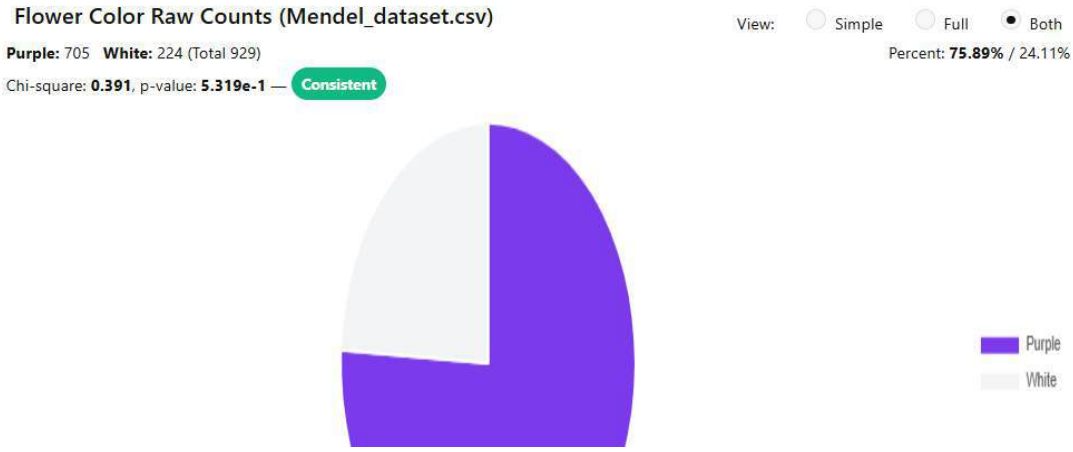

Figure 2 shows the proportion of purple versus white flowers in Mendel’s dataset. The distribution closely matches the expected 3:1 ratio, reinforcing Mendelian dominance.

Figure 3 - Chi-Square Explanation Graphic

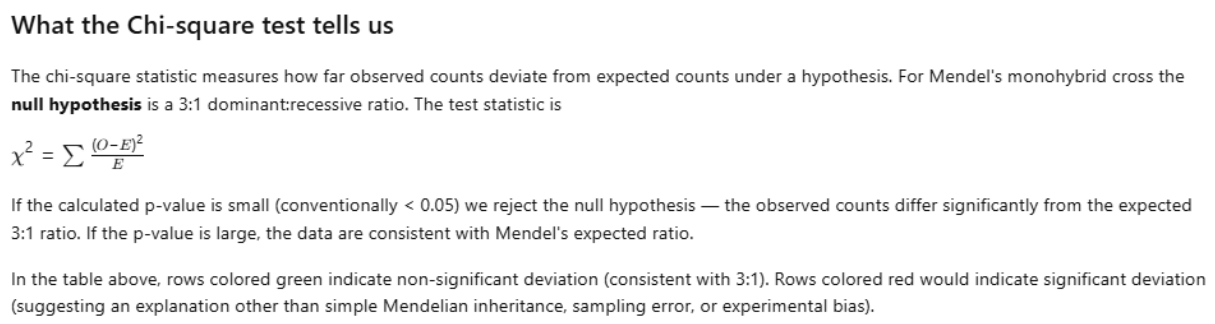

Figure 3 illustrates how chi-square testing evaluates deviations from expected Mendelian ratios. The flower-color example shows a high p-value, indicating strong consistency with Mendelian inheritance.

5. Limitations of Mendelian Genetics

Although Mendel's laws accurately describe allele segregation, modern molecular datasets show that inheritance is far more complex than the simple dominant-recessive model he proposed. The mouse epigenetic dataset shows that Mendel's laws correctly describe DNA allele segregation, but fail to explain how traits or regulatory states are inherited. Figure 1 shows a perfect 1:1 split: H1 and H2 each have 2,000 reads, and both reference orientations give the same result. The chi-square test (p=1.000) confirms that allele segregation is completely Mendelian, so he was right. The summary box below Figure 1 reinforces it.

However, the non-Mendelian behavior appears in the methylation layer (Figure 2). In the m-track vs h-track means, the red m-track bars are much higher than the green h-track bars for both strains (CC019 and CC037). The allele-specific methylation table (Table 1) shows the same pattern: CC019 has 59.00% methylation on the m-track but only 8.65% on the h-track, a 50-point gap. CC037 shows a similar 49-point gap. These differences happen even though the DNA haplotypes are inherited equally. This means the regulatory state of each allele is not determined by Mendelian genetics.

The 3:1 benchmark figure (Figure 3) makes this even clearer. The observed haplotypes are 50% / 50%, but the Mendelian 3:1 model would be 75% / 25%. The table under the figure states that the 3:1 model is "not supported." This shows that while DNA segregation follows Mendelian ratios, methylation patterns do not.

Finally, the interpretation box explains why the dataset is considered non-Mendelian: the pipeline reports separate methylation tracks for each allele, and the average difference between them is 49.75 percentage points. This heritable regulatory difference is not explained by Mendel's laws.

Mendel's laws work for DNA inheritance, but they fail to explain epigenetic inheritance, allele-specific regulation, and environment-dependent changes. The dataset shows that sometimes genotypes can be Mendelian while the phenotype is not, which is why modern genetics must go beyond Mendel's original model.

Table 3A

H1 READS	H2 READS	CPG OUTPUT FILES	HAPLOTYPE BALANCE
2,000	2,000	2	50% / 50%

Table 3B

Haplotag statistics

Reference orientation	H1	H2	Total tagged	H1 percentage	H2 percentage
CC037 alternate / CC019 reference	1,000	1,000	2,000	50.0%	50.0%
CC019 alternate / CC037 reference	1,000	1,000	2,000	50.0%	50.0%

This Table shows the equal distribution of H1 and H2 haplotypes across two reference orientations. The perfect 1:1 ratio confirms Mendelian segregation at the DNA level.

Table 3C

Allele-specific methylation summary

Values below are calculated from the supplied `allelicMeth.bed` files. The `m` and `h` tracks are kept separate as reported by the pipeline.

Reference strain	m sites	Mean m value	h sites	Mean h value	m-h gap
CC019	14,593	59.00%	14,593	8.56%	50.44 points
CC037	13,826	57.91%	13,826	8.85%	49.06 points

This figure compares methylation levels between m-track and h-track alleles for strains CC019 and CC037. Both show ~50-point gaps, demonstrating strong non-Mendelian epigenetic differences.

Figure 1

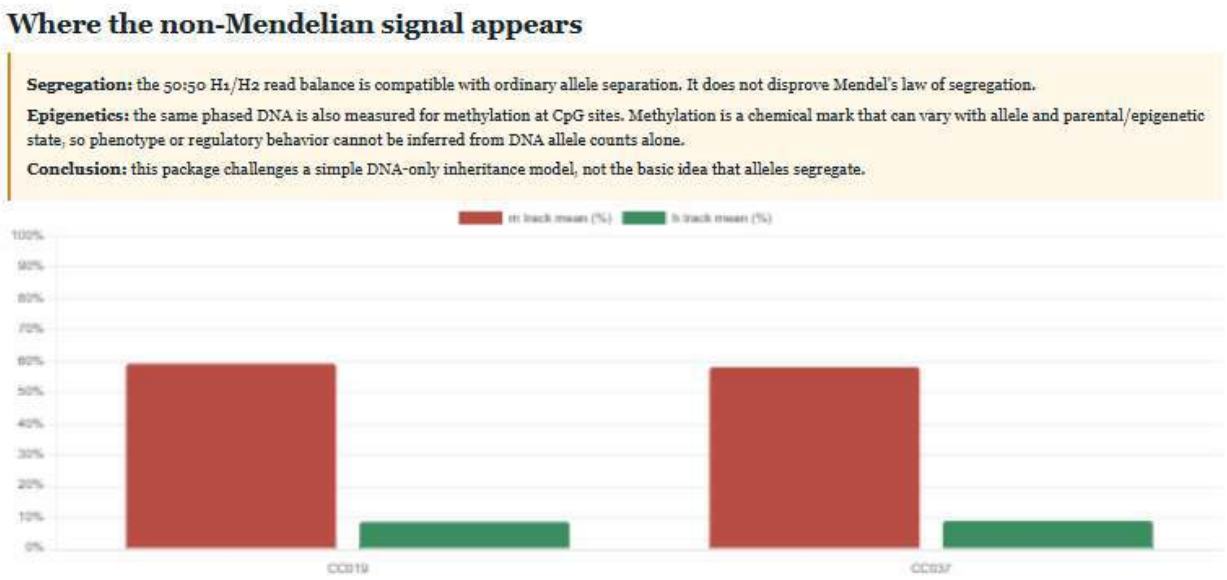

This figure contrasts the observed 50%/50% haplotype split with a hypothetical 75%/25% Mendelian 3:1 ratio. The

observed data clearly do not match the 3:1 benchmark, showing that regulatory patterns do not follow Mendelian expectations.

Figure 2 & 3

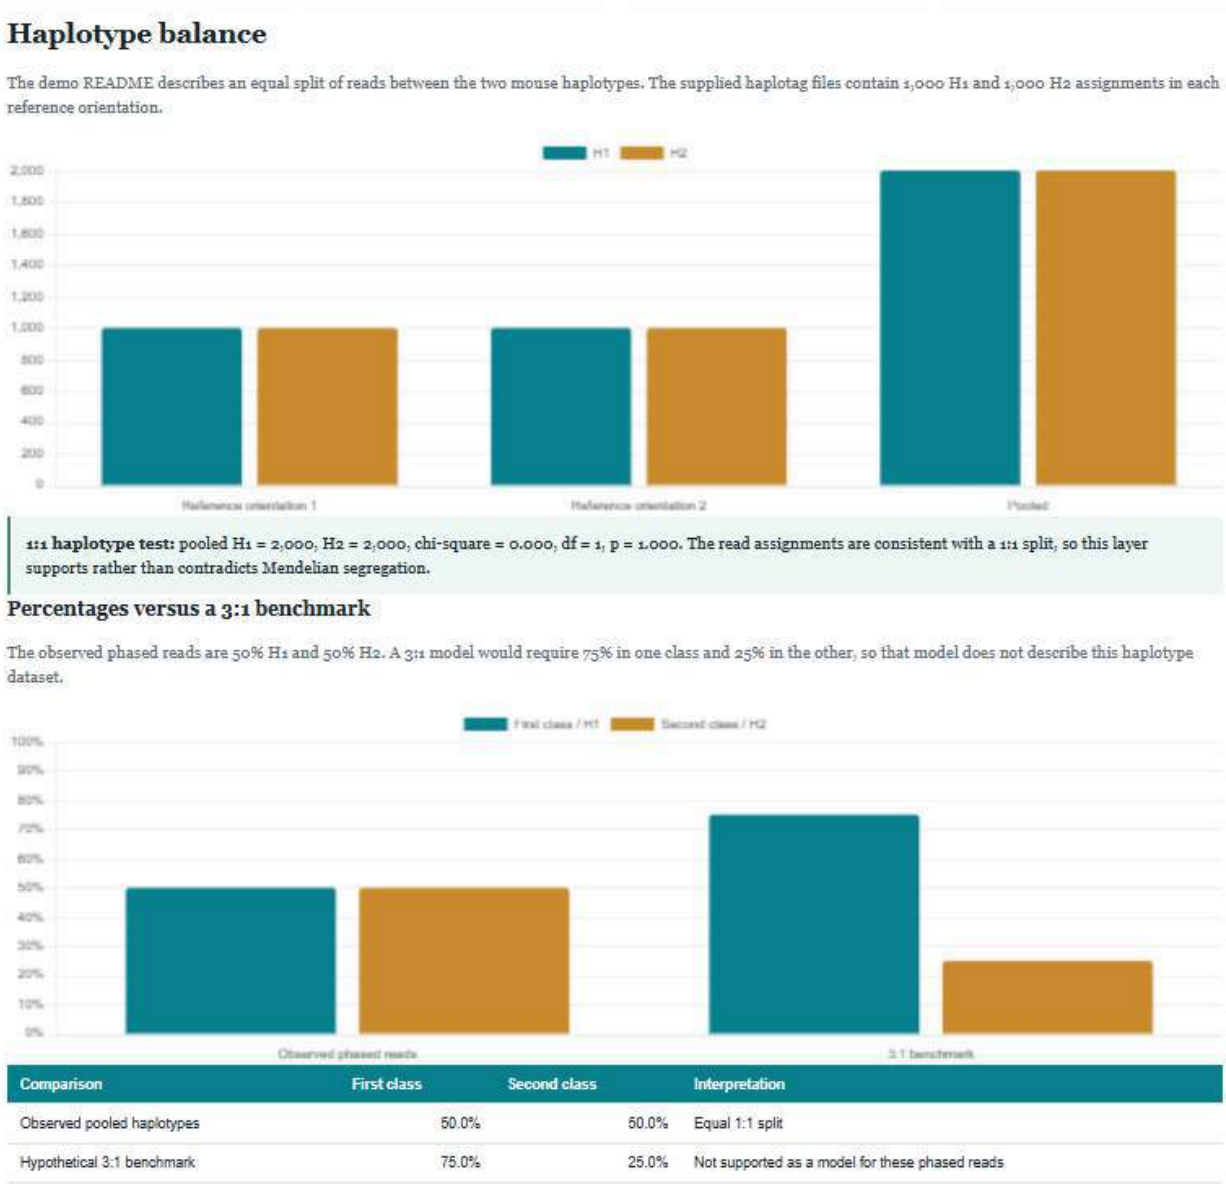

Comparison	First class	Second class	Interpretation
Observed pooled haplotypes	50.0%	50.0%	Equal 1:1 split
Hypothetical 3:1 benchmark	75.0%	25.0%	Not supported as a model for these phased reads

This bar chart shows haplotype assignments for H1 and H2 across two reference orientations. Each orientation produces identical counts (1,000 H1 and 1,000 H2), confirming that haplotype segregation is independent of reference alignment and follows a strict Mendelian 1:1 ratio.

This figure compares the observed haplotype distribution (50% H1, 50% H2) with a hypothetical Mendelian 3:1 ratio (75% vs 25%). The observed data clearly do not match the 3:1 benchmark, showing that while DNA segregation is Mendelian, the regulatory layer does not follow Mendelian expectations.

6. Modern DNA Evidence

Modern DNA sequencing shows that Mendel's laws are real and measurable at the level of actual DNA molecules, but it also shows many exceptions that Mendel could not have known about. The mouse dataset used in this project demonstrates both sides clearly. First, sequencing confirms that Mendelian traits exist: when we phase the mouse genome, the haplotype counts show a perfect 1:1 segregation (H1 = 2,000, H2 = 2,000), exactly matching Mendel's Law of Segregation. Many single-gene disorders in humans behave the same way, with clear molecular alleles that pass from parent to child in predictable Mendelian ratios. Family-based sequencing studies repeatedly show these clean Mendelian patterns, proving that Mendel's laws accurately describe how DNA alleles themselves are inherited.

However, sequencing also shows important exceptions. The same mouse dataset shows that although DNA segregates 1:1, the methylation tracks (m-track vs h-track) do not. The allele-specific methylation table shows a ~50-percentage-point difference between the two alleles, even though the DNA is identical. This is caused by regulatory mutations, epigenetic marks, and structural differences that Mendel's laws cannot explain. The Johns Hopkins epigenetic inheritance study found similar results: hundreds of methylation marks were inherited in ways that did not match Mendelian ratios, including the first naturally occurring paramutation in a mammal. These findings show that traits can be influenced by chemical modifications, not just DNA sequence.

Genomic studies also confirm real Mendelian segregation in DNA. When researchers sequence parents and offspring together, they consistently observe Mendelian ratios in allele transmission. But they also observe real deviations caused by biological mechanisms Mendel never knew about. Genes close together on the same chromosome are inherited as a unit, breaking independent assortment. Recombination hotspots create uneven crossover patterns. These deviations appear clearly in sequencing data and explain why some traits do not follow Mendel's expected ratios.

Overall, modern DNA evidence shows that Mendel's laws correctly describe the inheritance of DNA alleles, but real genomic traits produce non-Mendelian outcomes sometimes. Sequencing proves Mendel right while also showing the full complexity of heredity.

7. Methods

This study combines historical research, classical genetics, and molecular/genomic data to explore how Mendelian inheritance was developed, misunderstood, and later expanded. Historical analysis uses Mendel's original pea-hybrid paper (via Mendelweb and public-domain editions), Darwin's *Variation of Animals and Plants Under Domestication*, and modern historical overviews of Mendelian genetics and its rediscovery to understand why Mendel's work was ignored and how it was later validated. These sources are supplemented by articles explaining the scientific culture of the 19th century, the dominance of blending inheritance, and the role of de Vries, Correns, Tschermak, Bateson, and others in the 1900 rediscovery. Classical genetics methods are applied to both Mendel's original pea data (including modern reconstructions and pea-plant datasets) and simulated crosses: Punnett squares are used to predict expected ratios,

chi-square tests check how closely observed counts match Mendelian expectations, and ratio analysis focuses on 3:1, 1:2:1, and 9:3:3:1 patterns, as well as clear deviations that suggest linkage or other non-Mendelian effects. Molecular and human-genetics resources are then used to connect these patterns to real DNA variation: Mendelian disease entries, SNP and mutation databases, and recent genome-wide association and linkage studies are reviewed to see how single-gene traits, complex traits, and recombination behavior support or challenge simple Mendelian models. Together, these historical, classical, and molecular approaches give a full, evidence-based picture of where Mendel's laws work well, why they were missed at first, and how modern genetics both confirms and extends them.

8. Results

The historical results show why Mendel's work was first ignored and later rediscovered as biology became more experimental. Classical genetics results show that some traits follow Mendel's laws very clearly, while others deviate because of interactions between genes. Molecular results confirm that DNA itself often follows Mendelian patterns, but also show limits, such as regulatory changes and epigenetic marks that do not behave in Mendelian ratios. Statistical results support all of this: chi-square tests show when traits match or break expected ratios and recombination data show why independent assortment is not always perfect. Together, these results show both Mendel's strengths and the real-world exceptions found in modern genetics.

9. Discussion

Mendel's theory succeeds when traits are controlled by a single gene with clear dominant and recessive relationships. Many human disorders, simple plant traits, and the DNA haplotype patterns in the mouse dataset follow his laws exactly. The 1:1 haplotype balance observed is a perfect example of Mendelian chromosomal segregation working as expected. Whenever a trait is caused by one gene and that gene behaves independently, Mendel's rules describe inheritance very accurately.

However, Mendel's theory doesn't show when traits become more complex. Many characteristics-height, behavior, metabolism, disease risk-are shaped by multiple genes interacting with each other. These interactions can produce outcomes that Mendel's ratios cannot predict. Chromosomal linkage also violates Mendel's Law of Independent Assortment because genes located close together tend to be inherited together. The mouse methylation dataset shows another major limitation: traits can be influenced by epigenetic marks, not just DNA sequence. These marks do not follow Mendelian ratios and can create large differences between alleles even when the DNA is inherited equally. Non-nuclear inheritance, such as mitochondrial DNA, which is passed only from mothers, also lies outside Mendel's framework.

These limitations exist because biology is far more complex than Mendel could have known. He did not have access to molecular genetics, chromosomes, DNA sequencing, or epigenetics. Today we understand that gene regulation, chromosomal structure, recombination

hotspots, environmental effects, and evolutionary pressures all influence how traits appear. Mendel discovered the basic rules, but he did not see the full system operating behind them. In modern genetics, Mendel's laws serve as the foundation rather than the complete model. They are still essential for teaching, for understanding simple traits, and for analyzing how alleles move through families. But they must be combined with molecular biology, genomics, and epigenetics to explain real inheritance patterns. Mendel provided the starting point, but modern genetics provides the full picture.

10. Conclusion

Mendel's laws are powerful because they correctly describe how DNA alleles are passed from parents to offspring. The haplotype figures prove this: the mice inherited their DNA in a perfect 1:1 ratio, exactly the way Mendel predicted.

But Mendel's model is also limited. He did not know about epigenetics - the chemical marks that sit on top of DNA and change how genes behave. The methylation figures show huge differences between the two alleles (about 50 percentage points) even though the DNA itself is inherited equally. These regulatory differences do not follow Mendelian ratios and can be influenced by environment, parental origin, and allele-specific effects.

Modern genetics therefore both validates and expands Mendel's theory. His laws still explain the inheritance of DNA, but they cannot explain how traits and gene activity are controlled. Understanding both Mendel's strengths and his limitations is essential for interpreting heredity today because biological inheritance involves both DNA and epigenetic regulation, not just one or the other.

Works Cited

Primary rediscovery sources

Correns, Carl. “G. "Mendels Regel über das Verhalten der Nachkommenschaft der Rassen Mischlinge." Bericht der Deutschen Botanischen Gesellschaft, 1900.

de Vries, Hugo. “Sur la loi de disjonction des hybrides.” Comptes de l’Académie des Sciences, 1900.

von Tschermak, Erich. “Über künstliche Kreuzung bei Pisum sativum.” Zeitschrift für das landwirtschaftliche Versuchswesen in Österreich, 1900.

Mendel, Gregor. Versuche über Pflanzen-Hybriden. 1866. MendelWeg, http://www.mendelweb.org/MWtoc.html.

Darwin, Charles. The Variation of Animals and Plants Under Domestication. Project Gutenberg, https://www.gutenberg.org/ebooks/40854.

Historical analyses

Gasking, Elizabeth B. “Why Was Mendel’s Work Ignored?” Journal of the History of Ideas, vol. 20, no. 1, 1959, pp. 60–84.

Gasking, Elizabeth B. “Why Was Mendel’s Work Ignored?” Journal of the History of Ideas, vol. 28, no. 2, 1967, pp. 167–192.

Olby, Robert. The Origins of Mendelism. University of Chicago Press, 1979.

“Why Was Mendel’s Work Not Accepted at the Time?” ScienceInsights.org, https://scienceinsights.org/why-was-mendels-work-not-accepted-at-the-time/.

“The Rediscovery of Mendel’s Laws of Heredity.” Encyclopedia.com, https://www.encyclopedia.com/science/encyclopedias-almanacs-transcripts-and-maps/rediscovery-mendels-laws-heredity.

“Mendelian Genetics.” EBSCO Research Starters, https://www.ebsco.com/research-starters/history/mendelian-genetics.

“How Many Pairs of Pea Plants Did Gregor Mendel Study?” OGGardenOnline, https://oggardenonline.com/how-many-pairs-of-pea-plants-did-gregor-mendel.html.

Modern DNA evidence

Bhattacharyya, M. K., et al. “The Wrinkled Seed Phenotype of Mendel’s Peas Is Caused by a Transposon-like Insertion in a Gene Encoding Starch-Branching Enzyme.” Cell, vol. 60, no. 1, 1990, pp. 115–122.

“Mendel and Darwin: Untangling a Persistent Enigma.” Proceedings of the National Academy of Sciences, 2022.

van Dijk, P., et al. “How Did Mendel Arrive at His Discoveries?” Nature Genetics, 2022.

“Mendelian Gene Discovery and Molecular Signatures.” Nature Reviews Genetics, 2020.

“The American Journal of Human Genetics Article on Mendelian Trait Analysis.” AJHG, https://www.cell.com/ajhg/fulltext/S0002-9297(19)30274-5.

bioRxiv. “Preprint on Mendelian Trait Modeling.” bioRxiv, https://www.biorxiv.org/content/10.1101/443515v1.full.

NCBI GEO. “GSE225336 Dataset.”

https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE225336.
EBI ENA. “PRJEB3860.”
https://www.ebi.ac.uk/ena/browser/view/PRJEB3860.
NCBI dbSNP. “ABO Blood Group SNPs.”
https://www.ncbi.nlm.nih.gov/snp/?term=ABO+Blood.
OMIM. “Entry 147510 - Mendelian Alleles.”
https://omim.org/entry/147510.
Zenodo. Dataset Record.
https://zenodo.org/records/17245276.
PMC. “Article on Mendelian Genetics.” *PubMed Central*,
https://pmc.ncbi.nlm.nih.gov/articles/PMC12072684/.
Nature. “Mendelian Inheritance PDF.” *Nature*,
https://www.nature.com/articles/s41588-022-01020-3.pdf.
Modern DNA evidence
Kaggle. “Pea Plant Dataset.”
https://www.kaggle.com/datasets/zunorain/pea-plant-dataset.
Datalumos. Project 250593.
https://www.datalumos.org/datalumos/project/250593/version/V1/view.
Google Docs (User Provided).
https://docs.google.com/document/d/1l60xAlXOfX1ZXQ8RFxfI9FIPqo9g-ZOqlylNOrkitrY/edit.
Local HTML File (User Provided).
file:///C:/Users/sport/OneDrive/Desktop/Gregor%20Mendal/Mendel%20Pea/index.html.
Johns Hopkins Medicine. “Lab Study shows Patterns of Inheritance That Defy Mendel’s Laws.”
Johns Hopkins Medicine, 2024,
https://www.hopkinsmedicine.org/news/newsroom/news-releases/lab-study-shows-patterns-of-inheritance-that-defy-mendels-laws.
“When Mendel’s Rules Don’t Apply: Mouse Study Shows Hidden Epigenetic Inheritance.”
Phys.org, 2024,
https://phys.org/news/2024-01-mendel-mouse-hidden-epigenetic-inheritance.html.
“Beyond Mendel: New Study shows the Complexity of Epigenetic Inheritance.” Scientific
Inquirer, 18 Jan. 2024,
https://scientificinquirer.com/2024/01/18/beyond-mendel-new-study-shows-the-complexity-of-epigenetic-inheritance/.
Morgan, Thomas Hunt. The Theory of the Gene. Yale UP, 1926. Internet Archive,
https://archive.org/details/theoryofgene00morg.
Griffiths, Anthony J. F., et al. An Introduction to Genetic Analysis. 11th ed., W. H. Freeman,
2015. NCBI Bookshelf, https://www.ncbi.nlm.nih.gov/books/NBK21949/.
Miscellaneous
Microsoft Corporation. *Visual Studio Code*. Version 1.90, Microsoft, 2024,
https://code.visualstudio.com/.

The Role of Social Connectedness in Veteran Reintegration in the United States
By Seira Watanabe

Introduction

When soldiers leave the military, they first take off their uniform, which is a symbol of rank, role, and identity. Losing it can mean losing the sense of belonging that came with military life. Veterans transitioning beyond the uniform, face challenges rebuilding their identity in civilian life. One of the first challenges that is widely recognized as a major struggle during reintegration is PTSD. According to Mission Roll Call, veterans experience post-traumatic stress at rates up to fifteen times the level seen in civilians. However, reintegration is more than a single clinical condition. Social isolation is one of them, which is linked to serious health consequences that can be as dangerous as "smoking up to 15 cigarettes a day" (Office of the Surgeon General). This makes isolation a critical risk for veterans, as the sudden separation from the military makes them particularly vulnerable to a loss of social connection. In the United States, support for veterans is largely provided through the Department of Veteran Affairs (VA), which provides medical care, disability compensation, housing, employment assistance, education benefits, and insurance. Although these programs are extremely important, there is a concern that it heavily focuses on measurable clinical or structural support. Therefore, challenges such as social connectedness may receive less recognition. In order for veterans to integrate into their community, social interactions play a significant role in fostering a sense of belonging. Repeated attendance in community spaces facilitates social interactions, resulting in a sense of belonging and a 'public familiarity' (Ferreira et al.). Therefore, this paper evaluates the extent to which social connectedness is fundamental to the reintegration of veterans through both social and clinical lenses. The clinical lens provides insight into the strengths and limitations of relying solely on clinical models to address reintegration challenges. The social lens is significant in evaluating how social connectedness is interconnected to veteran wellbeing. Although social isolation during transition from military into civilian life is a global issue, this research will be limited in the scope of veterans in the United States, where despite significant efforts toward reintegration, challenges remain regarding its effectiveness.

Overall, social connectedness is important to a significant extent in the United States, as a clinical approach alone fails to address identity
loss and isolation, which are linked to poor wellbeing outcomes, and community-based initiatives that enhance social connectedness demonstrate improved belonging and increased engagement with support services among veterans.

The Limitations of a Purely Clinical Approach

While current reintegration evaluations focus on measurable processes such as clinical approaches, they do not address all dimensions of reintegration. The current VA reintegration model emphasizes clinical and practical services, including mental health treatment and medical care. By providing only clinical and practical services, the model suggests that current

reintegration services are mainly about symptom treatment and practical support. This overemphasis on clinical approaches overlooks the social factors underlying reintegration, as Keeling, a researcher specializing in military mental health and veteran transition, argues that evaluating transition solely through measurable factors neglects veterans experiencing subclinical mental health challenges. By magnifying clinical and practical challenges, services assume that reintegration difficulties exist only as measurable forms, such as clinical diagnostics. This results in veterans being treated simply as patients rather than individuals who are undergoing complex social and subclinical aspects of transition. Thus, reintegration progress is evaluated mainly through clinical treatment, the social components of reintegration are overlooked.

One significant subclinical challenge veterans face is the loss of identity and belonging when adjusting from a military system into community life. Research conducted by the University of Exeter found that many veterans experience a loss of military identity, leading to lost purpose and sense of belonging, which is often neglected due to the dominant focus on practical support (Maddern). This finding complements Keeling, who argues that solely measurable evaluation overlooks subclinical challenges, by providing evidence that veterans experience challenges rooted by social factors that are neglected by practical support. Therefore, a purely clinical approach to reintegration does not address all dimensions of reintegration, as belonging within a community remains a major challenge veterans face.

The Consequences of Social Isolation

Because a purely clinical approach does not address all aspects of reintegration, veterans continue to experience challenges that stem from social factors. This is shown by data revealing that, despite access to reintegration services, veterans continue to face significant challenges. In a survey of 754 combat veterans, 25% of those utilizing healthcare provided by the VA reported clear difficulties in “social functioning, community involvement, productivity, and self-care” (Gorman et al., "Using Veterans Socials"). Ongoing challenges persist, suggesting treatment alone can’t resolve all factors complicating reintegration. One contributing factor to these ongoing challenges is social isolation. As discussed previously, social isolation is a result of the loss of military identity and social networks that clinical models fail to address. The failure to acknowledge loss of belonging leads to isolation, as seen in veterans who lose established connections after military service (Flack and Kite). It suggests that the structure of military life demands consistent social interaction, which leaves veterans feeling detached when those networks disappear. Research shows that isolated veterans face greater risk of harmful wellbeing outcomes. The Office of the Surgeon General of the United States reports that a lack of social connectedness is directly interconnected with physical health, psychological wellbeing, and an individual’s sense of purpose in life. This complements Flack and Kite who argue that loss of belonging leads to isolation. It further adds that social isolation is strongly associated with veteran wellbeing; however, the root cause, specifically loss of belonging, is often overlooked in clinical evaluation of reintegration progress. If loss of belonging contributes to isolation, and

isolation significantly affects overall wellbeing, then fostering belonging and social connections must be prioritized due to its ripple effect.

However, it can be argued that a clinical model alone is sufficient to treat these social challenges, as they root from neurological or psychological symptoms. Research by Swerdlow et. al. published in the peer reviewed Journal of Traumatic Stress, showed that clinical trials treating PTSD revealed indirect improvements in "Interpersonal functioning," defined as the "quality, satisfaction, and day-to-day functioning within close relationships, such as those with family, intimate partners, and friends." This suggests that clinical models may be sufficient, as treating the cognitive symptoms could indirectly improve social functioning. By framing social isolation as an outcome of clinical symptoms, this perspective implies that the veteran's ability to socially connect will naturally recover as a result of treating the brain. However, this argument is weakened by Keeling, who states that veterans often find reconnecting with their families challenging due to the military's unique norms and time away, which makes adjustment complex. This contrasts with Swerdlow et. al. by highlighting the assumption hidden in the claim that resolving symptoms in the brain naturally resolves complex social challenges, such as adjusting to norms and fitting back into their family. While strong social connections may be hard to achieve without recovery from clinical symptoms, veterans would regardless struggle to connect if they experience a mismatch with the community they return to. This demonstrates that for many veterans, regardless of whether they are experiencing clinical symptoms, the difficulty is in finding a community to reintegrate. Therefore, because social isolation is strongly interconnected with multiple aspects of wellbeing, reintegration would remain limited when it is viewed only through a clinical lens, as it extends beyond what happens in the brain, and is shaped by social connectedness.

Community-based Programs Improving Reintegration Outcomes

If social isolation leads to negative wellbeing outcomes, then initiatives that strengthen social connection require careful examination. Community based programs are emerging initiatives designed as informal gatherings led by veterans that foster social connectedness. One example is the Veterans' Breakfast Clubs, an initiative that provides opportunities for in-person interaction, fosters an inclusive environment through "recognition, good natured gossip, and shared socializing" (McDermott). Through informal gatherings, these programs satisfy the need to belong, which Gorman et al. describes as frequent, "positive interactions with the same individuals over a long period of time". By increasing meaningful interaction, these programs support the sense of belonging, which as a result helps prevent the isolation that stems from the loss of belonging. Because isolation is closely related to negative wellbeing outcomes, satisfying the need to belong indirectly supports mental and social wellbeing.

Another way community based programs support reintegration is by providing a connection to existing care. Currently, less than 50% of the 200,000 veterans returning to civilian life are enrolled in VA healthcare (Mission Roll Call). This suggests that the limitation of existing services is more than their design but also in their accessibility. Even if services

effectively support wellbeing, that alone cannot achieve their intended impact if they are not accessed. However, community based initiatives have shown the ability to increase access to existing services. This was demonstrated in a similar community based intervention, Veteran Coffee Socials, where peer specialists facilitated support groups in local coffee shops. In this community, a veteran who reported substance abuse challenges but had never accessed care through the VA system was able to enroll with the assistance of peers in the group (Gorman et al., "Veteran Coffee Socials"). This complements evidence on low enrollment rates which indicated that although services are available, their lack of reach reduces their impact. By demonstrating that in contrast, community programs that have social connections as their main framework enabled access to a service that would've not have reached them otherwise, it highlights the role of social connectedness in supporting veterans wellbeing. Overall, the success of community based reintegration programs demonstrates that increasing social connectedness satisfies veterans' sense of belonging while also assisting the enrollment of veterans previously unengaged in accessing support. By strengthening belonging, these programs improve wellbeing by preventing isolation. By assisting enrollment among veterans not engaged in support, these initiatives allow reintegration services to achieve their intended impact on veteran wellbeing.

Conclusion

Veteran reintegration in the United States often overlooks subclinical challenges such as loss of belonging. This loss of belonging can lead to isolation, which harms wellbeing. Community based programs show that social connections can prevent isolation and push veterans toward existing services. Ignoring this issue means reintegration efforts remain clinically focused, while clinically undiagnosed veterans struggling socially are overlooked. Ultimately, for successful reintegration, the US must build opportunities for social connectedness in communities.

Works Cited

Ferreira, J., et al. "Spaces of Consumption, Connection, and Community: Exploring the Role of the Coffee Shop in Urban Lives." Geoforum, vol. 119, 2021, pp. 21–29.

Flack, M., and L. Kite. "Transition from Military to Civilian: Identity, Social Connectedness, and Veteran Wellbeing." PLOS ONE, vol. 16, no. 12, 2021, e0261634, https://doi.org/10.1371/journal.pone.0261634.

Gorman, J. A., et al. "Using Veterans Socials to Build a Community: Feasibility of the VOICES Intervention." Community Mental Health Journal, vol. 58, no. 8, 2022, pp. 1544–1553, https://doi.org/10.1007/s10597-022-00969-z.

Gorman, J. A., et al. "Veteran Coffee Socials: A Community-Building Strategy for Enhancing Community Reintegration of Veterans." Community Mental Health Journal, vol. 54, no. 8, 2018, pp. 1189–1197, https://doi.org/10.1007/s10597-018-0288-y.

Keeling, M. "Stories of Transition: U.S. Veterans' Narratives of Transition to Civilian Life and the Important Role of Identity." Journal of Military, Veteran and Family Health, vol. 10, no. 4-EN, 2024, pp. 59–67, https://doi.org/10.3138/jmvfh-0718-0011.

Maddern, K. "New Military Focus on 'Leaving Well' and Strengthening Social Connections Should Be Core Part of Defence Strategy." University of Exeter, 19 Nov. 2025, https://news.exeter.ac.uk/university/new-military-focus-on-leaving-well-and-strengthening-social-connections-should-be-core-part-of-defence-strategy/.

McDermott, J. "'It's Like Therapy but More Fun': Armed Forces and Veterans' Breakfast Clubs: A Study of Their Emergence as Veterans' Self-Help Communities." Sociological Research Online, vol. 26, no. 3, 2020, pp. 512–528, https://doi.org/10.1177/1360780420905845.

Mission Roll Call. "The State of Veterans' Mental Health." 24 Sept. 2024, https://missionrollcall.org/veteran-voices/articles/the-state-of-veterans-mental-health/.

Office of the Surgeon General of the United States. "How Social Connection Impacts Individual Health and Well-Being." Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community, U.S. Department of Health and Human Services, 2023, pp. 23–35.

Swerdlow, B. A., et al. "The Impact of Trauma-Focused Psychotherapy for Posttraumatic Stress Disorder on Interpersonal Functioning: A Systematic Review and Meta-Analysis of Randomized Clinical Trials." Journal of Traumatic Stress, vol. 36, 2023, pp. 496–510, https://doi.org/10.1002/jts.22906.

Agroforestry and Nature-Based Solutions as Tsunami Protection strategies: a Case Study of Sendai By Rio Ito

Abstract

A tsunami is one of the most horrific natural disasters that can cause devastating damage to society and human-made debris. While many tsunami-protection systems are constructed in tsunami-prone countries, some of them are only focused on tsunami protection and are not useful under normal conditions. This paper aims to propose ways to establish more disaster resilient and sustainable systems by introducing nature-based solutions and Agroforestry systems. This paper uses Sendai, where a severe earthquake and tsunami struck in 2011, as a case study to develop a framework that can be adapted to other tsunami-prone areas. Since this paper focuses on sustainable urban planning, several limitations need to be further investigated in future research, such as the lack of qualitative analysis and limited perspectives and evaluation of implementation. This research has the potential to encourage local communities to construct multifunctional tsunam-protection systems based on local conditions and needs. It is essential to make tsunami protection systems not only resilient, but also environmentally sustainable and practical.

1. Introduction

Tsunami disasters, large sea waves caused mainly by earthquakes, are among the most severe natural disasters in the world, such as the 2004 Indian Ocean Tsunami, the 2011 Tohoku Tsunami, the 2018 Palu Tsunami, and the 2018 Sunda Strait Tsunami. (Takabatake et al.) To protect coastal areas from tsunamis, many tsunami protection systems have been constructed, such as protection forests, coastal embankments, evacuation centres and elevated roads. (Takabatake et al.) However, they are often focused only on tsunami mitigation and are not sustainable.

For instance, embankments are often made from steel-reinforced concrete, causing concerns such as poor ecological performance, corrosion of steel in embankments (Hosseinzadeh et al.) and high construction costs. (Lunghino et al.) Protection forests are sometimes monocultures, which are less resilient and have lower biodiversity compared with mixed-species protection forests. (Iwachido et al.) The placement of evacuation centres and elevated roads can become more effective tsunami protection systems by introducing additional vegetation. According to these examples, some of the tsunami protection systems can be improved by developing their sustainability and multifunctionality.

To propose these improvements, this paper aims to answer a question: How can tsunami protection systems protect tsunami-prone areas from tsunamis more effectively while providing
environmental benefits and everyday value? In this paper, two main approaches are proposed: nature-based solutions and Agroforestry systems.

Agroforestry can be defined as a technique to integrate trees, crops, livestock and fisheries

to maximize the benefits of agriculture and livestock industries and improve biodiversity. (Manono and Mwami) This paper focuses on its natural and agricultural aspects and proposes ways to introduce this system in tsunami-prone areas. This can offer everyday value by efficiently using the land (VijayKumar et al.) and providing farmers and residents with additional agricultural fields and incomes.

Nature-based solutions are methods to integrate natural elements like vegetation with tsunami protection systems, such as embankments. (Lunghino et al.) The introduction of nature will offer not only tsunami protection, but also social, economic and environmental services. (Lunghino et al.)

In this paper, some systems related to Agroforestry and nature-based solutions that can be combined with tsunami protection systems are mentioned first. After that, the Sendai area in Japan is chosen as a case study to illustrate proposals for sustainable urban development and to make it easier to apply to other areas. As the term 'sustainable development' can have a wide range of meanings, this paper focuses on three goals: cities that are disaster-resilient, environmentally sustainable, and produce more of local food.

Sendai is one of the areas where the Great East Japan Earthquake, with the highest magnitude of 9.0 (Imamura et al.) and severe tsunamis with the highest height of 13m (Mimura et al.) struck in 2011. To protect the coastal areas from upcoming earthquakes, some tsunami protection systems are constructed. However, there are several limitations and potential improvements, which are discussed in this paper.

Although this research has potential to guide developing sustainable and disaster-resilient areas, many limitations can be considered. Firstly, this paper does not conduct a quantitative analysis, being unable to determine whether advantages outweigh disadvantages. Secondly, while this paper conducts a qualitative analysis of individual design elements, it does not evaluate the benefits, synergy and conflict of combining them. Thirdly, since this paper only considers nature-based and Agroforestry systems, the effectiveness of proposals can be improved by considering other perspectives. Finally, factors that need to be considered to implement the plans, such as regulations, policies and social conditions, are not analyzed in this paper. Those limitations need to be further investigated in future research to fully develop the theory.

This paper is structured as follows. Section 2 briefly describes the research method. Sections 3 and 4 mention some of the Agroforestry and nature-based tsunami protection systems that can form the basis for design elements proposed in Section 5. Section 5 then applies the ideas to the Sendai area by reviewing existing systems and spatially integrating them into the area. Section 6 compares the advantages and disadvantages of the individual proposed design elements and conducts cross-element analysis. Section 7 evaluates the applicability of the proposals to other areas and clarifies some limitations and challenges of this paper that need to be investigated in future research.

2. Methodology

Field research in Sendai helped identify the key issues addressed in this paper: disaster resilience, environmental sustainability, and food production. Agroforestry can be used to make disaster-protection infrastructure useful to local residents both during disasters and under normal conditions by developing food production. In contrast, introducing nature-based solutions can develop disaster resilience in an eco-friendly way.

This paper explored Agroforestry and nature-based solutions first to select elements that can contribute to developing resilience and be integrated in tsunami-protection systems. Since the research problem is quite broad, this paper introduced a case study (Sendai area) first to make it more approachable to propose solutions and generalized the approach through creating a framework. Sustainable systems in the Sendai area were proposed by evaluating conventional tsuami-protection systems in the Sendai area and spatially combining elements based on Agroforestry and nature-based solutions. This paper then conducted qualitative analysis of individual elements to identify potential advantages and disadvantages and potential how these elements can be combined.

3. Agroforestry systems

Agroforestry has a principle that optimizes different resources by complementing trees, agriculture and livestock, leading to high land use efficiency. (VijayKumar et al.) In this section, some of the Agroforestry systems that can help provide everyday values by integrating into tsunami-prone areas are listed in Table 1.

Table 1 shows six Agroforestry systems. Each column represents the name of the systems, their brief explanations, how they can be applied to tsunami protection systems and some countries where each system has already been applied. As livestock-based planning is beyond the scope of this paper, this table excludes Agroforestry systems involving livestock.

Name	Description	Applicability	Case study
1 Alley Cropping	Short-term crops are planted between rows of long-term trees and shrubs. (Manono and Mwami)	This system can effectively grow trees by addressing the problem of the long tree-growing period. As many trees are required in sustainable urban planning, this system can be a method of growing trees. (Manono and Mwami)	Germany (van Ramshorst etal.) Hungary (Honfy et al.)
2 Shelterbelts/ Windbreakers	Trees and shrubs are planted in a row. (Manono and Mwami)	The row of trees can be used not only in preventing soil erosion when a tsunami comes but also for protecting people and crops from strong sea winds under normal conditions. (Manono and Mwami)	California (USA) (Thompson et al.) Southern Quebec (Canada) (Mathieu et al.)
3 Riparian Forest Buffers	Trees, plants and shrubs are planted alongside a body of water like a river, stream and wetland. (Manono and Mwami)	This system can help protect riverside areas from river flooding and conserve aquatic biodiversity. This can protect the aquatic ecosystem by filtering pollutants from agricultural runoff and regulating stream temperature through shading. (Manono and Mwami)	Ontario (Canada) (Gold)

4 Leguminous Cover Crops (LCCs)	Cover Crops, which cover the ground surface, are planted near waterways in most cases. (Suliman and Kassim)	This system can protect soil from the impact of rain and wind and prevent surface runoff and water erosion. (Suliman and Kassim)	California (USA) (Thompson et al.)
5 Homegardens	Trees, crops and sometimes animals are integrated near a dwelling, producing food, fuel and income. (Manono and Mwami)	This system can be implemented easily, compared with other systems. It can provide healthier soils, greater biodiversity and additional agricultural fields and incomes. (Manono and Mwami)	Bangladesh (Gifawesen et al.) Ethiopia (Yaregal and Sime)
6 Hedgerows	Woody vegetation is planted in a row around the agricultural fields. (García De León et al.)	This system can be used as boundaries between paddy fields and provide food and habitats for wildlife. (Clark et al.)	England (Clark et al.) France (Clark et al.) Italy (Clark et al.)

Table 1. Comparison of common Agroforestry systems, how they can be applied to tsunami protection systems and their case studies.

According to this table, each Agroforestry system can offer different everyday values, and some of them can even strengthen tsunami protection. Therefore, the introduction of Agroforestry systems in tsunami-prone areas is likely to develop sustainability and multifunctionality of tsunami protection systems.

4. Nature-Based Tsunami Protection

Nature-based tsunami protection can have a positive effect on the environment while strengthening tsunami protection. Section 4 briefly describes four main nature-based tsunami protection systems: green embankments, mixed-species protection forests, coastal vegetation planting objects, and paddy fields.

4.1 Green Embankment

Embankments are often made from concrete and function as water wave absorption and coastline erosion and flooding alleviation. (Hosseinzadeh et al.) However, concrete embankments often use steel reinforcement to provide extra tensile strength, leading to negative effects on marine ecosystems and water quality caused by corrosion. (Hosseinzadeh et al.) To address these issues, a nature-based embankment called a green embankment (Fig 1) has been proposed after the Great East Japan Earthquake. (Miyawaki) This embankment functions as both environmental protection and disaster mitigation (Miyawaki) by integrating nature with a tsunami protection system. In addition, mixing soil with debris from the earthquake in a mound will help reuse debris. (Miyawaki) This mound planting was first conducted in Miyagi Prefecture, where the Sendai area is located. (Miyawaki) More productive embankments can be built by selecting tree species effectively. (Miyawaki) Short trees and arborescent species are planted to support tall trees in the middle part, and salt-wind-resilient trees are planted along the coastline.(Miyawaki) If a concrete embankment is replaced with a green embankment, this will

mitigate damage from both tsunamis and climate change, thereby being helpful both during disasters and under normal conditions.

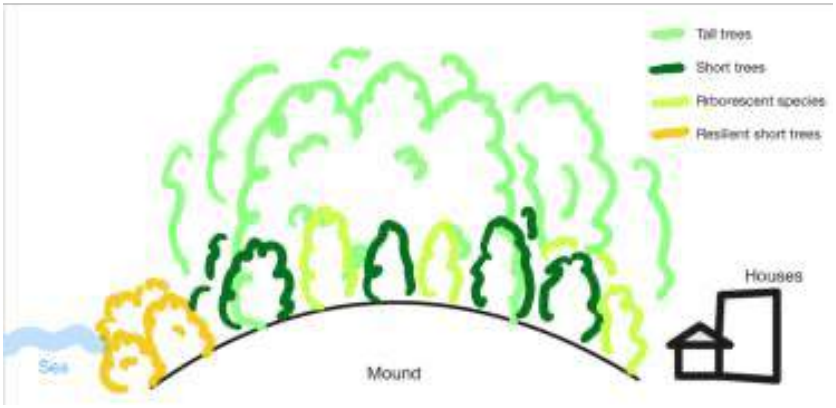

Figure 1. Illustration of green embankment (inspired by(Miyawaki)))

4.2 Mixed-species Protection Forests

A protection forest is one of the common tsunami protection infrastructures. Many trees and shrubs are often planted alongside coastal lines to protect the built environment from natural disasters such as water, soil erosion, floods and landslides. (Zhu and Song) However, some of the protected forests are black pine monocultures, which have less biodiversity and high vulnerability to extreme tsunamis. (Iwachido et al.) Mixed forests (Fig 2) of black pine and broad-leaved tree species (i.e. *Fagus crenata* and *Quercus serrata*) are less vulnerable to damage from tsunamis, conserve biodiversity and improve the landscape. (Iwachido et al.) This can be because each tree species has different root morphology and growth patterns, leading to less competitive and more effective use of underground spaces, light and soil water. (Iwachido et al.) (Thomas et al.) In addition, mixed protection forests with complicated spatial structure are less vulnerable than those with systematic spatial structure, due to higher physical stability caused by greater spatial complementarity, root tensile strength and root biomass across species. (Iwachido et al.) Moreover, to build productive forests, it is essential to carefully select tree species which are suited to site conditions, exhibit healthy growth, and are unlikely to die back. (Zhu and Song)

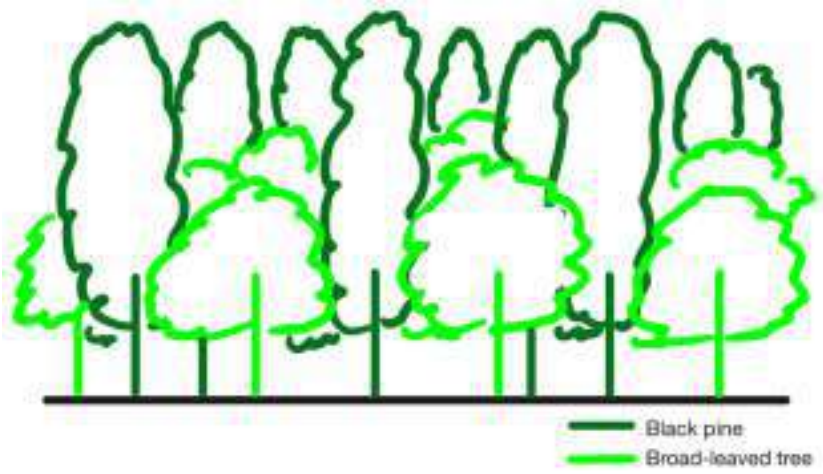

Figure 2**.** Illustration of mixed-species protection forests

4.3 Coastal Vegetation Planting Projects

A coastal vegetation plan can protect human lives and the environment from natural disasters like tsunamis by incorporating trees in coastal areas. (Tanaka et al.) This may provide multiple benefits both during natural disasters and under normal conditions. During a tsunami, the benefits will include (Fig 3): (1) trapping effect: large-diameter trees may prevent human-made debris from washing away and injuring people by trapping them (mangrove forest is one of the examples), (2) escaping effect: some tree species can enable people to climb and escape from a tsunami (tree height requirement depends on the building height), (3) soft landing effect: dense leaves and branches can reduce damage to people caught in water waves and increase probability of survival. (Tanaka et al.) Under normal conditions: (1) beautiful forest landscapes can provide people with shade and adventure activities, and attract tourists and visitors, (2) woody vegetation can protect the shoreline from soil erosion and wave energy, (3) planting food-producing trees will offer farmers agricultural income. (Tanaka et al.) Therefore, trees with dense leaves and branches, large diameter and suitable height can increase the possibility of survival and minimise the damage from disasters and generate several benefits compared with monoculture protection forests. However, additional scientific guidance and data are required to effectively organise those systems. (Tanaka et al.) Integration of agricultural trees with this coastal system can create multifunctional tsunami-protection systems, which can be further discussed in future research, including plant species.

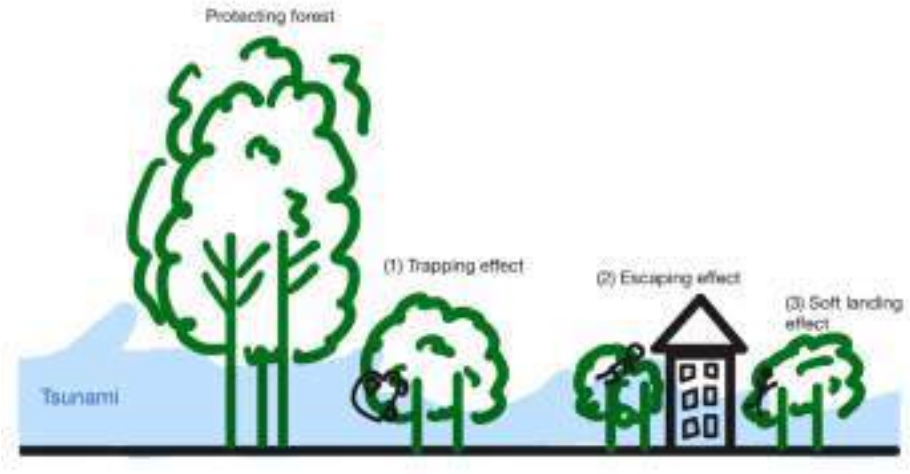

Figure 3. Coastal Vegetation Planning Projects when a tsunami comes (inspired by (Tanaka et al.))

4.4 Paddy Fields

Paddy fields play roles not only I producing rice, but also in other benefits: (1) flood control and soil erosion prevention by increasing water storage capacity during heavy rain, (2) climate mitigation due to vegetation and water surfaces, (3) beautiful traditional landscape, (4) biodiversity by offering habitats for water birds, frogs and fish. (Natuhara) In addition, the use of ecological engineering can enhance the aquatic ecosystem in paddy fields. (Natuhara) A fish ladder is one of them, creating small drops to make elevation changes between drainage ditches and paddy fields more gradual, so that fish can run across one another. (Natuhara) Therefore, the efficient distribution of paddy fields and fish ladders around water resources, such as rivers and the sea, can mitigate tsunami damage and protect local culture. (Natuhara)

In conclusion, the integration of nature with tsunami protection systems can not only provide environmental benefits but also further strengthen tsunami protection. Therefore, introducing those systems can be one of the approaches to improve existing tsunami protection systems.

5. Sendai-focused proposal

This section proposes how to improve tsunami protection systems by using a case study. Section 5.1 mentions representative tsunami protection systems that already exist and are used in the coastal Sendai area. In Section 5.2, ten design elements are proposed based on the systems mentioned in Section 3 and 4, and they are spatially integrated in the Sendai area.

5.1 Tsunami protections in Sendai

The coastal area of the case study is shown in Figure 4, where a severe tsunami was caused by the Great East Japan Earthquake in 2011. To prepare for upcoming earthquakes, several tsunami protection systems have been constructed in this area, including: (1) Concrete Embankment, (2) Protection Forest, (3) Evacuation Hill, (4) Elevated Road and (5) Evacuation

Centre. Below, these five representative systems are discussed.

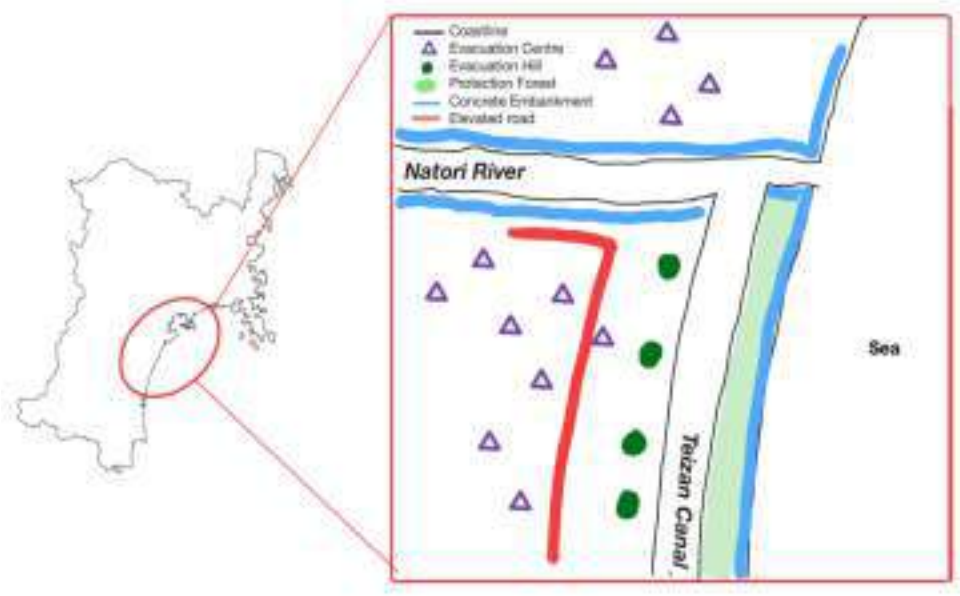

Figure 4. The map of Miyagi prefecture (where Sendai area is located) and simplified illustration of tsunami protection systems in the coastal area (inspired by (Sendai City))

5.1.1. Concrete Embankment

The Embankment in Sendai is about 8.5 km long and its height is 7.2m above T.P. (Tokyo Peil). As shown in Figure 5, air holes and gaps are created between the concrete blocks to prevent them from floating and potentially injuring people. Although this embankment can provide additional time for evacuation, it may not completely prevent a large tsunami like the Great East Japanese Earthquake, whose highest water level was T.P. 13m (Mimura et al.) in the Sendai area. (Sendai City)

Figure 5. Concrete embankment located in Sendai area

5.1.2. Protection Forests

Behind the concrete embankment, there are protection forests (Fig 6), composed mainly of black pine (*Pinus thunbergii*) according to the information board. Black pine is one of the common trees that are used in protection forests in Japan, as they are early colonizers of barren and dry land ecosystems and can grow in harsh coastal environments because of its high tolerance to tidal conditions. (Iwachido et al.)

Figure 6**.** Protection forests besides concrete embankment and an information board showing tree species planted in the forests

5.1.3. Evacuation Hill

In Sendai city, each Evacuation Hill is about 10- 15m high and constructed behind the protection forests. People can use the hills as evacuation sites during a tsunami (Sendai City) (Fig 7).

Figure 7. Evacuation hill located in Sendai area

5.1.4 Elevated Road

An Elevated Road is a roadway built on an embankment-like structure with a length of 10.2 km. (Sendai City) This road is located between large paddy fields, which can also mitigate damage from a tsunami.

5.1.5 Evacuation Centre

In coastal areas, multiple facilities are constructed to help people evacuate from a tsunami: six evacuation towers with a height of about 6m (Fig 8), two outdoor stairs at elementary and junior high schools, and five evacuation buildings incorporating fire brigade stations. (Sendai City) These facilities are strongly built to withstand tsunami force, and emergency supplies are stored inside the buildings, such as generators, blankets, drinking water, food and portable toilets. (Sendai City)

Figure 8. Evacuation tower located in Sendai area

As discussed above, a variety of tsunami protection systems have been planned and constructed, including Green Embankment (mentioned in Section 4.1) which could protect both the environment and coastal areas from a tsunami. However, most of the systems are mainly focused on tsunami protection. They can be improved by introducing Agroforestry and nature-based systems, which can develop their sustainability and multifunctionality.

5.2 Proposals of Design elements

Based on Agroforestry systems (Section 2) and nature-based tsunami protection systems (Section 3), 10 design elements are proposed and spatially integrated in the coastal Sendai area as shown in Figure 9.

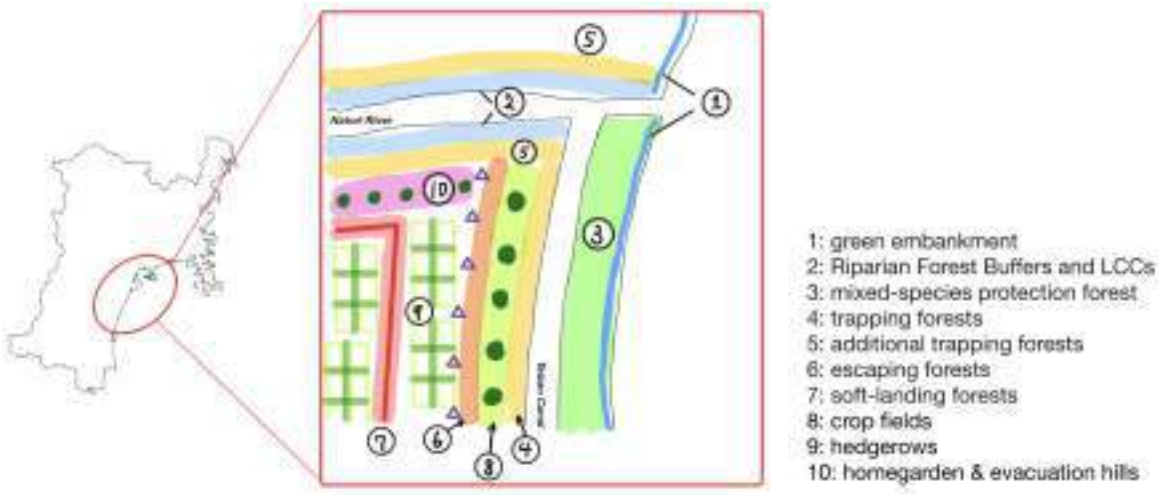

Figure 9. Illustration of 10 design elements spatially integrated in Sendai area

Design elements 1 to 7 are based on Nature-based protection systems, whereas 8 to 10 are based on Agroforestry systems. Design element 1 is green embankment, 2 is Riparian Forest Buffers and LCCs, 3 is mixed-species protection forest, 4 is trapping forests, 5 is additional trapping forests, 6 is escaping forests, 7 is soft-landing forests, 8 is crop fields, 9 is hedgerows, and 10 is homegarden and evaluation hills. Detailed descriptions for each element are listed in section 6.1.

These design elements are not intended to be considered as ten independent elements, but they have the potential to be selected and integrated according to local conditions and objectives when they are applied to disaster-prone areas, including Sendai.

6. Results and Discussion

This section evaluates and analyses the proposed design elements. Section 6.1 qualitatively evaluates individual elements in terms of description, advantages, disadvantages, potential challenges and factors that need to be further investigated. Based on the evaluation, Section 6.2 conducts a comparison across different design elements to identify functions and trade-offs and hypothesises the potential of complementarity of the elements to achieve the aim of this paper and illustrates how they could be combined under different planning priorities.

6.1 Qualitative Evaluation of Individual Elements

Table 2 compares brief descriptions, main advantages and disadvantages that can be caused by introducing each design element, and typical challenges and requirements that need to be addressed to implement the elements

Design elements	Description	Advantage	Disadvantage	Requirements

E1	Replacement of concrete embankment with green embankment (Section 4.1)	More environmentally sustainable (Miyawaki) Biodiversity conservation (Miyawaki) Can recycle debris by incorporating it into the mound (Miyawaki) May provide additional income if agricultural trees are planted (Miyawaki)	Long tree-growing period (Miyawaki) High replacement costs	Accurate estimate of the extent to which green embankment decreases vulnerability to tsunami and impacts of climate change Tree species should be carefully selected to be suited to local conditions
E2	Replacement of concrete embankment with Riparian Forest Buffers and LCCs (Section 3)	More environmentally sustainable (Suliman and Kassim) Reinforce soils by root systems and prevent soil erosion (Suliman and Kassim) Biodiversity conservation (Suliman and Kassim) Can protect the aquatic ecosystem by filtering pollutants (Suliman and Kassim) Can maintain lower water temperature (Manono and Mwami) Can mitigate wind erosion (Manono and Mwami)	High replacement costs Long tree-growing period	Careful selection of cover crop species Accurate estimate of whether combining Riparian Forest Buffers and LCCs provides greater tsunami resilience and environmental benefits than a concrete embankment to determine whether those benefits justify replacement costs
E3	Addition of broad-leaved species into protection forests (Section 4.2) Plant broad-leaved trees with black pines with complicated spatial structures (Iwachido et al.)	Biodiversity conservation (Iwachido et al.) Likely to be less vulnerable to tsunamis (Iwachido et al.) Can trap debris more effectively (Iwachido et al.)	High replacement costs Long tree-growing period	Careful selection of broad-leaved species Determine the most effective number and placement of broad-leaved trees to create strong tsunami protection forests(Iwachido et al.)
E4	Replacement of parks with trees which have a trapping effect (Section 4.3) along Teisan Canal	More environmentally sustainable and greater biodiversity than parks (Tanaka et al.) Can trap and prevent tsunami debris from washing away and potentially injuring people who are on the evacuation hills (Tanaka et al.) Can function as mini protection forests and windbreaks, and protect coastal areas from wind erosion and tsunami (Tanaka et al.)	High replacement costs Long tree-growing period May destroy public recreational places provided in parks May cause further damage to people on the evacuation hills if trees are destroyed by tsunamis	Suitable tree species selection that has large diameter, a high level of resilience and a strong root system that can reduce the risk of uprooting Accurate estimate of the extent to which this system can mitigate the damage from tsunamis and offer additional benefits

E5	Introduction of trees which have an escaping effect (Section 4.3) besides evacuation centres (Section 5.1.5)	Environmentally sustainable Biodiversity conservation Can enable people to ride off from water waves and easily access the evacuation centres by climbing the trees (Tanaka et al.) Can function as extra trapping forests and trap tsunami debris that is not caught by the trapping forests (Tanaka et al.)	High construction costs Long tree-growing period May cause further damage to the evacuation centres and people inside if trees are destroyed by tsunami	Suitable tree species selection with a strong root system, dense foliage and a height similar to that of evacuation centres. (Tanaka et al.) The heights may need to be adjusted if the centres have different heights. Accurate estimate of the extent to which this system can help reduce the damage from a tsunami, including the trapping effect
E6	Introduction of trees which have a soft-landing effect (Section 4.3) alongside elevated road (Section 5.1.4) Planting soft-landing forests on both sides of the elevated road	Can enable people to be intercepted by trees and reach the elevated road, and prevent them from being carried further away (Tanaka et al.) Can function as trapping forests and protect people from being injured by tsunami debris (Tanaka et al.)	High construction costs Long tree-growth period May cause further damage to people evacuated on the road if trees are destroyed by tsunamis	Suitable tree species selection with strong root systems, wide-spreading branches, dense foliage and a height similar to that of the elevated road (Tanaka et al.) Accurate estimate of the extent to which this system can mitigate tsunami and intercept debris and people caught in water waves
E7	Introduction of additional trees that have a trapping effect (Section 4.3) inland of E2	Can function as extra windbreaks and habitats for wildlife, in addition to system 2 Can play a role in protecting residential areas from river flooding caused by tsunamis (Tanaka et al.) Can protect people reached on the evacuation hills from tsunami debris by trapping them (Tanaka et al.)	High construction costs Long tree-growing period May cause further damage to evacuation hills and residential areas if trees are destroyed by tsunami May intersect infrastructure connecting the sea and the town	Infrastructure should be carefully designed to maximise the effectiveness of tsunami mitigation Suitable tree species selection that has large diameter, a high level of resilience and a strong root system that can reduce the risk of uprooting (Tanaka et al.) Accurate estimate to determine whether the combination of E2 and 9 reduces damage from tsunami sufficiently compared with E2 alone to justify the construction costs
E8	Replacement of parks with crop fields between E4 and E5, while retaining evacuation hills This system is inspired by Alley Cropping (Section 3)	Effective use of land by incorporating additional green spaces and evacuation hills (Manono and Mwami) Can allow farmers to earn extra income while waiting for trees (E4 and 5)to mature (Manono and Mwami) E4 and 5 can protect crops and hills from wind and soil erosion	May not have a significant impact on tsunami mitigation Likely to be destroyed when a tsunami comes and causes financial damage to farmers Crop growth rate may decrease due to competition with trees for light and soil water without careful species selection (Manono and Mwami) May destroy public recreation space provided by parks	Careful selection of crop species that is resilient and likely to result in positive consequences with tree species of E4 and 5, regarding competition for light and soil water Accurate estimate of the likelihood of being destroyed once a tsunami comes

E9	Introduction of hedgerows (Section 3) and fish ladders (Section 4.4) Although hedgerows can have a variety of forms and structures, (García De León et al.) they are used as paddy field boundaries in this case	Can show clear boundaries of paddy fields by using green (Clark et al.) Can provide additional habitats and shelters (Clark et al.) Likely to enhance rice growth in paddy fields by storing carbon above-and below-ground and offering other various ecological benefits (Clark et al.) Can help fish run across drainage ditches and paddy fields	Can reduce rice growth by shading and reducing soil moisture (Clark et al.) May not have a significant effect on tsunami mitigation	Proper management and careful species selection are required to minimise the downsides
E10	Introduction of homegarden (Section 3) and evacuation hills (Section 5.1.3) in residential areas	Biodiversity conservation (Manono and Mwami) Can help cultivate sufficient local products, such as crops, fruits and medicine, for self-sufficiency and improve sustainability (Manono and Mwami) May provide farmers with additional agricultural fields and incomes Additional evacuation hills near the houses can increase the likelihood of survival during tsunamis	May not have a significant effect on tsunami mitigation High construction costs, especially evacuation hills	Suitable tree species selection that is easy to grow and produce crops, fruits and medicine. Determine the placement of evacuation hills so that people living anywhere within the area can access them easily

Table 2. Description and comparison of pros and cons of each proposed system and its challenges or requirements

6.2 Cross-Element Analysis

According to the qualitative analysis, each design element has different advantages and disadvantages, meaning that different design elements have different balances of tsunami protection, environmental and everyday functions, with different trade-offs. Trade-offs of each design element can be simplified by categorising based on Agroforestry and nature-based protection.

6.2.1 Trade-offs

The trade-offs can be mainly categorised according to the two approaches. For design elements based on nature-based systems, although these elements have the potential to tackle both tsunami disasters and environmental challenges, their replacement costs, tree-growing period and species selection need to be carefully considered and further investigated in future research. For design elements based on Agroforestry, while they have the potential to provide various benefits under normal conditions, such as food production, biodiversity and financial benefits, they may have limitations in tsunami protection. Understanding these trade-offs can be key to proposing solutions for sustainable urban planning.

6.2.2 Proposal of Integrated Systems

Since this analysis suggests that all design elements have both advantages and disadvantages, this indicates that applying only a single design element is less likely to achieve the aim of this paper, proposing the system that improves disaster resilience, environmental sustainability and local production. From this conclusion, it can be inferred that the diversity of functions has the potential for complementarity among several design elements. However, since complementarity may cause new challenges, such as cost, spatial conflict, ecological interaction and maintenance complexity, this hypothesis needs to be further investigated in future research. Nevertheless, if this hypothesis is assumed to be valid, two main example integrated systems can be proposed: (1) resilience-focused system: the integration of E1, E2, E3, E4, E5, E6 and E7, which prioritises gathering tsunami protection-focused elements near the waterways. (2) multifunctionality-focused system: the integration of E1, E3, E4, E6, E8, E9 and E10, which emphasize everyday-value as well as tsunami protection. These two integrated systems demonstrate the potential flexibility of ten design elements, rather than the advantages of either system.

Overall, the analysis indicates that each design element has different functional strengths and trade-offs, and it is difficult to address multiple social issues by introducing a single element. These findings suggest the potential to complement some elements, but potential benefits, conflicts and synergy need to be evaluated in future research to validate this hypothesis. However, by creating example combinations of the elements under different planning objectives can indicate the flexibility of the design framework that can be applied to other areas, not only the Sendai area.

7. Conclusion

This section concludes the research by summarising its main findings, identifying the factors required to be investigated in future research and considering the applicability of the proposed design elements to other areas.

7.1 Summary

This paper mainly aims to review Agroforestry and nature-based protection and propose ways to incorporate the two approaches to improve the disaster resilience and sustainability of tsunami-prone areas. Systems that can be integrated into sustainable and disaster-resilient proposals are discussed in Sections 3 and 4. In Section 5, 10 design elements are proposed by carrying out a case study. By conducting a qualitative analysis comparing each design element in Section 6, it is indicated that each design element has different strengths and trade-offs, which can be simply categorised as Agroforestry and nature-based protection. From this conclusion, the potential complementarity is hypothesised, and examples of integrated systems are proposed. However, since this paper does not determine the superiority or synergy of the integrated systems, it can be further explored in future research.

7.2 Limitations and Future Research

Although this paper has the potential to improve sustainability and disaster resilience of tsunami-prone areas, there are still many limitations and challenges: (1) Quantitative analysis: There is an insufficient amount of quantitative data that potentially shows the extent to which the mentioned disadvantages of each design element are significant. For example, the numerical data on the length of the tree-growing period, construction or replacement costs, and the reduction in effectiveness of tsunami protection can be measured in future research. (2) System interactions: While this paper proposes and evaluates individual design elements, the potential benefits, synergy and conflict of complementarity and integrated systems should be evaluated in future research. (3) Perspectives: Since this paper mainly considers Agroforestry and nature-based protection as approaches to sustainable urban planning, other perspectives, such as livestock, tidal energy and urban green infrastructure, (Pauleit et al.) can help improve the design elements further. (4) Implementation: Since this paper does not consider factors required for implementation of proposed elements, such as regulations, community engagement, farmers' knowledge, financial resources and policy support, they should be included in future research.

7.3 Broader Applicability of the Design Framework

Since the proposed design elements and integrated systems are adapted to the Sendai area, the broader applicability lies not in introducing them directly, but in adapting them to suit local factors, such as planning purposes, local hazards, regulations, climate, topography and social conditions. Therefore, this paper proposes not a solid solution or design element, but a framework that can be adapted to various factors, namely a design framework. Although the two examples of integrated systems are different according to their planning priorities, both are created through the design framework. The design framework is proposed as follows: (1) identify the current objectives and problems, (2) select the suitable design elements, (3) spatially combine them, (4) evaluate the potential benefits and downsides. The first step should be conducted since tsunami protection systems have the potential to address multiple problems, not only tsunami resilience. In the second step, design elements are selected based on their planning objectives, local conditions and problems. The third step considers how the selected elements can be spatially integrated into specific areas. Finally, the fourth step allows planners to examine whether the proposed configuration is effective by identifying potential advantages, disadvantages, costs and spatial conflicts. This framework has the potential to be used as an iterative process. For example, if the evaluation identifies important trade-offs and limitations, the selected design elements or their spatial integration can be reconstructed or modified by iterating. Applying this framework before implementation could help people carefully plan tsunami protection systems and could increase likelihood of effective improvements in local sustainability and disaster resilience in tsunami-prone areas. However, the effectiveness and suitability of this design framework have yet to be fully evaluated and require further investigation. Nevertheless, this framework provides a basis for adapting multifunctional tsunami protection systems to different tsunami-prone areas while responding to their local

conditions and planning objectives.

In conclusion, this research paper proposes several design elements that can address not only tsunami disasters, but also other social issues, such as climate change, biodiversity and food production. The main contribution of this paper is the proposal of design elements that can potentially be adapted as an integration of design elements and a design framework for selecting and combining them to apply in various areas through providing disaster resilience, environmental benefits and agricultural functions.

8. Acknowledgements

This paper has been completed with the valuable guidance, support and advice from indigo research and my supervisors, Dr. Andrea Giusti and Mr. Marcin Kedziera.

Works Cited

Clark, Katherine, et al. ‘The Multiple Benefits Delivered by Hedgerows: Where Is the Evidence and Does It Meet Current Knowledge Needs?’ People and Nature, vol. 7, no. 10, 2025, pp. 2437–50. Wiley Online Library, https://doi.org/10.1002/pan3.70114.

García De León, David, et al. ‘Contributions of Hedgerows to People: A Global Meta-Analysis’. Frontiers in Conservation Science, vol. 2, Dec. 2021, p. 789612. DOI.org (Crossref), https://doi.org/10.3389/fcosc.2021.789612.

Gifawesen, Sisay Taye, et al. ‘Review on Role of Home Garden Agroforestry Practices to Improve Livelihood of Small Scale Farmers and Climate Change Adaptation and Mitigation’. Journal of Plant Sciences, vol. 8, no. 5, Sept. 2020, pp. 134–45. www.sciencepublishinggroup.com, https://doi.org/10.11648/j.jps.20200805.15.

Gold, Michael A. ‘Preface Developments in North American Agroforestry Research: Selected Papers from the Sixth Biennial North American Agroforestry Conference, 1999’.
Agroforestry Systems, vol. 53, no. 3, Nov. 2001, pp. 3–4. DOI.org (Crossref), https://doi.org/10.1023/A:1013340922134.

Honfy, Veronika, et al. ‘The Effect of Tree Spacing on Yields of Alley Cropping Systems—A Case Study from Hungary’. Plants, vol. 12, no. 3, Feb. 2023, p. 595. www.mdpi.com, https://doi.org/10.3390/plants12030595.

Hosseinzadeh, Nima, et al. ‘Concrete Seawalls: A Review of Load Considerations, Ecological Performance, Durability, and Recent Innovations’. Ecological Engineering, vol. 178, May 2022, p. 106573. ScienceDirect, https://doi.org/10.1016/j.ecoleng.2022.106573.

Imamura, Fumihiko, et al. ‘Recent Occurrences of Serious Tsunami Damage and the Future Challenges of Tsunami Disaster Risk Reduction’. Progress in Disaster Science, vol. 1, May 2019, p. 100009. DOI.org (Crossref), https://doi.org/10.1016/j.pdisas.2019.100009.

Iwachido, Yuki, et al. ‘Mixed Coastal Forests Are Less Vulnerable to Tsunami Impacts than Monoculture Forests’. Natural Hazards, vol. 120, no. 2, Jan. 2024, pp. 1101–12. Springer Link, https://doi.org/10.1007/s11069-023-06248-8.

Lunghino, Brent, et al. ‘The Protective Benefits of Tsunami Mitigation Parks and Ramifications for Their Strategic Design’. Proceedings of the National Academy of Sciences, vol. 117, no. 20, May 2020, pp. 10740–45. pnas.org (Atypon), https://doi.org/10.1073/pnas.1911857117.

Manono, Bonface O., and Boniface Mwami. ‘Agroforestry and Soil Health: A Review of Impacts and Potential for Sustainable Agriculture’. Earth, vol. 7, no. 1, Feb. 2026, p. 31. www.mdpi.com, https://doi.org/10.3390/earth7010031.

Mathieu, Antoine, et al. ‘Drivers of Tree Establishment in Planted Windbreaks and Riparian Buffers: A Case Study of Farms in Southern Quebec, Canada’. Geoderma Regional, vol. 37, June 2024, p. e00788. DOI.org (Crossref), https://doi.org/10.1016/j.geodrs.2024.e00788.

Mimura, Nobuo, et al. ‘Damage from the Great East Japan Earthquake and Tsunami - A Quick Report’. Mitigation and Adaptation Strategies for Global Change, vol. 16, no. 7, Oct. 2011, pp. 803–18. Springer Link, https://doi.org/10.1007/s11027-011-9297-7.

Miyawaki, Akira. ‘The Japanese and Chinju-No-Mori Tsunami-Protecting Forest after the Great East Japan Earthquake 2011’. Phytocoenologia, vol. 44, nos 3–4, Oct. 2014, pp. 235–44. DOI.org (Crossref), https://doi.org/10.1127/0340-269X/2014/0044-0571.

Natuhara, Yosihiro. ‘Ecosystem Services by Paddy Fields as Substitutes of Natural Wetlands in Japan’. Ecological Engineering, vol. 56, July 2013, pp. 97–106. DOI.org (Crossref), https://doi.org/10.1016/j.ecoleng.2012.04.026.

Pauleit, Stephan, et al. ‘Advancing Urban Green Infrastructure in Europe: Outcomes and Reflections from the GREEN SURGE Project’. Urban Forestry & Urban Greening,

Urban green infrastructure – connecting people and nature for sustainable cities, vol. 40, Apr. 2019, pp. 4–16. ScienceDirect, https://doi.org/10.1016/j.ufug.2018.10.006.

Sendai City. Sendai City Earthquake Disaster Reconstruction Plan -Digest Version-. 2011, https://www.city.sendai.jp/shinsaifukko/shise/daishinsai/fukko/kanren/documents/fukkokeikakugaiyou.pdf.

Suliman, Nurul Shazuwani, and Nur Qursyna Boll Kassim. ‘Enhancing Soil Chemical Properties at Plantation Riparian Buffer Zone Using Leguminous Cover Crops: A Review’. IOP Conference Series: Earth and Environmental Science, vol. 1019, no. 1, Apr. 2022, p. 012003. Institute of Physics, https://doi.org/10.1088/1755-1315/1019/1/012003.

Takabatake, Tomoyuki, et al. ‘Simulated Effectiveness of Coastal Forests on Reduction in Loss of Lives from a Tsunami’. International Journal of Disaster Risk Reduction, vol. 74, May 2022, p. 102954. ScienceDirect, https://doi.org/10.1016/j.ijdrr.2022.102954.

Tanaka, Norio, et al. ‘Coastal Vegetation Planting Projects for Tsunami Disaster Mitigation: Effectiveness Evaluation of New Establishments’. Landscape and Ecological Engineering, vol. 7, no. 1, Jan. 2011, pp. 127–35. Springer Link,

https://doi.org/10.1007/s11355-010-0122-3.

Thomas, Anaïs, et al. ‘Growth Dynamics of Fast-Growing Tree Species in Mixed Forestry and Agroforestry Plantations’. Forest Ecology and Management, vol. 480, Jan. 2021, p. 118672. DOI.org (Crossref), https://doi.org/10.1016/j.foreco.2020.118672.

Thompson, Jennifer B., et al. ‘Remote Sensing of Hedgerows, Windbreaks, and Winter Cover Crops in California’s Central Coast Reveals Low Adoption but Hotspots of Use’. Frontiers in Sustainable Food Systems, vol. 7, Jan. 2023. Frontiers,

https://doi.org/10.3389/fsufs.2023.1052029.

van Ramshorst, Justus G. V., et al. ‘Reducing Wind Erosion through Agroforestry: A Case Study Using Large Eddy Simulations’. Sustainability, vol. 14, no. 20, Oct. 2022, p. 13372. www.mdpi.com, https://doi.org/10.3390/su142013372.

VijayKumar, R., et al. ‘Agroforestry Systems: A Pathway to Resilient and Productive Landscapes’. International Journal of Environment and Climate Change, vol. 14, no. 12,

Dec. 2024, pp. 177–93. journalijecc.com, https://doi.org/10.9734/ijecc/2024/v14i124617.
Yaregal, Yihunie, and Getachew Sime. ‘Traditional Home Garden Agro-Biodiversity Dynamics, Agro-Ecosystem Services, and Management Practices in Smallholder Farmers Setting, South-Central Ethiopia’. Food and Energy Security, vol. 13, no. 4, 2024, p. e569. Wiley Online Library, https://doi.org/10.1002/fes3.569.
Zhu, Jiaojun, and Lining Song. ‘A Review of Ecological Mechanisms for Management Practices of Protective Forests’. Journal of Forestry Research, vol. 32, no. 2, Apr. 2021, pp. 435–48. DOI.org (Crossref), https://doi.org/10.1007/s11676-020-01233-4.

Why Mid-Sized Firms Lag: Barriers to AI-Based Facility Management Adoption in Indian Commercial Real Estate — Maharashtra as a Case Study By Aadith Palresha

ABSTRACT

Artificial intelligence is starting to change how commercial buildings are run, helping predict equipment breakdowns, cut energy use, and handle tenant requests automatically. Large real estate firms are already adopting these tools, but mid-sized commercial real estate firms in Maharashtra have not been adopting at the same rate. This paper shows what stops mid-sized firms from adopting AI-based facility management tools, and how their reasons differ from large firms. Existing research studies only small or only large firms, so mid-sized firms in India have barely been studied despite managing most of the state's buildings. The paper uses the Technology-Organisation-Environment (TOE) framework, explaining adoption through its three parts: how well the technology fits the firm, how ready the firm is inside, and the impact of outside pressures on it.The research shows, mainly a lack of trained staff and low trust in AI, matter more than technology or outside pressure, while cost is present in every barrier. It also explains what AI facility management offers, and argues mid-sized firms face different problems than large firms, not in size but degree.

1. INTRODUCTION

Why Mid Sized Firms Lag: AI Facility-Management Adoption Barriers in Maharashtra

Aadith Palresha - 2026

Artificial intelligence is beginning to transform how commercial buildings are managed, from predicting equipment failures to optimising energy use and automating tenant services (Sampaio, Costa, and Flores-Colen 2022; van Roosmale et al. 2024). Yet across Maharashtra's commercial real estate sector, the midsized firms managing much of the commercial space aren't familiar with these tools leaving them behind from large firms (Rizwan and Hussain 2024; Patil and Swaroop 2025).

AI based facility management covers predictive maintenance, energy optimisation, and the automation of tenant services (Sampaio, Costa, and Flores-Colen 2022; Hakimi, Liu, and Abudayyeh 2024). Existing research on AI adoption in small and medium enterprises has found that the main obstacles are not simply cost related but have other criterias which this paper identifies using the TOE

(Technology-organizational-environmental) framework (Rizwan and Hussain 2024; Sánchez, Calderón, and Herrera 2025).Moreover, research shows that technological, organizational and environmental differences are significant between large and small firms, with larger firms having an advantage due to various different reasons the paper discusses later leading them to adopt AI quicker (Rizwan and Hussain 2024; Sánchez, Calderón, and Herrera 2025). Research on AI adoption in small and medium-sized enterprises identifies organisational readiness, employee capabilities and access to technical expertise as important determinants of

successful adoption (Sánchez, Calderón, and Herrera 2025; Ganguly et al. 2025). In India, these challenges are compounded by limited digital infrastructure, RERA (Real Estate Regulatory Authority) regulation, and data privacy concerns (Ken Research 2025), while Indian PropTech research points to low trust and high service costs as leading concerns (Darshan 2024).

This literature paper investigates the main barriers preventing mid-sized commercial real estate firms in Maharashtra from adopting AI based facility management tools and how these are different from the challenges facing larger firms. The question matters because, existing research concentrates on either large firms or SMEs in Western and Gulf economies, leaving mid sized firms in Indian markets like Maharashtra largely unexamined, even though these mid sized providers account for the majority of the Indian facility management market (Mordor Intelligence 2026). If the barrier is misidentified, proptech aimed only at lowering cost will fail as cheaper or more accessible tools cannot solve a problem unless examined thoroughly, and large firms that do adopt AI are pulling ahead while mid sized firms are left further behind increasing the gap (Mathagu 2024).

Although plenty of Indian mid sized firms are familiar with the advantages of AI (Patil and Swaroop 2025), adoption among mid sized commercial real estate firms in Maharashtra remains low, a gap this paper wants to solve. This paper also focuses on

not only to showcase the barriers of adopting AI in property management but also to show how AI can be used in commercial real estate which will help firms consider using AI and also help Tech firms develop tools to solve the problem rooted into our system.

Some barriers preventing AI adoption are government policies, low expertise in AI based tools, limited availability and quality of past operational data and lack of trust in AI driven solutions. However, past research suggests that the cost equation is also a recurring barrier most mid sized commercial real estate firms face as they are incapable of handling huge tool costs or hiring specialised AI employees.

2. THEORETICAL FRAMEWORK

2.1 Introduction of TOE

This paper analyses adoption through the Technology–Organization–Environment framework, which explains technology adoption at the level of the firm rather than the individual user (Tornatzky and Fleischer 1990). The framework suggests that adoption is based on three criterias which include the characteristics of the technology, the internal conditions of the organisation, and the external environment—and that the framework does not apply wherever any one of these is unlikely, however likely the other two may be. The depth of the strategy is what makes it appropriate for my present research question as it has a high success rate.Even though there are alternative models such as TAM (Technology Acceptance Model) and UTAUT (Unified Theory of Acceptance and Use of Technology) explain whether individual users will accept a system once it has been installed in the firm which is a different question from whether a firm will install it at all, and adoption studies in emerging markets have

increasingly moved to TOE for this reason (Chittipaka et al. 2023). The decision to deploy AI based facility management in a mid sized commercial real estate firm is not made by the facility manager who will use the tool but is made by senior management involving a sufficient amount of capital (Soomro et al. 2025). TOE is built for decisions of that shape helping firms adopt and use AI to their advantage.

2.2 Its three dimensions: Technology, Organisation, Environment

The technological dimension refers to how a company views a new technology before deciding whether to adopt it (Tornatzky, Fleischer, and Chakrabarti 1990). It checks the cost of the new PropTech, complexity, and how compatible it is with existing models in the firm. For AI facility management this is largely a question of whether new tools can

reach the building management as most existing tool aren't compatible with HVAC (Heating, Ventilation, and Air Conditioning) and metering systems a firm already operates, and the sector's long standing network problem makes that connection difficult (Domingues et al. 2016; van Roosmale et al. 2024). The organisational dimension concerns internal conditions inside the firm such as firm size, available resources, in house technical skill, management orientation, and how the firm's existing digital infrastructure is structured. Moreover, organizational barriers are an important dimension that influences AI adoption more strongly than either the technology itself or environmental pressure (Sánchez et al. 2025). The environmental dimension concerns conditions outside the firm such as vendor ecosystem, competitive pressure, government regulations. These carry particular weight in Maharashtra, where government or RERA policies and unstable internet and power supply mainly contribute to external pressure limiting the adoption of AI in property management (Ministry of Communications 2026). Moreover, low trust and high service costs remain the leading concerns within Indian PropTech (Darshan 2024).

3. LITERATURE REVIEW

3.1 Technological barriers

One of the primary technological barriers preventing mid-sized commercial real estate firms in Maharashtra from adopting AI based facility management tools is the limited interoperability between existing building infrastructure and newly constructed AI Proptech. AI applications such as predictive maintenance, energy optimisation, and automated fault detection depend on daily high quality data collected from existing building management systems (BMS), HVAC equipment and elevators. However, many commercial buildings continue to rely on legacy systems that use proprietary communication protocols, making it difficult for AI platforms to access and integrate operational data (Sampaio, Costa, and Flores-Colen 2022). To fix compatibility requires costly middleware, software customisation, or replacement of existing equipment before AI solutions can function effectively (van Roosmale et al. 2024). These costs are burdensome for mid-sized firms which manage smaller properties and have fewer financial

resources than large real estate companies, leaving mid-tier providers with less resources to invest in technology than larger firms (Mordor Intelligence 2026; Ganguly et al. 2025). As a result, many firms delay AI adoption because their existing infrastructure cannot support advanced digital technologies without substantial investment incentivising AI adoption (Sampaio, Costa, and Flores-Colen 2022).

A second technological barrier is the limited availability and quality of historical operational data required to train AI models (Hakami 2024). Predictive maintenance systems rely on large datasets containing equipment performance, maintenance histories, and failure records to accurately identify patterns and forecast future breakdowns (Hakami 2024). Mid sized commercial real estate firms generally manage fewer buildings than larger organisations, resulting in fewer equipment operating hours and a smaller number of recorded failures from which AI systems can learn. This creates the challenge of data scarcity reducing the efficiency and trust in AI predictions (Hakami 2024). Furthermore, even mid sized firms with sufficient financial resources may struggle to implement AI successfully because inadequate data limits and system performance can also lead to reduced confidence in AI systems (Deloitte 2025). Beyond the technology itself, the literature points to an even bigger obstacle inside the firm:organisational barriers.

3.2 Organisational barriers

Organisational barriers, as research suggests, are the most significant factors limiting the adoption of AI based facility management systems in mid sized commercial real estate firms in Maharashtra. Research shows that organisations being ready, including organizational support, digital resources and employee skills, have a greater influence on AI adoption than technological or environmental factors (Sánchez et al. 2025). Research on AI adoption in SMEs identifies limited technical expertise, organisational readiness and insufficient digital capabilities as important barriers to adoption (Sánchez, Calderón, and Herrera 2025; Ganguly et al. 2025). In the commercial real estate sector, the shortage of internal AI knowledge is often considered a bigger obstacle than cost, making firms unlikely to invest in AI's (Dealpath 2025; Rizwan and Hussain 2024).

Another important barrier is the lack of trust in AI-driven decisions (Dietvorst et al. 2015; Mahmud et al. 2023). Facility managers mainly make decisions based on their own experience and are less likely to trust AI solutions, especially when some errors are seen.This reflects the concept of tech distrust, as people are more likely to reject advice on AI adoption after seeing errors (Dietvorst et al. 2015). Moreover, around 29% of companies report a shortage of talent skilled enough to manage AI tools, and without employees who know how to use and manage these systems, organisations find it difficult to trust them, delaying their adoption in facility management (Dealpath 2025). While these gaps talk about the firm from the inside, environmental barriers come from outside entirely.

3.3 Environmental barriers

In Maharashtra, there are plenty of environmental barriers affecting adoption of proptech AI in mid sized commercial real estate for property management. Firstly due to government policies, laws like India's Digital Personal Data Protection Act (2023) restrict how real estate platforms capture and utilise consumer behavioural data, and compliance with cybersecurity rules and digital infrastructure standards increasing implementation costs of PropTech.(Ministry of Electronics and Information Technology 2023). Moreover, as India's AI governance framework is still developing, the absence of a dedicated AI law creates regulatory uncertainty that discourages organisations from making long-term investments in AI technologies (Gupta 2026). Moreover, poor digital connectivity across Maharashtra continues to limit AI adoption, particularly in the remote areas of Maharashtra (Ministry of Communications 2026). AI based facility management systems require consistent high speed internet to give real time data from IoT sensors and building management systems (Hakimi, Liu, and Abudayyeh 2024). Poor network reliability disrupts predictive maintenance reducing the effectiveness of AI solutions (Ministry of Communications 2026).

3.4 The cost equation

Outside the TOE framework, cost remains one large barrier quietly contributing to all three sections of the framework (Ganguly et al. 2025; Rizwan and Hussain 2024). The first barrier is technology, as this paper discusses, it is really expensive to adopt AI based facility management systems, making it hard for mid sized firms in Maharashtra to purchase these tools and delaying adoption. Research on Indian MSMEs identifies the price of automation as one of the driving barriers to AI adoption rather than a secondary consequence of it (Ganguly et al. 2025), and a NITI Aayog assessment found that 82% of medium enterprises still rely on outdated technology, limiting their capacity to invest in AI tools, computing infrastructure, and training alike (NITI Aayog and Institute for Competitiveness 2025).

For organisation, firms need to hire specialised workers who can manage the tools for them. Hiring specialised workers in the PropTech field is usually costly and hard for mid sized firms demotivating them from adopting new AI tools. Similarly, research on Indian SMEs found that 70% lack access to skilled AI professionals, while 72% showed the need for dedicated AI training programs (Patil and Swaroop 2025).

Lastly, unreliable connectivity in rural parts of Maharashtra forces firms to invest in stable network infrastructure before AI systems can operate without disturbance (Ministry of Communications 2026) .Compliance with the Digital Personal Data Protection Act requires consent management, breach reporting and security safeguards that carry an ongoing operational cost (Ministry of Electronics and Information Technology 2023).

4. What AI facility management offers?

4.1. PREDICTIVE MAINTENANCE

Predictive maintenance is a strategy that helps the industries to detect the failure of

equipment before it occurs and to know exactly when it is going to happen and which part of the system is going to be affected (Es-sakali et al. 2022). AI-based predictive maintenance works in two main stages (Es-sakali et al. 2022).

In the first stage, small sensors are attached to building equipment such as air conditioning units and lifts. These sensors constantly measure things like how much the machine is shaking, how hot it is running, and what sounds it makes and sends this information to a computer system along with records of past repairs and breakdowns (Es-sakali et al. 2022).

In the second stage, the computer briefly studies this information processed by the small sensors and learns how the machine looks when it is working normally. Once it knows the basic pattern, then it can alert the facility manager once something starts to change in a way that usually happens before the equipment breaks down, helping warn the facility manager early, so the repairs can be done before the equipment actually fails (Es-sakali et al. 2022; Hakami 2024).

Recently, a review of predictive maintenance systems applied to HVAC systems found that these systems can reduce maintenance costs by 25% to 35%, cut breakdowns by 70% to 75% and shorten breakdown duration by 35% to 45% helping the business save time and decrease costs (Es-sakali et al. 2022).

Lastly, in a commercial building equipment failures such as air conditioning or lift breakdowns affect the people living there directly, and research previously conducted on office buildings has found that the conditions of a building, including its services and management, greatly influences tenant satisfaction, which in turn drives the decision to renew their lease. (Cheah et al. 2014; Lee et al. 2021).

4.2 ENERGY OPTIMIZATION

Energy optimization means running a building's systems in a way that uses the least energy possible while still keeping the people inside comfortable. It matters because buildings account for about 40% of the world's energy use and air conditioning takes up the largest share of that (Ekanayaka Gunasinghalge, Alazab, and Talukder 2025).

Normally, a building runs on fixed settings. Under fixed settings, an air conditioner works from 8am to 10pm at the temperature of 24 degrees throughout the day. At 6pm half the staff has gone home, the 3rd floor barely consists of 20 percent of the full capacity. Still,

the air conditioner is still working at 24 degrees, the system cools those spaces exactly as efficiently as full when it's not needed. AI changes this by using sensors to check how many people are inside, how hot it is outside, and how efficiently the equipment is working. The system then adjusts the cooling and lighting continuously through the day to match what the people inside the building actually need at that moment, instead of following a fixed schedule (Ekanayaka Gunasinghalge, Alazab, and Talukder 2025).

A systematic review of 126 research papers covering 45 countries and more than 2.8 million hours of building data found that AI-based optimisation reduced air-conditioning energy consumption by approximately 19.7% and lighting energy consumption by 15.3%. The highest-performing systems reported overall energy savings of between 22% and 28%

(Ekanayaka Gunasinghalge, Alazab, and Talukder 2025). However, these findings represent international evidence rather than measured results from commercial buildings in Maharashtra. Although they demonstrate the potential of AI-based energy optimisation, the actual savings achievable by Maharashtra's commercial real estate firms would depend on building design, climate conditions, occupancy patterns and existing infrastructure.

4.3 TENANT SERVICES

Tenant services are the everyday support and communication a building's management team provides to the businesses renting the space, things like fixing something that's broken, letting tenants book shared spaces, and keeping them updated on requests they've already made. Basically, it's how a landlord or management team looks after tenants day to day, beyond just collecting rent (Cove n.d.).

Normally this is all done by people. A tenant calls or emails the facility manager during office hours, someone writes the problem down and passes it on to a technician. If a tenant reports a problem at night or on a weekend, nothing happens until the office opens again, which could significantly make the problem worse.

AI changes this by dealing with the requests on its own. A tenant sends a message at any time of day and the AI chatbot reads it, works out what kind of problem it is, creates a solution, sends it to the right technician, and keeps the tenant updated while the work is being done (Sampaio, Costa, and Flores-Colen 2022). This matters for business reasons too, because studies of office buildings show that how well a building is run affects how happy the tenants are, and happy tenants are more likely to stay and renew their lease (Cheah et al. 2014; Lee et al. 2021).

There is another benefit as well. Because every request is saved in one place, the commercial building slowly builds up a record of which machines break most often and how long repairs usually take. This is exactly the kind of past information that predictive maintenance needs in order to work. So these three tools are not really three separate systems, but three parts of one system that feed each other information and work collaboratively (Hakami 2024; Hakimi, Liu, and Abudayyeh 2024).

5. MID SIZED VS LARGE FIRMS IN REAL ESTATE

5.1 What counts as mid-sized

A mid sized commercial real estate firm is large enough to manage its own buildings competently with its own staff but not large enough to keep specialist teams for things such as property management or technology. It works across one city or state rather than the whole country (Sánchez, Calderón, and Herrera 2025; Soomro et al. 2025).

Most mid sized firms have three features.They usually work within one city or state and operate by being close to their clients rather than by being big and also avoiding diseconomies of scale. They do not have separate technology teams or specialists, so any specialized work is handled by the same people who run the daily tasks. Lastly, mid sized firms in commercial real

estate do not have spare money set aside to invest in new technology the way larger companies do (Mordor Intelligence 2026).

The Indian government treats a company as medium sized if it has invested up to ₹125 crore and earns up to ₹500 crore a year (Ministry of Micro, Small and Medium Enterprises 2025).Yet, these limits were written mainly for factories where the money goes into machines rather than buildings. For a property firm, the number of buildings it manages tells you more about its size than its earnings do. Having defined what counts as mid-sized, the real question is what separates these firms from large ones.

5.2 Resources large firms have that mid sized firms don't

Large firms consistently adopt AI at higher rates than mid-sized ones. OECD data from 2025 puts average AI adoption among firms at around 20%, with larger, better resourced firms adopting at notably higher rates than smaller and mid-sized ones (OECD 2025). The reason is not that large firms are simply richer but that they hold two things mid sized firms do not separate them.

The first resource is specialist staff. A study of 836 firms found that those employing IT specialists and running their own IT training for its employees adopted AI at higher rates than those that did not (Sánchez, Calderón, and Herrera 2025). Large commercial real estate firms keep technology teams whose only job is to run digital systems. In a mid sized firm, the person running daily facility management is also the person expected to handle any new Proptech.

The second resource is recorded information. Because large firms manage far more buildings, they have recorded far more equipment breakdowns over the years giving them access to more usable data. AI systems learn from these past records and their

accuracy drops when the data size is small, since these learning models require large volumes of data to identify accurate failure patterns. (Hakami 2024; Ucar, Karakose, and Kırımça 2024).

5.3 Why the barriers differ in kind, not just degree

The problems mid sized firms face come in two types. Some are problems of degree, which means the same problem a large firm has however felt more strongly. Moreover, higher amounts of money will solve this problem, these problems are also called money related problems. Others are problems of some kind which come from the size of the firm, so money cannot solve the problem as it isn't the only missing thing.

Cost and a lack of trained staff are problems of degree. If a mid-sized firm had more money, it could buy the software tomorrow and pay someone to run it. The size of the firm does not stop this from working. A large firm has to pay for the same two things, it just has more money to pay with (NITI Aayog and Institute for Competitiveness 2025).

Past records are a problem of kind. AI systems learn by studying what happened before a machine failed so they need a long history of breakdowns to work from and their accuracy reduces sharply when that history is not sufficient (Hakami 2024). A smaller portfolio produces

fewer breakdowns simply because fewer machines have been running for fewer hours. Money cannot shorten this, because the records build up only through years of actually operating buildings. These are organizational based problems based on kind rather than degree.

6.0 CONCLUSION

6.1 The main barriers the literature identifies.

This review found barriers across all three parts of the TOE framework but they are not equal in weight. The organisational barriers carry the highest weight as a firm's internal skills and it being ready affect adoption more strongly than the technology itself or outside pressure (Sánchez, Calderón, and Herrera 2025), and in commercial real estate a shortage of internal knowledge is treated as a bigger obstacle than budget (Dealpath 2025). The technological barriers are older building systems that AI tools cannot easily connect to (van Roosmale et al. 2024) and a shortage of past failure records for those tools to learn from (Hakami 2024). The environmental barriers are unsettled AI regulation and uneven connectivity in rural areas of Maharashtra (Ministry of Electronics and Information Technology 2023; Gupta 2026; Ministry of Communications 2026). Cost runs through all three rather than standing apart from them (Ganguly et al. 2025).

6.2 The remaining research gap

Literature generally explains why small firms and large firms adopt AI differently, but it barely talks about the firms in between. Studies of SMEs treat mid sized firms as scaled up small businesses and studies of large firms treat them as scaled down large ones. Neither asks whether the middle has its own problems.This distinction is particularly relevant to Indian commercial real estate, where the facility management sector includes firms operating at different scales. However, the precise contribution of mid-sized providers to the overall market remains insufficiently established in the literature reviewed here.

The second gap relates to the first. Because mid-sized firms are treated as a smaller version of large ones, their barriers are assumed to be the same barriers with less money used to solve them. This paper suggests otherwise. Nobody has tested where the line between these two types of barrier falls, or how many buildings a firm needs before AI tools become reliable for it.

6.3 Future Research

Future research's next step is to measure the threshold. Research should establish how many buildings or how many years of operating records a firm needs before predictive maintenance tools become highly accurate and reliable. Lastly, this paper works from existing literature, it has been really hard to find specific interviews or articles about the Maharashtra commercial real estate market. Future research should carry out primary interviews or survey work with mid sized firms in Maharashtra to test whether the barriers identified here actually work realistically.

Works Cited

Cheah, Jun-Hwa, Siew-Imm Ng, Chern Li Lee, and Kevin Guan Choo Teoh. 2014. "Assessing Technical and Functional Features of Office Buildings and Their Effects on Satisfaction and Loyalty." International Journal of Economics and Management 8 (S): 137–76. https://core.ac.uk/works/31086392.

Chittipaka, Venkataiah, Satish Kumar, Uthayasankar Sivarajah, Jana Lay-Hwa Bowden, and Manish Mohan Baral. 2023. "Blockchain Technology for Supply Chains Operating in Emerging Markets: An Empirical Examination of Technology-Organization-Environment (TOE) Framework." Annals of Operations Research 327 (1): 465–92. https://doi.org/10.1007/s10479-022-04801-5.

Cove. n.d. "Building Operations." Cove Glossary. Accessed August 21, 2026. https://cove.is/glossary/building-operations.

Darshan, P. 2024. "Proptech: A New Era of Real Estate Industry." International Journal of Research Publication and Reviews 5 (11): 4696–4706. https://ijrpr.com/uploads/V5ISSUE11/IJRPR35270.pdf.

Dealpath. 2025. "Survey Reveals CRE's AI Paradox: Leaders Accelerating AI Adoption Despite Persistent Challenges." Business Wire, September 30, 2025. https://markets.financialcontent.com/observernewsonline/article/bizwire-2025-10-7-survey-reveals-cres-ai-paradox-leaders-accelerating-ai-adoption-despite-persistent-challenes.

Deloitte. 2025. 2025 Commercial Real Estate Outlook. Deloitte Insights. https://www.deloitte.com/us/en/insights/industry/financial-services/commercial-real-estate-outlook.html.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey. 2015. "Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err." Journal of Experimental Psychology: General 144 (1): 114–26. https://doi.org/10.1037/xge0000033.

Domingues, Pedro, Paulo Carreira, Renato Vieira, and Wolfgang Kastner. 2016. "Building Automation Systems: Concepts and Technology Review." Computer Standards and Interfaces 45: 1–12. https://doi.org/10.1016/j.csi.2015.11.005.

Ekanayaka Gunasinghalge, Lakpriya Udayanga Gunasena, Ammar Alazab, and Md. Alamin Talukder. 2025. "Artificial Intelligence for Energy Optimization in Smart Buildings: A Systematic Review and Meta-analysis." Energy Informatics 8: 135. https://doi.org/10.1186/s42162-025-00592-8.

Es-sakali, Niima, Mohamed Cherkaoui, Moha Oualid Mghazli, and Zakaria Naimi. 2022. "Review of Predictive Maintenance Algorithms Applied to HVAC Systems." Energy Reports 8: 1003–12. https://www.sciencedirect.com/science/article/pii/S2352484722013944.

Ganguly, Boudhayan, Sahana Roy Chowdhury, Rajarshi Ghosh, and Dona Ghosh. 2025. "What Hinders AI Adoption in MSMEs of India?" Indian Growth and Development Review. https://doi.org/10.1108/IGDR-05-2024-0070.

Gupta, S. 2026. “India’s AI Regulation at the Crossroads: A Comprehensive Law or a Sectoral Mosaic?” Journal of Development Policy and Practice 11 (1): 120–32. https://doi.org/10.1177/24551333251382704.

Hakami, Ali. 2024. “Strategies for Overcoming Data Scarcity, Imbalance, and Feature Selection Challenges in Machine Learning Models for Predictive Maintenance.” Scientific Reports 14: 9645. https://doi.org/10.1038/s41598-024-59958-9.

Hakimi, Obaidullah, Hexu Liu, and Osama Abudayyeh. 2024. “Digital Twin-Enabled Smart Facility Management: A Bibliometric Review.” Frontiers of Engineering Management 11 (1): 32–49. https://doi.org/10.1007/s42524-023-0254-4.

India. Ministry of Communications. 2026. “12th State Broadband Committee Reviews Telecom Infrastructure and Digital Connectivity across Maharashtra.” Press Information Bureau, July 3, 2026. https://pib.gov.in/PressReleasePage.aspx?PRID=2280679.

India. Ministry of Electronics and Information Technology. 2023. The Digital Personal Data Protection Act, 2023. Act No. 22 of 2023. https://www.meity.gov.in/static/uploads/2024/06/2bf1f0e9f04e6fb4f8fef35e82c42aa5.pdf.

Ken Research. 2025. India PropTech Market Outlook. https://www.kenresearch.com/india-proptech-portals-and-broker-tech-market.

Lee, Chun-Chang, Lou-Jung Ting, Wen-Chih Yeh, and Zheng Yu. 2021. “The Influence of the Technical Dimension, Functional Dimension, and Tenant Satisfaction on Tenant Loyalty: An Analysis Based on the Theory of Planned Behavior.” International Journal of Strategic Property Management 25 (6): 469–84. https://doi.org/10.3846/ijspm.2021.15566.

Mahmud, Hasan, A. K. M. Najmul Islam, and Ranjan Kumar Mitra. 2023. “What Drives Managers towards Algorithm Aversion and How to Overcome It? Mitigating the Impact of Innovation Resistance through Technology Readiness.” Technological Forecasting and Social Change 193: 122641. https://doi.org/10.1016/j.techfore.2023.122641.

Mathagu, Samuel Wandeto. 2024. “Artificial Intelligence in Small and Medium Enterprises: An Empirical Analysis of Critical Factors.” Premier Journal of Science. https://premierscience.com/pjs-24-263/.

Ministry of Micro, Small and Medium Enterprises. 2025. “Notification S.O. 1364(E): Amendment to the Criteria for Classification of Micro, Small and Medium Enterprises.” https://egazette.gov.in/writereaddata/2025/261838.pdf.

Mordor Intelligence. 2026. India Facility Management Market: Industry Analysis, Size and Growth Outlook. https://www.mordorintelligence.com/industry-reports/india-facility-management-market.

NITI Aayog and Institute for Competitiveness. 2025. Enhancing Competitiveness of MSMEs in India. New Delhi: NITI Aayog. https://www.niti.gov.in/sites/default/files/2025-05/Enhancing_Competitiveness_of_MSMEs_in_India.pdf.

OECD. 2025. AI Adoption by Small and Medium-Sized Enterprises. Paris: OECD Publishing.

https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/12/ai-adoption-by small-and-medium-sized-enterprises_9c48eae6/426399c1-en.pdf.

Patil, Rakesh, and Diksha Swaroop. 2025. “Artificial Intelligence in SMEs in India: Implementation, Benefits, and Challenges.” International Journal for Multidisciplinary Research 7 (4). https://www.ijfmr.com/papers/2025/4/50209.pdf.

Rizwan, Rana, and Atif Hussain. 2024. “Strategic AI Adoption in SMEs: A Prescriptive Framework.” Preprint, submitted August 21, 2024. https://arxiv.org/abs/2408.11825.

Sampaio, Rodrigo Pedral, António Aguiar Costa, and Inês Flores-Colen. 2022. “A Systematic Review of Artificial Intelligence Applied to Facility Management in the Building Information Modeling Context and Future Research Directions.” Buildings 12 (11): 1939. https://doi.org/10.3390/buildings12111939.

Sánchez, Esther, Reyes Calderón, and Francisco Herrera. 2025. “Artificial Intelligence Adoption in SMEs: Survey Based on TOE–DOI Framework, Primary Methodology and Challenges.” Applied Sciences 15 (12): 6465. https://doi.org/10.3390/app15126465.

Soomro, Raheem Bux, Waleed Mugahed Al-Rahmi, Nisar Ahmed Dahri, Latifah Almuqren, Abeer S. Al-mogren, and Ayad Aldaijy. 2025. “A SEM–ANN Analysis to Examine Impact of Artificial Intelligence Technologies on Sustainable Performance of SMEs.” Scientific Reports 15: 5438. https://doi.org/10.1038/s41598-025-86464-3.

Tornatzky, Louis G., Mitchell Fleischer, and Alok K. Chakrabarti. 1990. The Processes of Technological Innovation. Lexington, MA: Lexington Books.

Ucar, Aysegul, Mehmet Karakose, and Necim Kırımça. 2024. “Artificial Intelligence for Predictive Maintenance Applications: Key Components, Trustworthiness, and Future Trends.” Applied Sciences 14 (2): 898. https://doi.org/10.3390/app14020898.

van Roosmale, Sophie, Peter Hellinckx, Jasmine Meysman, Stijn Verbeke, and Amaryllis Audenaert. 2024. “Building Automation and Control Systems for Office Buildings: Technical Insights for Effective Facility Management — A Literature Review.” Journal of Building Engineering 97: 110943. https://doi.org/10.1016/j.jobe.2024.110943.

Air Quality Perceptions and Public Transport Attitudes: A Literature Review of Delhi

By Ira Chinmay Dhamne

Abstract

Air pollution is a major environmental and public health issue in Delhi, with pollution levels frequently exceeding recommended limits and posing significant health risks. Public transportation plays an important role in reducing traffic emissions; however, its effectiveness is influenced by how commuters perceive air quality. Factors such as visible pollution, health concerns, travel conditions, and awareness of air pollution can shape public attitudes towards using public transport. The purpose of this review is to analyze how perceptions of air quality influence public attitudes towards the use of public transportation in Delhi. By synthesizing research on commuter perceptions, environmental psychology, and transport policy, this review examines how perceived air pollution affects travel behaviour and the acceptance of sustainable transportation. Understanding these perceptions can support the development of policies that encourage greater public transport use while improving urban air quality.

Keywords Delhi; public transport; perceptions; air quality.

Introduction

Air pollution is one of the leading environmental health risks worldwide, contributing to an estimated 7 million premature deaths annually (World Health Organization [WHO]). Rapid urbanization, industrialization, and increasing vehicle emissions have caused air quality to deteriorate in many cities, although the severity and sources of pollution vary depending on geography, climate, population density, and transportation systems (Sicard et al; Apte et al.). As cities continue to expand, understanding both the physical impacts of air pollution and how people perceive it has become increasingly important for developing effective environmental and public health policies (Lu et al.)

Delhi is one of the world's most polluted megacities where transport emissions, industrial activities, construction dust and seasonal crop residue burning contribute to persistently poor air quality (Sicard et al.). Despite major investments in sustainable transport infrastructure, particularly the Delhi Metro, private vehicle use remains widespread. Transportation is therefore both a significant contributor to air pollution and a key opportunity for reducing emissions through greater public transport use (Ray et al; Padma et al; Meena et al.). However, the relationship between air quality and transportation extends beyond measured pollution levels, as commuters also interpret and respond to the quality of the air they experience during travel. These interpretations or perceptions of air quality may influence how commuters evaluate different transportation options and their willingness to use public transportation.

Perception and attitude represent two distinct but related concepts in this context. Perception refers to how individuals interpret or understand the quality and potential effects of the air around them, whereas attitude refers to their evaluation or preference towards using

public transportation (Romarate et al.). Commuters may form perceptions of air quality based not only on actual pollution levels but also on visible smog, health concerns, travel experiences and awareness of environmental issues (Boso et al.). These perceptions may subsequently influence travel behavior and attitudes towards public transportation (Meena et al; Padma et al; Romarate et al.) Understanding this relationship is therefore important for determining whether concerns about air pollution translate into more favorable attitudes towards public transportation in cities such as Delhi.

This review paper examines how perceptions of air quality influence public attitudes towards the use of public transportation in Delhi, drawing on six research papers selected for detailed analysis from an initial screening of 56 papers. The literature includes studies on commuter perceptions, environmental psychology, transport behavior, urban air quality, and public health to identify the factors that influence perceptions of air quality and their relationship with transportation attitudes and behavior (Pinakidou; Rasmussen et al.) As air pollution is expected to remain a significant challenge in Delhi, understanding these relationships can help inform policies that improve both urban mobility and air quality.

Methodology

This study was conducted as a narrative literature review examining how perceptions of air quality influence public attitudes towards the use of public transportation in Delhi. Google Scholar was used to identify relevant academic literature concerning air pollution, public perceptions, commuter experiences, and transportation behaviour. The search strategy used combinations of keywords relating to air quality, perception, transportation, and Delhi. The initial search used the following search string: “Delhi” AND (“perceived air quality” OR “air pollution perception”) AND (“public transport” OR “Delhi Metro”) NOT (“India” OR Mumbai OR Kolkata OR Chennai OR Bengaluru). Additional search terms included “air quality perception,” “air pollution perception,” “perceived air pollution exposure,” “public transport,” “transportation behaviour,” “commuter perceptions,” and “mode choice.” Geographic terms, including “Delhi” and “New Delhi,” were combined with these terms to identify research relevant to the study area. Broader searches were also used to identify comparative research examining the relationship between air-quality information, pollution perceptions, and travel behaviour.

These records were assessed against the review's inclusion and exclusion criteria. Screening was conducted in stages, beginning with an assessment of titles and abstracts against the inclusion and exclusion criteria, followed by examination of the full texts of potentially eligible studies. Eligible studies were required to be peer-reviewed journal articles and to provide empirical evidence relevant to the relationship between air quality, its perception, and transportation attitudes or behaviour. Studies focusing on Delhi were prioritised. Particular attention was given to the distinction between perception and attitude: perception refers to how individuals interpret or understand the quality and potential effects of the air around them, whereas attitude refers to their evaluation or preference towards using public transportation.

Studies examining transportation attitudes without considering air-quality perceptions were excluded.

Following the initial screening, the full texts of potentially eligible papers were examined to evaluate their relevance to the research question and the suitability of their methods and findings. Papers that did not contribute sufficiently to understanding the relationship between air-quality perception and transportation were excluded. Six papers were ultimately selected for detailed analysis. Zhang et al. was retained as a comparative study because it experimentally examines how air-quality information influences pollution perception and travel willingness. However, because its empirical research was conducted outside Delhi, its findings are used to discuss a possible explanatory mechanism rather than as direct evidence of the behaviour of Delhi commuters.

The selection process is presented in Figure 1 using a PRISMA-informed flow diagram. The six selected papers were analysed thematically to identify recurring factors influencing air-quality perceptions, including environmental conditions, perceived health effects, travel experiences, information and awareness, and socioeconomic characteristics. The analysis then examined how these factors related to transportation preferences and behaviour, while distinguishing between reported attitudes, intended behaviour, and observed or reported travel choices.

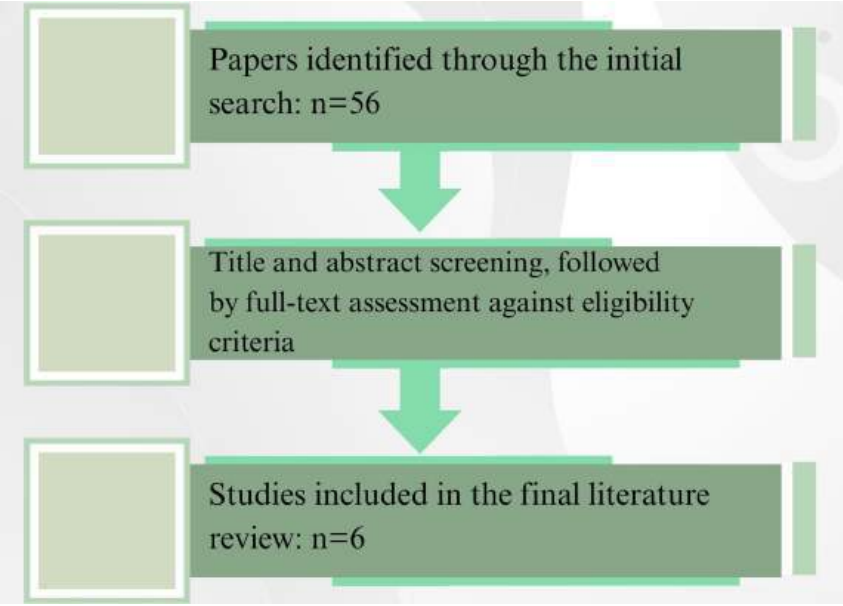

Figure 1. Literature Search and Study Selection Process
Source: Created by author using Canva.

Analysis

Table 1 summarises the research designs, samples or units of analysis, and geographical locations of the six studies included in this review. The studies use different methodological approaches, providing both direct evidence concerning air-quality perceptions and transportation attitudes and contextual or comparative evidence relevant to the research question.

Table1. Characteristics of the six studies included in the review

Study	RESEARCH DESIGN	SAMPLE/UNIT ANALYSIS	LOCATION
Bhatnagar, Jain and Sachdeva [2022]	Quantitative spatial analysis using land-use regression and a transit-oriented development[TOD] index	Neighbourhood-level built-environment indicators and air-quality data; no individual survey sample reported in the article summary	New Delhi, India
Meena, Singh, and Agarwal [2023]	Quantitative, face-to-face commuter survey	346 completed commuter responses	Delhi, India
Mondschein [2024]	Survet-based observational study	994 respondents from two neighbourhoods	Batla House and Bhalswa Dairy, Delhi, India
Padma, Shukla, and Dogra [2025]	Intercept questionnaire survey combined with experimental air-pollution measurements	389 questionnaires collected; 317 valid responses	Delhi Metro's Magenta Line and associated transport environments, Delhi, India
Venkatesh et al. [2018]	Community-based cross-sectional survey using simple random sampling	384 adults from seven localities	South Delhi, India
Zhang, Yu, and Lou [2025]	Multistage emperical study combining media analysis, behavioural experiments, field experiments and agent-based modelling	Multiple datasets and experimental samples across study stages; the study does not have one single sample size	Multiple locations, including cities in China, not Delhi

Air Quality and Public Transportation in Delhi

Delhi's severe air pollution is closely connected to its transportation system, making public transport an important part of the discussion surrounding exposure and air-quality improvement. The literature indicates that commuters are exposed to substantially different levels of pollution depending on their mode of travel. (Padma et al.), for example, compared commuters' perceptions of pollution exposure across different modes with measured pollutant concentrations. The measured $PM_{2.5}$ concentration was approximately 34.2 μg/m³ in the Metro, compared with 113.2 μg/m³ in auto-rickshaws and 68.9 μg/m³ in open-window buses. Commuters also generally perceived the Metro as having lower pollution exposure than other modes.

The relationship between transport and air pollution is therefore not simply about the amount of pollution produced by vehicles; it also concerns how commuters understand their exposure. (Meena et al.) found that many Delhi commuters had limited understanding of the Air Quality Index (AQI), with only around 27% reporting awareness of it. Nevertheless, 68% reported avoiding travel when pollution reached very poor or severe levels. This suggests that commuters can respond to perceived pollution even when their understanding of formal

air-quality measurements is limited.

The importance of public transportation is further demonstrated by research on travel choices during poor air-quality conditions. (Meena et al.) found that as air quality deteriorated, commuters tended to reduce their use of more exposed modes and increase their preference for modes such as the Metro. However, these changes were not uniform across socioeconomic groups. Higher-income commuters had greater opportunities to avoid open modes of transport, while lower-income commuters were more likely to continue using modes that exposed them to pollution, demonstrating that transport choice is also constrained by economic circumstances.

Public Perceptions of Air Quality

The reviewed literature indicates that perceptions of air quality are shaped by more than objective measurements. In Delhi, people commonly assess pollution through their immediate experiences, awareness of pollution, perceived health effects, and characteristics of their surrounding environment (Venkatesh et al; Meena et al; Padma et al; et al.; Mondschein). (Venkatesh et al. 22) found that 57.6% of participants believed that air quality had become worse compared with five years earlier, while 95% (Venkatesh et al. 23) reported that pollution had affected their lives in some way. Breathlessness and irritation of the eyes, nose, or throat were among the commonly perceived effects.

These findings demonstrate an important distinction between perception and attitude. Perception refers to an individual's interpretation of the quality and healthiness of the surrounding air, whereas attitude refers to an individual's evaluation or preference regarding public transportation. A person may perceive Delhi's air as highly polluted without necessarily changing their preferred mode of transport, particularly when they have limited alternatives. This distinction becomes particularly important when interpreting the relationship between air-quality perceptions and transportation behaviors.

Factors influencing perception

Several factors appear to influence how Delhi residents perceive air pollution.

Visible pollution can make poor air quality more noticeable. Smog and haze provide direct sensory cues that allow individuals to recognise pollution without accessing formal measurements which is particularly relevant in Delhi, where seasonal pollution episodes can make air pollution highly visible (Venkatesh et al; Meena et al; Mondschein).

Health concerns are another important factor. (Venkatesh et al.) found that many participants perceived physical symptoms associated with air pollution, while their survey also showed substantial concern about the health consequences of pollution. (Padma et al.) provides further evidence of why these concerns may be significant, as the study examined commuters' perceived exposure alongside measured $PM_{2.5}$ concentrations across different modes of transport. (Meena et al.) similarly found that health concerns were an important part of commuters' responses to poor air quality and that many respondents were willing to reconsider their travel choices when air quality deteriorated (Venkatesh et al; Padma et al; Meena et al.).

Travel conditions also influence perceptions. (Padma et al.) found that characteristics such as waiting time, trip purpose, education, occupation, distance and mode of travel were associated with differences in commuters' perceptions of pollution exposure. (Meena et al.) similarly demonstrated that commuters' travel preferences were affected by deteriorating air-quality conditions, with many respondents reporting a willingness to change their travel choices when pollution became poor or worse. (Mondschein) provides an important additional perspective, showing that environmental perceptions do not always translate directly into behavioural changes, as residents sometimes continued walking despite perceiving their surrounding environmental conditions negatively. Thus, perception is influenced not only by the surrounding air but also by the conditions under which people experience that air (Padma et al; Meena et al; et al.; Mondschein).

Information and awareness are particularly important. Meena et al. found that awareness of the AQI was relatively limited, meaning that many commuters may not interpret air pollution through formal AQI measurements. Venkatesh et al. similarly demonstrated that people's understanding of air pollution was shaped by their experiences and awareness of its effects on their daily lives. Zhang et al. provides further evidence that information can influence perception: their experiment demonstrated that presenting people with misleadingly positive air-quality information reduced their perceived level of pollution and increased their willingness to travel. Together, these findings suggest that the information available to individuals can influence not only how seriously they perceive air pollution but also how they respond to it (Meena et al; Venkatesh et al; Zhang et al.).

Finally, personal experience and socioeconomic characteristics can shape perceptions. Venkatesh et al. identified differences in perceptions according to gender, education and income, while (Padma et al.) similarly found socioeconomic and travel characteristics to be associated with perceptions of pollution exposure.

Perceived and Measured Air Quality

An important finding across the literature is that perceived air quality does not always correspond perfectly with measured pollution.

Padma et al. provide evidence of relatively strong correspondence between perception and objective exposure. Their measurements showed considerably lower $PM_{2.5}$ concentrations inside the Metro than in several road-based modes, and commuters similarly ranked the Metro as the mode with lower pollution exposure. In this case, commuters' perceptions were broadly consistent with the physical environment.

However, other findings demonstrate that perception can differ from objective conditions. Meena et al. found that many commuters lacked sufficient knowledge of the AQI even though they were able to identify periods of severe pollution and reported changing their travel behaviour during these periods. Zhang et al. provides an even clearer demonstration: when participants were given misleading information suggesting that air quality was better than it actually was, their perceived pollution decreased and their willingness to travel increased.

Mondschein further demonstrates that negative perceptions do not necessarily result in behavioural change. The study identified a "walking dissonance" in which some Delhi residents continued to walk frequently despite perceiving their environmental conditions as poor. Higher walking dissonance was associated with poorer perceived air quality and lower subjective well-being. Together, these studies suggest that perceived air quality can correspond with measured pollution, but it can also be influenced by information, experience and practical constraints. Therefore, perception should not be treated as a direct substitute for objective measurements.

Impact of Air-Quality Perceptions on Public Transport Attitudes

Travel Behaviour

The evidence suggests that perceptions of air quality can influence transportation behaviour, although the effect is not always straightforward.

Meena et al. found that 68% of commuters reported avoiding travel when air pollution was very poor or severe. The study also found that when air quality deteriorated, commuters changed their preferences between transport modes. In particular, use of motorised two-wheelers decreased while Metro use increased. This indicates that when pollution is perceived as sufficiently severe, commuters may prefer modes that they believe provide greater protection from exposure.

Padma et al. supports this finding from a different perspective. Commuters generally perceived the Metro as less polluting than road-based modes, and the measured exposure levels supported this perception. This suggests that the perception of public transport as a cleaner mode can contribute positively to its attractiveness.

However, the relationship between perception and behaviour is constrained by socioeconomic circumstances. Meena et al. found that lower-income commuters were more likely to remain dependent on walking or open modes during poor air quality, whereas higher-income commuters had greater flexibility to change their transport choices. et al. similarly found that some residents continued walking despite poor perceived environmental conditions.

Therefore, negative perceptions of air quality may encourage people to seek cleaner modes, but perception alone does not determine behaviour. Accessibility, income, trip purpose and available alternatives can either strengthen or weaken its effect.

Psychological and Social Factors

Perceived health risk provides an important psychological pathway between air quality and transport attitudes. When individuals believe that pollution poses a serious threat to their health, they may become more concerned about exposure and more willing to alter their travel behaviour.

Venkatesh et al. found substantial public awareness of the effects of air pollution, with

many respondents reporting perceived symptoms and lifestyle changes related to pollution. Almost half reported restricting outdoor activities because of air pollution. Bhatnagar et al. further demonstrates that exposure differs according to transport mode, providing a physical basis for commuters' concerns about exposure during travel. Mondschein similarly found that some residents continued walking despite poor perceived environmental conditions.

Environmental awareness also appears to be associated with support for pollution-control measures. Venkatesh et al. found generally high support for measures intended to control air pollution. However, this should not automatically be interpreted as support for public transportation specifically as the study measured broader attitudes towards pollution control, rather than directly establishing that perceptions of air quality caused a positive attitude towards public transport.

Information is another psychological factor. Zhang et al. demonstrates experimentally that changing the information people receive can alter both their perception of pollution and their willingness to travel. This suggests that public understanding of air quality can influence behavioural intentions even when the physical environment itself has not changed.

Policy Acceptance

The literature provides some evidence that public perceptions are relevant to transportation policy, but this area is less directly established than travel behaviour.

Padma et al. argues that commuters' perceptions should be considered in policy decisions and supports the promotion of segregated public transportation from a health perspective. Venkatesh et al. also found relatively strong support for air-pollution control measures and recommended public transport as one component of broader pollution-control efforts.

However, the six paper evidence base does not provide strong enough evidence to conclude that perceptions of air quality directly determine support for specific policies such as congestion charges. Therefore, rather than claiming that commuters would necessarily support congestion pricing or vehicle restrictions, the literature supports a more cautious conclusion: greater awareness of pollution and perceived health risks can create conditions that make cleaner transport interventions more acceptable, but specific policy preferences require further research.

Discussion

The findings from the reviewed literature indicate that perceptions of air quality have a meaningful relationship with attitudes and behaviour surrounding public transportation in Delhi. However, this relationship is conditional rather than deterministic and three major patterns emerged:

First, commuters generally recognise differences in pollution exposure between transport modes. Padma et al. found that the Metro was perceived as less polluting than several road-based modes, and objective measurements largely supported this perception. This is important because it suggests that public transportation can become more attractive when commuters associate it with lower pollution exposure.

Second, perceptions of air pollution can influence travel behaviour independently of complete knowledge of measured pollution. Meena et al. found limited AQI awareness alongside substantial reported avoidance of travel during severe pollution. This suggests that people do not need precise technical knowledge of pollution levels to respond to their perception that conditions are unsafe.

Third, perceptions do not always translate into behavioral change. Mondschein demonstrates that people may continue walking despite poor perceived environmental conditions, while (Meena et al.) shows that income affects the ability to change modes. This is particularly important for Delhi because it shows that improving perceptions alone cannot guarantee greater public transport use. Commuters need accessible and practical alternatives.

The literature therefore suggests that perception acts as one factor within a broader decision-making process. Health concerns, information, socioeconomic circumstances, trip purpose, travel time and mode availability interact with perceptions to shape transportation decisions. Padma et al., for example- identified several socioeconomic and travel related variables associated with perceptions of pollution exposure.

The distinction between perceived and measured air quality is also central. In some circumstances, perceptions reflect objective differences, as demonstrated by the Metro exposure results. In others, perceptions can be changed by information or can fail to change behaviour because of economic necessity. Zhang et al. demonstrates the first situation experimentally, while Mondschein illustrates the second in Delhi.

For Delhi, this suggests that transportation policy should address both actual exposure and perceived exposure. Improving the environmental performance of public transport is important, but communicating those improvements clearly may also influence whether commuters recognise public transportation as a safer alternative. At the same time, policies must address affordability, accessibility and connectivity so that people who are concerned about pollution actually have viable alternatives.

Research Gaps and Recommendations

Despite the evidence provided by the six papers, several gaps remain.

First, relatively few studies directly examine the complete pathway from air-quality perception to public transport attitude to actual transport behaviour. Many studies measure perception and behaviour separately, making it difficult to establish whether changes in perception directly cause changes in transport choice.

Second, more longitudinal research is needed. Much of the existing evidence captures perceptions or behaviour at a particular point in time. Research conducted across different seasons could determine whether commuters' perceptions and attitudes change as Delhi moves between relatively cleaner and more polluted periods.

Third, socioeconomic differences require further investigation. Meena et al. and Mondschein indicate that income and access to alternatives can influence responses to poor air quality. Future research should examine how income, gender, education, occupation and access

to Metro stations influence the relationship between perceptions and transport choices.

Fourth, future research should investigate the role of air-quality information more directly. Zhang et al. demonstrates that information can alter perceptions and willingness to travel, but further Delhi-specific research could examine whether accurate, accessible real-time AQI information encourages commuters to choose lower-exposure public transport modes.

Finally, research should distinguish more clearly between attitudes and intentions and actual behaviour. A commuter may express a positive attitude towards public transport without actually using it because of cost, travel time, convenience or limited accessibility. Future studies should therefore combine questionnaires with observed travel behaviour.

Conclusion

This review examined the research question: What impact do perceptions of air quality have on public attitudes towards the use of public transportation in Delhi?

The evidence indicates that perceptions of air quality have a meaningful influence on transportation attitudes and behaviour, particularly when commuters perceive pollution as a serious health risk. Poor perceived air quality can encourage people to avoid travel, reduce their use of more exposed modes, or prefer modes they perceive as offering greater protection from pollution. In Delhi, the Metro is frequently perceived as a lower-exposure option, and objective measurements support this perception (Padma et al.).

However, the relationship is not straightforward. Perceptions do not always correspond with measured air quality, and negative perceptions do not always result in behavioural change. Limited AQI awareness, socioeconomic constraints, trip requirements and availability of transport alternatives can influence whether perceptions translate into action. Meena et al. and Mondschein particularly demonstrate that commuters may continue using less desirable or more exposed modes when alternatives are limited.

Therefore, perceptions should be understood as an important influence on public transportation attitudes rather than the sole determinant of transport behaviour. For Delhi, increasing public transport use requires both improvements in actual air-quality exposure and greater public awareness of these benefits. Accurate air-quality information, health-focused communication and accessible low-exposure transport options could help commuters make more informed choices.

Ultimately, the literature suggests that improving Delhi's air quality and increasing public transport use requires attention to both the physical environment and how that environment is perceived by the people who travel through it.

Works Cited

Apte, Joshua S., and Chirag Manchanda. “High-resolution urban air pollution mapping.” Science 385.6707 (2024): 380-385.

Bhatnagar, Smriti, Deepty Jain, and Kamna Sachdeva. "Effect of transit-oriented development on air quality in neighbourhoods of Delhi." World Development Sustainability 1 (2022): 100015. https://doi.org/10.1016/j.wds.2022.100015.

Boso, Àlex, et al. "Why do we misperceive air pollution? A scoping review of key judgmental biases." Air Quality, Atmosphere & Health 18.2 (2025): 447-460.

Lu, Jackson G. "Air pollution: A systematic review of its psychological, economic, and social effects." Current Opinion in Psychology 32 (2020): 52-65.

Meena, Kapil Kumar, Vikram Singh, and Amit Agarwal. "Perception of commuters towards air quality in Delhi." Journal of Transport & Health 31 (2023): 101643 https://doi.org/10.1016/j.jth.2023.101643.

Mondschein, Andrew. “Walking and walkability in Delhi: dissonance between environmental perception and behavior.” Transportation Research Record, vol. 2678, no. 7, 2024, pp.878-890 https://doi.org/10.1177/03611981231210541.

Padma, S., Anuradha Shukla, and Nitish Dogra. "An evaluation of commuter’s perception of exposure to air pollution on a segment of Delhi Metro’s Magenta line." Discover Atmosphere, vol.3, 2025, article 43. https://doi.org/10.1007/s44292-025-00071-5

Pinakidou, Sophia. “People’s Perceptions of Air Pollution and Their Awareness of Official Indexes at the Start of the Twenty-First Century: A review.” Discover Environment, vol.3, 2025, article 57. https://doi.org/10.1007/s44274-025-00213-x

Ray, Shomik, Ranjana Singh, and Sukanya Pathak. "Are citizens of Delhi willing to stop using private vehicles for their daily office commute in order to reduce air pollution and related ailments? A protocol for a contingent valuation study." International Journal of Community Medicine and Public Health 9.12 (2022): 4765-4773.

Romarate II, Rodolfo, et al. "When air quality matters: Awareness, perception, and attitude of the residents in Metro Manila, Philippines." Environment, Development and Sustainability 28.2 (2024): 5109-5126.

Sicard, Pierre, et al. "Trends in urban air pollution over the last two decades: A global perspective." Science of the Total Environment 858 (2023): 160064.

Venkatesh, U., et al. "Public perception survey on air pollution in South Delhi." International Journal of Preventive, Curative & Community Medicine, vol. 4, no. 2, 2018, pp. 20-27 https://doi.org/10.24321/2454.325X.201812

Zhang, Yue, Xiaoli Yu, and Siyang Luo. "Misleading air quality reports lower the public’s perception of pollution and increase travel behavior." Communications Earth & Environment 6.1 (2025): 689.

Which Basic Emotion Increases Brand Recall More: Happiness, Surprise or Fear? By Zeynep Nur Ayan

Abstract

This study examines how discrete emotional appeals may differentially influence brand recall in advertising. While existing research frequently treats emotion as a broad construct, this study focuses specifically on three emotional categories: happiness, fear, and surprise as they are among the most common emotions used in marketing. Using a theory-driven content analysis grounded in the valence-arousal framework, six video advertisements were analyzed to identify differences in emotional tone, narrative structure, and brand positioning. Emotional appeals were classified based on established theoretical definitions and evaluated in relation to their potential influence on brand recall. The findings suggest that although happiness, fear, and surprise are all associated with heightened arousal, they differ meaningfully in valence and structural presentation. Happiness-based advertisements tended to integrate brands within positive and engaging narratives, while fear-based advertisements positioned brands as solutions to perceived risks. Surprise-based advertisements captured attention through unexpected events, often linking brand presentation to moments of heightened cognitive engagement. These distinctions indicate that emotional appeals should not be considered interchangeable strategies. By examining discrete emotions rather than generalized emotional appeals, this study contributes to a more nuanced understanding of emotional differentiation in advertising research. The results highlight the importance of emotional specificity and suggest that the strategic selection of emotional appeals may shape how brands are processed and remembered.

Keywords: Emotional Appeals; Brand Recall; Advertising Effectiveness; Valence-Arousal Framework; Content Analysis

Introduction

Advertising plays a critical role in how consumers perceive and remember brands. One of the most widely used ways to this idea is implemented is through brand recall, which refers to the customer's ability to remember the brand (Missaglia et al., 2017). Brand recall is important because when researchers want to examine whether an ad was successful, they can measure if consumers can remember the brand and how they remember it. This can be assessed through different approaches, for example, researchers can ask consumers if they recognize an ad when they see it again or what would change the way they perceive the brand (Mehta, 2006; Missaglia et al., 2017). Measuring brand recall examines whether an advertisement leaves a lasting impression because it reflects memory rather than just immediate reaction. This is important because the advertisements that have a high level of brand recall are more likely to influence customer decisions, making recall a key outcome in advertising research.

Emotional advertising has been shown to increase brand recall substantially (Missaglia et al., 2017). This is because consumers do not remember the brands based on information alone, but rather they remember the brand because of the emotion elicited by the ad which is associated with the memory of the brand - for example, television advertisements

with catchy songs are more memorable than those without them. Emotional reactions can capture attention and increase the involvement with the advertisement and branding which supports memory development. As a result, many advertising studies focus on understanding the emotional feedback from advertisements and how they influence the customer responses including attention, attitude, recognition and recall (Hussain, 2025).

The present study focuses on emotional branding as it is associated with three basic emotions: happiness, surprise and fear. Happiness, surprise, and fear were chosen because they are emotions that are encountered most frequently in daily life. For that reason, happiness is typically considered to be a safe emotion in emotional advertising because it is associated with positive feelings and memories, in turn making recall more effective and efficient (Mehta, 2006). Other emotions, like fear, differ from happiness because fear is linked to more negative feelings, memories, and situations which in turn makes the consumer more likely to forget the advertisement (Missaglia et al., 2017). Apart from that, when fear based emotional marketing is used carefully, and because it is a high-arousal emotion, it can strongly capture attention if used correctly. Finally, surprise is distinct in that it is an attention-driven emotion, meaning that it grabs the attention of the consumer even before it is fully processed and evaluated by the individual (Mehta, 2006). This element of surprise can lead to positive and/or negative feelings and memories through the disruption of expectations and increase the memorability. In sum, each of these three emotions differ in valence and arousal —direction and intensity— which means that they create a different influence on brand recall rather than having uniform effects.

Although prior research establishes that emotion in advertising increases the effectiveness in brand recall, existing literature also leaves an important question unanswered: which specific emotions affect brand recall most? Studies such as Aydın (2014) show the mechanism behind brand recall, how emotional responses help with the memory retrieval. While Panda (2013) emphasizes that different emotional arousal responses show different risks and effectiveness. However, much of the existing literature establishes that emotion matters for recall but treats emotion as a broad construct, rather than examining special emotional experiences. As a result, the gap in literature appears here concerning the comparative effects of specific emotions on brand recall. Rather than directly comparing individual emotions, prior studies often focus on general emotional appeals in branding. The current study directly addresses this limitation by examining how happiness, fear, and surprise differ in their influence on brand recall, providing a more detailed understanding of emotion-specific effects in advertising.

Methods

This study used a theory-driven content analysis to examine how different emotional appeals influence brand recall in advertising. Instead of treating emotion as a single, general construct, the analysis focused on three specific emotional appeals: happiness, fear, and surprise. This approach allowed for a more detailed comparison of how each emotion may differently affect advertising outcomes and was guided by established emotion theory.

Six video advertisements were selected for analysis, with two advertisements representing each emotional appeal category. The advertisements were chosen based on the clarity and dominance of their emotional framing, the visibility of the brand within the

advertisement, and their relevance to widely recognized advertising campaigns. To ensure consistency across stimuli, one representative still frame capturing the primary emotional moment was extracted from each advertisement (see Appendix A).

Emotional appeals were identified using the valence-arousal framework, which is commonly applied in emotion research. Within this framework, valence refers to whether an emotional experience is positive or negative, while arousal refers to the level of emotional activation. For each advertisement, emotional characteristics were assessed by examining visual and contextual cues, including facial expressions when present, situational context, narrative tone, and the overall intensity of the emotional scene. Advertisements characterized by positive valence and high arousal were classified as happiness based, those characterized by negative valence and high arousal were classified as fear based, and those associated with high arousal resulting from unexpected events or sudden attentional shifts were classified as surprise based. Final classifications were determined by ensuring consistency between the observed cues and established theoretical definitions of each emotional category.

Following the classification process, the advertisements were compared in terms of their potential influence on brand recall. By focusing on specific emotional appeals rather than generalized emotional content, the analysis aimed to address gaps in existing research concerning how discrete emotions may differ in their effects on advertising effectiveness.

Results

The content analysis revealed clear differences in emotional characteristics across the three advertising categories: happiness, fear, and surprise. Each category demonstrated distinct patterns in emotional valence, arousal, and brand presentation, suggesting that discrete emotions may influence brand recall in different ways.

Advertisements categorized as happiness-based were characterized by positive valence and high arousal. These advertisements frequently presented emotionally uplifting narratives, positive interpersonal interactions, and visually warm or engaging scenes. Brand elements in happiness-based advertisements were typically integrated smoothly into the narrative, often appearing alongside positive emotional moments. This alignment between positive emotional tone and brand visibility suggests a potential facilitative effect on brand recall, as the brand is associated with enjoyable and emotionally engaging content.

In contrast, fear-based advertisements exhibited negative valence combined with high arousal. These advertisements emphasized themes of risk, vulnerability, or potential threat, often creating a sense of urgency or concern. Brand presentation in fear-based advertisements was commonly positioned as a solution or source of reassurance in response to the depicted threat. As a result, brand recall in this category appeared closely tied to the resolution of the emotional tension, with the brand serving a functional and protective role within the narrative.

Surprise-based advertisements were primarily defined by high arousal resulting from unexpected events or sudden shifts in attention. Unlike happiness and fear, surprise was not consistently associated with a single valence but was instead characterized by abrupt attentional disruption. In these advertisements, brand presentation often followed the surprising moment or was closely linked to the unexpected action. This structure suggests

that surprise may enhance brand recall by capturing attention first and allowing the brand to benefit from heightened cognitive engagement immediately afterward.

Overall, the results of the content analysis indicate that while all three emotional appeals are associated with high arousal, they differ meaningfully in emotional tone and narrative structure. These differences suggest that happiness, fear, and surprise may influence brand recall through distinct psychological mechanisms rather than producing uniform effects.

Discussion

The findings from this study indicate that emotional appeals differ in their influence on brand recall. The results from the study challenge the assumption that all emotional appeals have the same effect and should be treated as a single category. Emotional appeals should not be treated as interchangeable strategies, as each emotion carries unique psychological functions that shape how advertising messages are interpreted and remembered. These are the differences in valence and arousal play an important role shaping how emotions influence brand recall.

Advertisements, often associated with the emotion of happiness, were found to have higher and stronger levels of brand recall in previous research (Aydın, 2014). This effect is explained by the positive valence of happiness, which improves memory encoding abilities, and may encourage more favorable feelings about the brand itself. Such positive emotional states are known to increase the engagement with advertisements, in turn enhancing recall. This increased brand recall may explain why happiness-based appeals are commonly used in brand-building campaigns. These findings align with existing literature highlighting the role of positive emotions in advertising effectiveness. Taken together, happiness contributes to brand recall by positive emotional engagement (Missaglia et al. 2017). Happiness-based emotional appeals may be particularly effective for brands that aim to create positive associations and long-term emotional connections with consumers, as positive valence emotions are associated with favorable evaluations and approach-oriented processing (Aydın, 2014; Missaglia et al., 2017). Tourism and travel brands, for example, frequently emphasize joy, relaxation, and enjoyment to frame their services as sources of pleasure and escape. Similarly, food and beverage brands often rely on happiness-centered messaging to associate their products with shared experiences, celebration, and social connection (Panda, 2013). Lifestyle and fashion brands may also use positive emotional framing to reinforce aspirational identities and feelings of confidence or satisfaction. In these contexts, the use of happiness aligns with the brand's objective of generating positive affect and strengthening favorable consumer perceptions (Mehta, 2006).

Among the emotional appeals examined, fear demonstrated a unique pattern in relation to brand recall. High-arousal emotions such as fear are known to increase attentional focus on the advertisement which may support memory formation. However, fear can also introduce feelings such as overwhelm, which may limit the effectiveness if not carefully managed. Unlike positive emotions, fear involves both a high attention and a potential resistance that highlights the complexity of its strategy (Missaglia et al., 2017). These findings suggest that fear must be carefully balanced to avoid negative brand associations. In this context, fear represents a powerful marketing tool while also posing potential challenges.

Fear-based emotional appeals may be particularly effective in contexts where risk, vulnerability, or potential loss is central to the message. Fear is associated with threat perception and protective motivation, which can encourage heightened attention and problem-focused processing (Mehta, 2006; Hussain, 2025). For this reason, safety-oriented brands such as insurance providers, home security companies, and health awareness campaigns often rely on fear-based strategies to highlight potential dangers and position the brand as a protective solution. In these contexts, fear serves a functional role by increasing perceived risk and motivating action. However, the same emotional strategy may be inappropriate for leisure, lifestyle, or enjoyment-oriented products, where excessive threat framing could create negative brand associations (Aydin, 2014; Missaglia et al., 2017). This further reinforces the idea that emotional appeals must align with the intended message and consumer expectations.

In regard to surprise-based advertising, the results suggest that surprise-based advertisements influence brand recall in a unique way. Unexpected elements in advertisements can increase the level of attention, by interrupting routine information processing (Metha, 2006; Missaglia et al., 2017). This suggests that surprise may be most effective when the unexpected element is clearly linked to the brand (Metha, 2006). Because surprise is not inherently positive or negative, its impact on brand recall may vary (Missaglia et al., 2017). In contrast to happiness or fear, surprise operates through mechanisms different from other emotional appeals. Overall, surprise appears to enhance brand recall by capturing attention and being distinctive when used strategically. Surprise-based emotional appeals may be particularly effective in contexts where capturing attention and creating memorability are central objectives. Surprise is characterized by heightened arousal and an interruption of expectations, which can increase cognitive engagement and information processing (Missaglia et al., 2017; Aydın, 2014). For this reason, brands in highly competitive markets, such as technology, entertainment, or fast-moving consumer goods, may use unexpected twists or unconventional storytelling to differentiate themselves and enhance recall. In digital and social media environments, where attention spans are limited, surprise-based strategies may be especially useful for interrupting scrolling behavior and drawing immediate focus. However, surprise must remain coherent with brand identity, as excessive unpredictability may reduce message clarity or weaken brand associations (Panda, 2013). These considerations further emphasize that the strategic use of emotion depends on contextual alignment and communicative purpose.

It is important to note that these findings should not be interpreted as indicating that one emotional strategy guarantees higher brand recall in all contexts (Aydın, 2014; Missaglia et al., 2017). For example, animal rescue and humanitarian aid advertisements in many countries often rely on sadness-based emotional appeals to evoke empathy and encourage engagement or donations. In these contexts, happiness-based appeals would likely weaken the intended emotional response (Panda, 2013). In contrast, tourism and vacation advertisements typically emphasize happiness, relaxation, and enjoyment, making fear-based appeals inappropriate and ineffective. Similarly, safety-related advertising, such as insurance or home security campaigns, frequently employs fear to highlight risk and vulnerability, whereas the same emotional strategy would be unsuitable for promoting leisure or lifestyle products. These examples illustrate that the effectiveness of emotional appeals depends

strongly on the contextual alignment between emotion, message purpose, and audience expectations. Therefore, the findings should be interpreted with caution rather than as definitive prescriptions for advertising strategy (Metha, 2006; Hussain, 2025). As a result, the findings reflect comparative tendencies rather than universal outcomes, emphasizing the value of emotion specific analysis in advertising research.

Implications

The findings of this study suggest that different emotional appeals may function in distinct ways within advertising rather than producing uniform effects. Although happiness, fear, and surprise are all associated with heightened emotional engagement, they differ in emotional tone and in how brands are positioned within advertisements. This indicates that emotional appeals should be considered in relation to their specific characteristics rather than treated as interchangeable strategies.

From a research perspective, these findings highlight the importance of examining emotions at a greater detail, as individual emotions can influence consumer responses in different ways. Much of the existing literature discusses emotional appeals broadly, which can obscure differences between individual emotions. By focusing on happiness, fear, and surprise separately, this study contributes to a more detailed understanding of how discrete emotions may relate to brand recall. The results also suggest that the valence-arousal framework can be a useful tool for organizing and comparing emotional appeals in advertising content.

From a practical perspective, the findings imply that advertisers may benefit from selecting emotional appeals based on their communication goals. Happiness-based appeals may be more appropriate when the goal is to create positive brand associations, while fear-based appeals may be more relevant in contexts where awareness of risk or problem-solving is emphasized. Surprise-based appeals may be particularly useful for attracting attention and increasing memorability. These implications should be interpreted cautiously, as the analysis focuses on advertising content rather than direct consumer responses.

Limitations

Several limitations of this study should be acknowledged. First, the analysis was based on a small, purposive sample of advertisements, which limits the generalizability of the findings. Although the selected advertisements were well-known and clearly represented their respective emotional appeals, they may not capture the full range of ways in which happiness, fear, and surprise are used in advertising across industries or cultural contexts. Second, the study relied on a theory-driven content analysis rather than direct measures of audience responses. As a result, conclusions about brand recall are inferred from the structure and emotional characteristics of the advertisements rather than from empirical recall data collected from consumers. Future studies incorporating participant-based measures could provide stronger evidence regarding the relationship between emotional appeals and actual brand recall outcomes. Third, emotional classification was based on the interpretation of visual and contextual cues using established theoretical frameworks. While this approach allows for transparent and replicable analysis, it may involve a degree of subjectivity,

particularly for emotions such as surprise, which can vary in valence depending on context. Although theoretical definitions were applied consistently, alternative interpretations of emotional content are possible.

Future Directions

Future research could expand this research by analyzing emotional appeals using larger and more varied selection of advertisements, providing a more comprehensive understanding of emotional strategies. Additional research could benefit from examining emotional appeals on different industries, cultural contexts and advertising formats which may help with the clarification of contextual differences with a deeper insight. For instance, comparing Western advertising contexts with East Asian markets may provide improved understanding into how cultural norms surrounding emotional expression influence the effectiveness of discrete emotional appeals.

. Additional future research should consider incorporating experimental studies involving consumers. Direct consumer reports would allow researchers to get a better understanding of the outcome and strengthen the empirical evidence. Survey-based or laboratory studies could garnish data-driven support for the findings of this content analysis. Another critical future direction includes investigating a broader spectrum of emotions beyond happiness, fear and surprise including sadness, anger or nostalgia. Analysis of additional emotions could provide a broader understanding by revealing patterns that are not captured in the current analysis limited by 3 emotions. The increased number of emotional appeals assessed increases, a more nuanced understanding could offer insight into how different emotional strategies influence consumer memory and perception. As a continuation of this study, I intend to further examine how discrete emotional appeals function across different cultural contexts, as emotional expression and audience expectations may vary internationally between different markets. A cross-cultural comparison may reveal whether different emotions function consistently or adapt to local norms and expectations. In addition, future studies may examine how emotional differentiation functions within digital and social media advertising formats, where attention dynamics and audience engagement differ from traditional video advertisements. Expanding analysis to these contexts may offer a more comprehensive understanding of emotional strategy in contemporary advertising.

Conclusion

In conclusion, this study set out to examine whether specific emotional appeals, namely happiness, fear, and surprise, may function differently in shaping brand recall within advertising contexts. Through a theory-driven content analysis grounded in the valence-arousal framework, the findings suggest that emotional appeals should not be understood as a single, uniform strategy. Although all three emotions involve heightened arousal, they differ in emotional tone, narrative structure, and the way brands are positioned within the advertisement, which may influence how audiences process and remember the content. By focusing on discrete emotions rather than general emotional appeals, this study attempts to contribute to a more detailed and differentiated understanding of emotional strategy in advertising research. The findings underscore the importance of examining emotional specificity and encourage a more careful consideration of how individual emotions

operate within communication contexts. Ultimately, this research represents an initial step toward a more nuanced exploration of emotional differentiation in advertising and highlights the value of continued inquiry into how distinct emotional experiences shape consumer perception and memory.

Appendix A: Advertisements Used in the Study

A1. Coca-Cola – Who will you Share A Coke with? (Happiness)

This advertisement portrays individuals sharing Coca-Cola products in positive social settings. The campaign emphasizes personal connection and enjoyment through the use of named bottles, smiling faces, and friendly interactions. The overall tone is warm and uplifting, with the Coca-Cola brand prominently displayed throughout the advertisement. Therefore, this ad was categorized under happiness due to its clear use of positive emotional cues.

A2. Google – 2022 — Year in Search (Happiness)

This advertisement presents a retrospective overview of global search trends from 2022, highlighting moments of resilience, achievement, and positive human connection. The narrative emphasizes hopeful and uplifting events, supported by emotional imagery and music. The Google brand is clearly presented at the conclusion of the video. The advertisement was categorized under happiness due to its overall positive emotional tone and emphasis on collective optimism.

A3. SimpliSafe – SimpliSafe: "Fear Is Everywhere" - 2019 Super Bowl Commercial (Fear)

This advertisement portrays a variety of threatening situations to emphasize how fear can be present in everyday life. The narrative highlights vulnerability and risk within domestic settings, creating a strong sense of threat and unease. Fear is elicited through the depiction of potential danger and loss of safety. The advertisement was categorized under fear due to its clear reliance on high-arousal, negative emotional cues, with the brand positioned as a protective solution.

A4. BMW – Autonomous Driving- Nothing To Fear (Fear)

This advertisement addresses fear related to autonomous driving technology by highlighting concerns about safety and loss of control. The narrative acknowledges technological anxiety before positioning the brand as a reliable and trustworthy solution. Fear is elicited through uncertainty and perceived risk rather than physical danger. The advertisement was categorized under fear due to its emphasis on high-arousal negative anticipation and safety concerns.

A5. Google – Parisian Love (Surprise)

This advertisement presents a sequence of Google search queries that gradually reveal a personal life narrative. The surprise emerges as the viewer recognizes the underlying story formed through seemingly unrelated searches. The minimalist, text-based format directs attention to the unexpected narrative structure rather than visual characters. The advertisement was categorized under surprise due to its reliance on sudden cognitive realization and attentional shift.

A6. Volvo – Volvo Trucks - The Epic Split feat. Van Damme (Surprise)

This advertisement features an unexpected performance in which Jean-Claude Van Damme performs a split between two moving trucks. The surprise emerges from the viewer's sudden realization of the difficulty and risk of the action being performed. Attention is captured through the violation of physical expectations rather than narrative or humor. The advertisement was categorized under surprise due to its reliance on abrupt attentional disruption and unexpected visual execution.

Works Cited

Missaglia, Anna Laura, et al. "The impact of emotions on recall: An empirical study on social ads." *Journal of Consumer Behaviour* 16.5 (2017): 424-433.

Mehta, Abhilasha, and Scott C. Purvis. "Reconsidering recall and emotion in advertising." *Journal of Advertising research* 46.1 (2006): 49-56.

Hussain, Shahid, Asim Qazi, and Abdul Salam. "Understanding the cascade effect: the role of emotions in opportunistic product recall outcomes." *Journal of Product & Brand Management*(2026): 1-17.

Panda, Tapan K., Tapas K. Panda, and Kamalesh Mishra. "Does Emotional Appeal Work in Advertising? The Rationality Behind Using Emotional Appeal to Create Favorable Brand Attitude." *IUP Journal of Brand Management* 10.2 (2013).

Aydýn, Duygu. "The relation between positive brand emotions and recall." *Online Journal of Communication and Media Technologies* 4.1 (2014): 173.

Why Some Words Stick: Contextual Diversity and Free-Recall Memorability Across 576 Words By Samaira Roy

Heritage Xperiential Learning School, Gurugram, Haryana, India

Abstract

Why are some words easier to bring back to mind than others? This study tested whether two prespecified word properties-contextual diversity, or the range of language contexts in which a word occurs, and concreteness-were associated with free-recall probability in a public 576-word dataset from the Penn Electrophysiology of Encoding and Retrieval Study (PEERS). Continuous values were analyzed using Pearson and Spearman correlations, followed by a standardized two-predictor ordinary least squares model and a log-transformed contextual-diversity sensitivity check. Recall probability ranged from 0.35 to 0.68 (M = 0.499, SD = 0.056). Contextual diversity was positively associated with recall (Pearson r = 0.210, 95% CI [0.131, 0.287], p < .001; Spearman ρ = 0.161, p < .001), whereas concreteness showed a weak bivariate relationship (Pearson r = 0.054, p = 0.199). In the standardized model, contextual diversity remained positive (β = 0.218, 95% CI [0.137, 0.298], p < .001), while concreteness was smaller (β = 0.076, p = 0.064); R^2 = 0.050. The same pattern remained after log transformation. The important point is not that contextual diversity explains memory by itself-it does not-but that a simple measure of how broadly a word appears across language contexts carried a consistent signal across several statistical checks, while concreteness did not. This raises a practical, testable question for future work: could vocabulary encountered across varied meaningful contexts become easier to retrieve later? The present data are associative, so that possibility requires direct experimental testing.

Keywords: word memorability; free recall; contextual diversity; concreteness; psycholinguistics; PEERS; secondary analysis; reproducibility; vocabulary learning; educational psychology

Introduction

Why do some words “stick” while others are harder to retrieve, even when people study them under similar conditions? Large word-list studies show that recallability is not random: lexical, semantic, and affective properties explain part of why some items are remembered more reliably than others (Aka et al., “Predicting Recall”). Broader work has also identified contributions from semantic and function-related features (Madan), extending classic findings that imagery, availability, and emotionality help distinguish more from less recallable nouns (Rubin and Friendly). Understanding these patterns matters to memory science, but it may also matter in more everyday settings-such as learning unfamiliar vocabulary, scientific terminology, or other information that has to be retrieved without a cue.

Memorability is also task-sensitive. In a large comparison of word recognition and free recall, the same 576 words showed only a weak rank-order relationship across the two tasks, even though broader semantic categories displayed more consistency (Aka et al., “Semantic

Determinants"). In other words, a word that is easy to recognize is not necessarily one that comes to mind easily without prompts. That distinction is important because real-world learning often depends on active retrieval-for example, recalling a term during an exam or producing a word while speaking-rather than simply recognizing the correct answer when it is shown.

Two properties are especially useful for a focused comparison. Contextual diversity indexes the number of language contexts in which a word appears and is a frequency-related measure developed from large subtitle corpora (Brysbaert and New). It can be thought of as a rough marker of how widely a word is distributed across different linguistic situations. Concreteness, by contrast, captures how strongly a word refers to something perceptible or imageable (Brysbaert et al.). These features represent two different routes by which a word might become memorable: one through breadth of linguistic experience and the other through the vividness of what the word represents. Frequency-related effects also vary across memory tasks: higher-frequency words can show an advantage in free recall (Hall), whereas lower-frequency words often show an advantage in recognition (Glanzer and Bowles). The 576-word PEERS pool itself was drawn from the University of South Florida free-association norms (Nelson et al.).

PEERS provides a strong behavioral foundation for examining these questions because it includes repeated memory testing across many sessions (Kahana et al.). Recognition findings from PEERS also illustrate why memory performance should be interpreted with task-specific measurement principles, including response confidence and bias-sensitive approaches (Weidemann and Kahana). At the same time, free recall is not identical to recollection, and recognition can reflect both recollection and familiarity (Yonelinas). Neuroimaging studies have reported partially separable correlates of these processes (Yonelinas et al.), while competing accounts and distributed retrieval models caution against reducing memory tasks to a simple one-region/one-function mapping (Wixted et al.; Wixted and Squire; Cabeza et al.; Rugg and Vilberg). Word memorability can also interact with neural retrieval signals (Xie et al.), but the present word-level behavioral analysis cannot identify a neural mechanism.

The present study asked a deliberately narrow question: across the 576-word PEERS free-recall dataset, how strongly are contextual diversity and concreteness associated with word-level recall probability? Three expectations were specified in advance: recall probability would vary across words; contextual diversity would show a positive association with recall; and concreteness would show a positive, but potentially smaller, association once contextual diversity was considered simultaneously. Rather than building a large predictive model, the analysis was intentionally kept small and interpretable. The aim was to see whether one simple feature would carry a stable signal across several checks-and, if it did, whether that pattern could generate a clearer hypothesis for future experimental work.

Methods

Study Design and Data Source

This study used public, aggregated word-level data and did not recruit new participants. The analysis file contained 576 unique nouns, complete free-recall probabilities, and

psycholinguistic feature values. The recall values trace to PEERS Experiment 4 and the 576-word materials used in published word-recall analyses (Aka et al., "Predicting Recall"; Kahana et al.). Public source materials were accessed through the Aka et al. project on the Open Science Framework (OSF) and the PEERS data archive.

OSF source project: https://osf.io/7phj6/

Recall probability was the proportion of participants who correctly recalled each word. In PEERS Experiment 4, 98 participants completed repeated delayed free-recall sessions using the same 576-word pool (Aka et al., "Predicting Recall"). The available analysis file contained no missing recall values.

Variables

The two prespecified predictors were concreteness and contextual diversity. Concreteness used the five-point rating framework reported for the word set, primarily based on the Brysbaert et al. norms (Brysbaert et al.; Aka et al., "Semantic Determinants"). Contextual diversity represented the number of different linguistic contexts in which a word appeared and followed the SUBTLX-US measure (Brysbaert and New; Aka et al., "Semantic Determinants").

Statistical Analysis

Analyses used continuous values throughout. Descriptive statistics summarized the recall distribution. Pearson correlations estimated linear relationships between each predictor and recall probability; Spearman correlations provided a rank-based check that was less sensitive to non-normality and influential values. A standardized ordinary least squares model then included concreteness and contextual diversity simultaneously, allowing direct comparison of coefficient magnitude.

Because contextual diversity was strongly right-skewed, a sensitivity analysis replaced raw contextual diversity with log(1 + contextual diversity), while retaining concreteness. This check tested whether the contextual-diversity result depended disproportionately on a small number of extremely high-frequency words. Two-sided p values and 95% confidence intervals were reported. Statistical verification was performed in Python 3.13.5 using NumPy 2.3.5 and SciPy 1.17.0. Because the two predictors and primary two-predictor model were specified in advance, no multiplicity adjustment was applied; interpretation emphasizes effect sizes, confidence intervals, and consistency across checks.

No causal interpretation was assigned to the associations. The analysis did not equate free recall with a pure measure of recollection, did not infer neural localization from word-level behavior, and did not conduct a new recall-recognition comparison because the exact recognition-probability vector was not part of the analysis file.

Ethical Considerations

No new participants were recruited. This secondary analysis used public, aggregated word-level data. Ethics approval and informed consent for the original PEERS studies are described in the cited source publications; no identifiable participant-level records were accessed.

Results

The 576-word dataset contained no missing recall values. Recall probability ranged from 0.35 to 0.68, with a mean of 0.499, median of 0.500, and standard deviation of 0.056 (Table 1; Fig. 1). The observed range closely matched the published PEERS recall range reported for the shared word set (Aka et al., "Semantic Determinants").

Table 1. Descriptive statistics for free-recall probability.

n	Mean	Median	SD	Minimum	Maximum
576	0.499	0.500	0.056	0.35	0.68

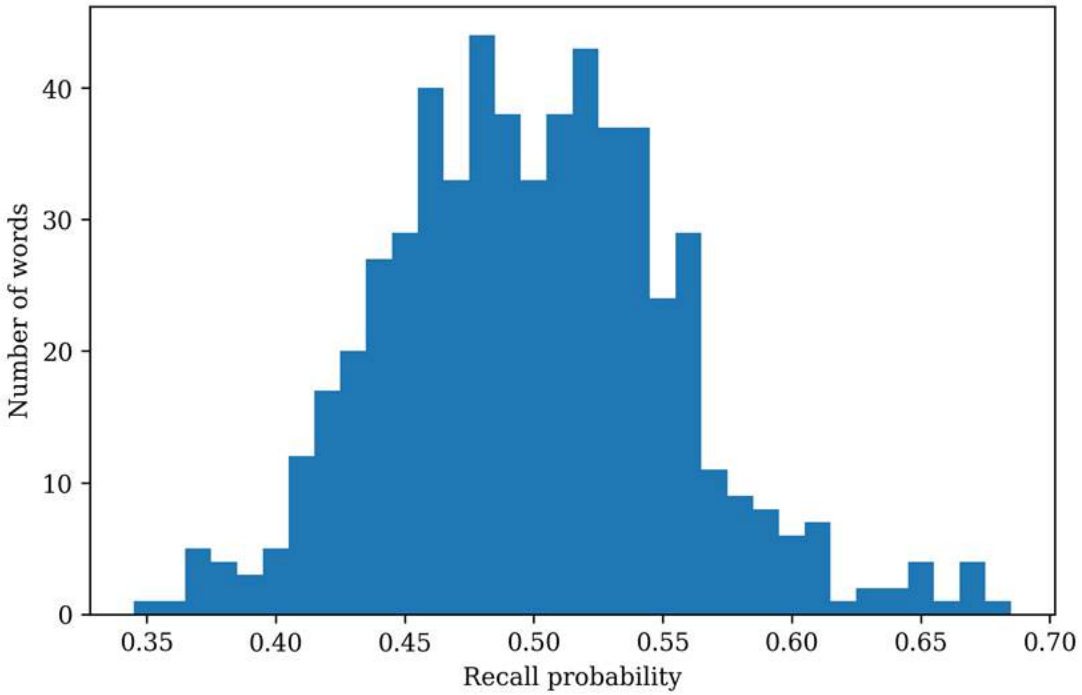

Fig. 1: Distribution of free-recall probability across the 576 words. Each bar shows the number of words within a recall-probability interval.

Contextual diversity showed a small-to-moderate positive linear association with recall probability (Pearson r = 0.210, 95% CI [0.131, 0.287], p < .001). The rank-based association was also positive (Spearman ρ = 0.161, p < .001) (Table 2; Fig. 2). Concreteness showed a substantially weaker relationship with recall (Pearson r = 0.054, 95% CI [-0.028, 0.135], p = 0.199; Spearman ρ = 0.075, p = 0.070) (Table 2; Fig. 3).

Table 2. Bivariate associations between word properties and recall probability.

Predictor	Pearson r	95% CI	p	Spearman ρ	p
Concreteness	0.054	[-0.028, 0.135]	0.199	0.075	0.070
Contextual diversity	0.210	[0.131, 0.287]	< .001	0.161	< .001

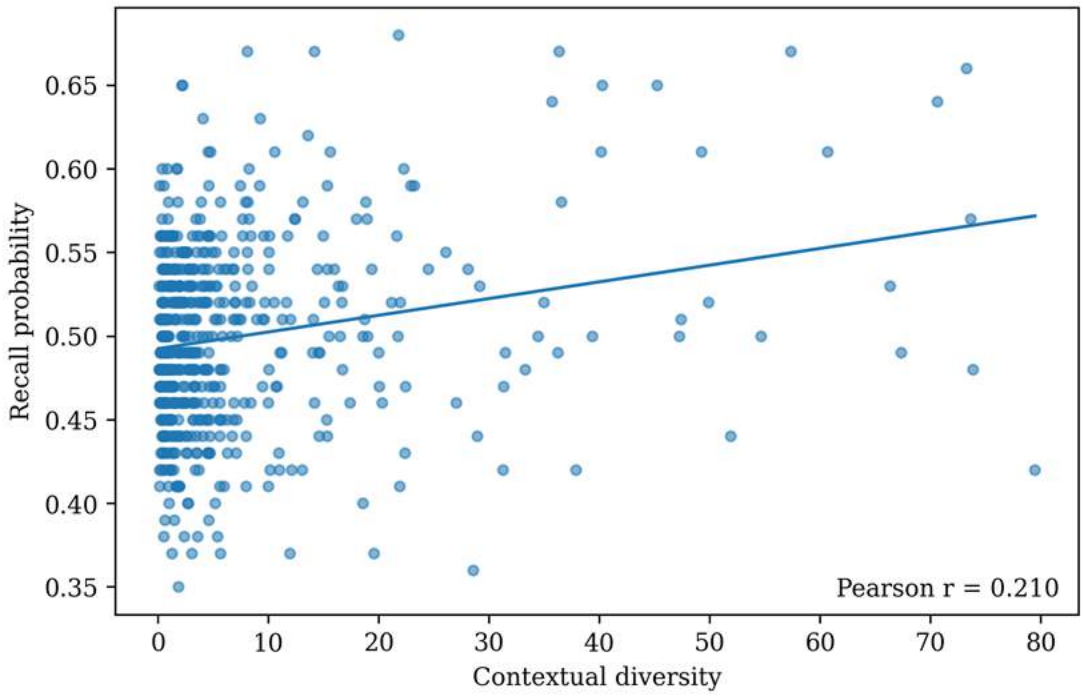

Fig. 2: Recall probability as a function of contextual diversity. The fitted line is shown for visual guidance; inference is reported in the text and Table 2.

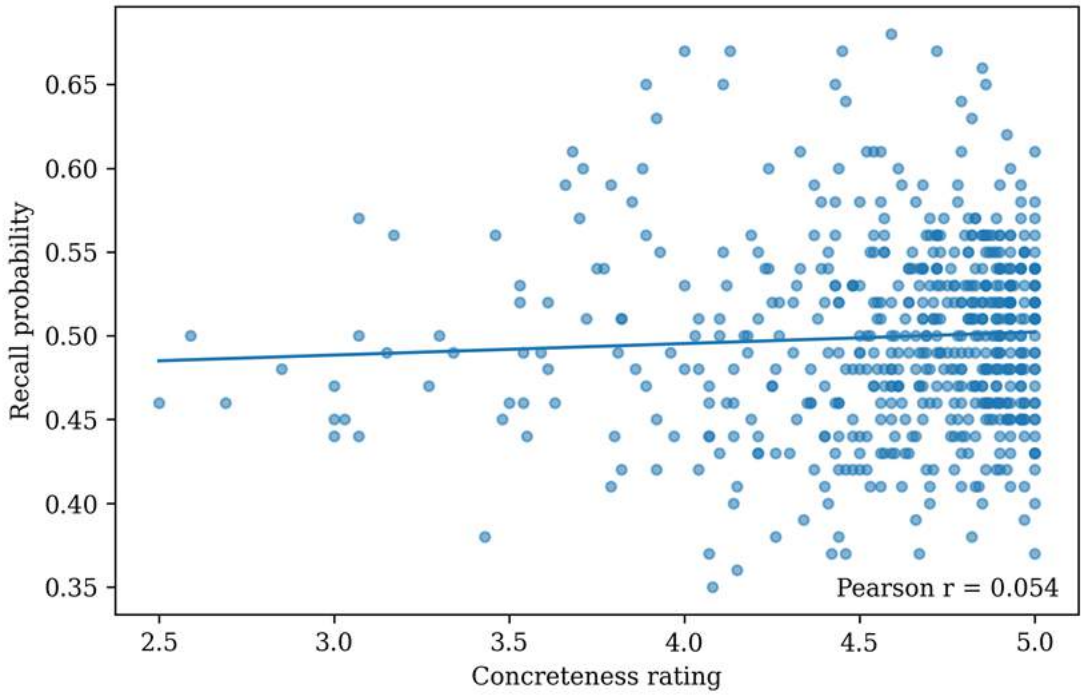

Fig. 3: Recall probability as a function of concreteness. The fitted line is shown for visual guidance; the bivariate association is weak.

In the standardized two-predictor model, contextual diversity remained positively associated with recall ($\beta = 0.218$, 95% CI [0.137, 0.298], $t = 5.32$, $p < .001$). The coefficient for

concreteness was smaller and did not meet the conventional .05 significance threshold (β = 0.076, 95% CI [-0.004, 0.156], t = 1.86, p = 0.064). The overall model was significant, F(2, 573) = 15.02, p < .001, with $R^2 = 0.050$ and adjusted $R^2 = 0.046$ (Table 3).

The sensitivity analysis produced the same substantive pattern. After log transformation, contextual diversity remained positive (β = 0.212, 95% CI [0.132, 0.293], p < .001), while concreteness remained smaller (β = 0.070, 95% CI [-0.011, 0.150], p = 0.088). The sensitivity model explained $R^2 = 0.048$ of recall variance.

Table 3. Standardized two-predictor model and sensitivity analysis.

Model / predictor	β	95% CI	t	p	R^2
Primary: concreteness	0.076	[-0.004, 0.156]	1.86	0.064	0.050
Primary: contextual diversity	0.218	[0.137, 0.298]	5.32	< .001	0.050
Sensitivity: concreteness	0.070	[-0.011, 0.150]	1.71	0.088	0.048
Sensitivity: log(1 + contextual diversity)	0.212	[0.132, 0.293]	5.20	< .001	0.048

Discussion

The most meaningful result was not simply that one statistical test reached significance. Contextual diversity pointed in the same positive direction across Pearson and Spearman correlations, the standardized two-predictor model, and the log-transformed sensitivity analysis. The effect was modest, but its consistency makes it less dependent on any single analytic choice. This pattern also agrees with larger PEERS-based analyses in which contextual diversity contributes positively to recall (Aka et al., “Predicting Recall”; Aka et al., “Semantic Determinants”). Together, these checks suggest that the breadth of contexts in which a word appears carries information about its free-recall memorability that is not captured by concreteness alone.

The concreteness result was more nuanced. Concreteness showed only a weak bivariate relationship and a small adjusted coefficient in the two-predictor model. This does not mean that imageable or perceptible words are generally unimportant to memory; broader studies have found positive effects of concreteness and related semantic properties (Madan; Aka et al., “Predicting Recall”; Rubin and Friendly). Instead, the comparison suggests that, within this particular word pool, the breadth of a word’s linguistic contexts was more informative for free recall than how concrete the word was. One possible interpretation is that contextual diversity partly reflects the number of different retrieval routes or associations available for a word, although the present analysis cannot test that mechanism directly.

The result also illustrates why frequency-related word properties should be interpreted within a specific memory task. Classic studies reported a high-frequency advantage in free recall and a low-frequency advantage in recognition (Hall; Glanzer and Bowles). More recent work on the shared PEERS word set likewise found opposite standardized frequency-related effects

across recognition and recall (Aka et al., “Semantic Determinants”). Contextual diversity therefore should not be treated as a universal memorability enhancer; its association depends on the retrieval conditions and outcome being studied.

The cognitive literature places an important boundary on that interpretation. Free recall requires self-initiated retrieval, but recollection is not identical to free recall, and recognition can involve both recollection and familiarity (Yonelinas). Although neuroimaging studies report differences between recollection- and familiarity-related activity (Yonelinas et al.), competing accounts of medial temporal lobe function and distributed retrieval models argue against reducing memory tasks to single anatomical modules (Wixted et al.; Wixted and Squire; Cabeza et al.; Rugg and Vilberg). Word memorability can interact with neural retrieval signals (Xie et al.), but the present behavioral associations cannot identify where or why such signals arise.

The model explained about 5.0% of recall variance. At first glance, that number may look small, but memorability is a multi-determined outcome shaped by many correlated lexical, semantic, affective, associative, and contextual properties (Madan; Aka et al., “Predicting Recall”; Aka et al., “Semantic Determinants”). With only two deliberately chosen predictors, the value of the model is not that it explains memory as a whole. Its value is that one simple, interpretable feature repeatedly carried information about recall while the other was much weaker. In a complex behavior such as memory, a modest but reproducible signal can be more useful than a larger effect that depends on a highly specific model.

This pattern is potentially relevant to learning design, but only as a hypothesis to test. In classrooms and digital study tools, new vocabulary is often introduced with a definition and a small number of repeated examples. The present result suggests a different experimental possibility: keep total exposure the same, but spread those exposures across varied, meaningful contexts. If future controlled studies show a causal benefit, this could inform vocabulary teaching, second-language learning, revision strategies, and memory-supportive educational apps-especially for technical terms that students need to produce from memory rather than simply recognize. The current study does not demonstrate such a benefit; it identifies a concrete and testable route from a basic memory finding to a possible real-world application.

Limitations

Several limitations should be kept in view. The analysis used aggregated word-level probabilities rather than participant-level trials, so it could not model individual differences or within-person dependencies. It was observational, so associations cannot establish that manipulating contextual diversity or concreteness would cause changes in recall. Only two predictors were tested, and both may correlate with unmodeled semantic or lexical properties. Recognition probabilities were not included in the analysis file, so published recognition findings were used only for context rather than presented as new results.

In addition, the ordinary least squares models treat aggregated recall probabilities as continuous outcomes; they do not model the underlying trial-level binomial process or

uncertainty in each word-level proportion. The regression coefficients are therefore best interpreted as descriptive item-level associations rather than participant-level causal effects.

Future Directions

The most immediate next step would be to test the practical hypothesis suggested by this analysis rather than assume it is true. A controlled learning experiment could hold the number of exposures constant while varying contextual diversity: for example, students could learn the same unfamiliar or low-familiarity academic terms either through repeated exposure in one narrow context or through the same number of exposures across several meaningful sentences, examples, or subject situations. Delayed free recall and recognition could then be compared. Such a design would directly test whether varied-context exposure causes better retrieval rather than merely being associated with it. A second extension would add a larger, prespecified set of lexical, semantic, affective, and associative predictors to the 576-word dataset to test whether contextual diversity retains incremental explanatory value. Adding recognition probabilities would allow a direct within-word comparison of whether the same features predict recall and recognition differently (Aka et al., "Semantic Determinants"). Finally, participant-level behavioral and electrophysiological PEERS data could be used in multilevel analyses to ask whether the item-level pattern generalizes across individuals and whether it relates to encoding or retrieval signals (Kahana et al.). Together, these studies could connect a simple word-level finding to broader questions about how language experience shapes memory and, eventually, how learning materials might be designed.

Conclusion

Across 576 words, contextual diversity carried a small but consistently positive signal of free-recall memorability, while concreteness was much weaker in the restricted two-predictor model. The broader implication is that memorability may depend not only on what a word means, but also on the range of linguistic environments in which it appears. This does not yet tell us how vocabulary should be taught, and the present analysis cannot establish causality. It does, however, provide a clear next question: if the number of exposures is held constant, does learning a word across varied meaningful contexts make it easier to retrieve later? For a deliberately simple secondary analysis, that is a useful outcome-not a final answer, but a reproducible clue about how language experience may shape what comes back to mind.

Data Availability

Public source materials are available through the Aka et al. memorability project on the Open Science Framework and the PEERS data resources cited in this manuscript. The present analysis uses aggregated word-level values and no identifiable participant-level records.

Open Science Framework: https://osf.io/7phj6/

Author Contributions

Samaira Roy conducted the secondary analysis, interpreted the results, prepared the manuscript, and is responsible for the final scientific content. Research mentorship and editorial support are acknowledged separately.

Acknowledgments

The project received research mentorship support from the scientists at Rapture Biotechnologies

AI-Assisted Editing Disclosure

AI tools were used for language editing, structural organization, and document-formatting assistance. The author remains responsible for the scientific interpretation, numerical results, source verification, and final submitted text.

Works Cited

Aka, A., S. Bhatia, and J. McCoy. "Semantic Determinants of Memorability." *Cognition*, vol. 239, 2023, article 105497. https://doi.org/10.1016/j.cognition.2023.105497

Aka, A., T. D. Phan, and M. J. Kahana. "Predicting Recall of Words and Lists." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 47, no. 5, 2021, pp. 765–784. https://doi.org/10.1037/xlm0000964

Brysbaert, M., and B. New. "Moving beyond Kučera and Francis: A Critical Evaluation of Current Word Frequency Norms and the Introduction of a New and Improved Word Frequency Measure for American English." *Behavior Research Methods*, vol. 41, no. 4, 2009, pp. 977–990. https://doi.org/10.3758/BRM.41.4.977

Brysbaert, M., A. B. Warriner, and V. Kuperman. "Concreteness Ratings for 40 Thousand Generally Known English Word Lemmas." *Behavior Research Methods*, vol. 46, no. 3, 2014, pp. 904–911. https://doi.org/10.3758/s13428-013-0403-5

Cabeza, R., et al. "The Parietal Cortex and Episodic Memory: An Attentional Account." *Nature Reviews Neuroscience*, vol. 9, no. 8, 2008, pp. 613–625. https://doi.org/10.1038/nrn2459

Glanzer, M., and N. Bowles. "Analysis of the Word-Frequency Effect in Recognition Memory." *Journal of Experimental Psychology: Human Learning and Memory*, vol. 2, no. 1, 1976, pp. 21–31. https://doi.org/10.1037/0278-7393.2.1.21

Hall, J. F. "Learning as a Function of Word-Frequency." *The American Journal of Psychology*, vol. 67, no. 1, 1954, pp. 138–140. https://doi.org/10.2307/1418080

Kahana, M. J., et al. "The Penn Electrophysiology of Encoding and Retrieval Study." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 50, no. 9, 2024, pp. 1421–1443. https://doi.org/10.1037/xlm0001319

Madan, C. R. "Exploring Word Memorability: How Well Do Different Word Properties Explain Item Free-Recall Probability?" *Psychonomic Bulletin & Review*, vol. 28, 2021, pp. 583–595. https://doi.org/10.3758/s13423-020-01820-w

Nelson, D. L., C. L. McEvoy, and T. A. Schreiber. "The University of South Florida Free Association, Rhyme, and Word Fragment Norms." *Behavior Research Methods, Instruments, & Computers*, vol. 36, no. 3, 2004, pp. 402–407. https://doi.org/10.3758/BF03195588

Rubin, D. C., and M. Friendly. "Predicting Which Words Get Recalled: Measures of Free Recall, Availability, Goodness, Emotionality, and Pronunciability for 925 Nouns." *Memory & Cognition*, vol. 14, no. 1, 1986, pp. 79–94. https://doi.org/10.3758/BF03209231

Rugg, M. D., and K. L. Vilberg. "Brain Networks Underlying Episodic Memory Retrieval." *Current Opinion in Neurobiology*, vol. 23, no. 2, 2013, pp. 255–260. https://doi.org/10.1016/j.conb.2012.11.005

Weidemann, C. T., and M. J. Kahana. "Assessing Recognition Memory Using Confidence Ratings and Response Times." *Royal Society Open Science*, vol. 3, no. 4, 2016, article 150670. https://doi.org/10.1098/rsos.150670

Wixted, J. T., L. Mickes, and L. R. Squire. “Measuring Recollection and Familiarity in the Medial Temporal Lobe.” *Hippocampus*, vol. 20, no. 11, 2010, pp. 1195–1205. https://doi.org/10.1002/hipo.20854

Wixted, J. T., and L. R. Squire. “The Medial Temporal Lobe and the Attributes of Memory.” *Trends in Cognitive Sciences*, vol. 15, no. 5, 2011, pp. 210–217. https://doi.org/10.1016/j.tics.2011.03.005

Xie, W., et al. “Memorability of Words in Arbitrary Verbal Associations Modulates Memory Retrieval in the Anterior Temporal Lobe.” *Nature Human Behaviour*, vol. 4, 2020, pp. 937–948. https://doi.org/10.1038/s41562-020-0901-2

Yonelinas, A. P. “The Nature of Recollection and Familiarity: A Review of 30 Years of Research.” *Journal of Memory and Language*, vol. 46, no. 3, 2002, pp. 441–517. https://doi.org/10.1006/jmla.2002.2864

Yonelinas, A. P., et al. “Separating the Brain Regions Involved in Recollection and Familiarity in Recognition Memory.” *Journal of Neuroscience*, vol. 25, no. 11, 2005, pp. 3002–3008. https://doi.org/10.1523/JNEUROSCI.5295-04.2005

The Chemistry Underlying Chronic Stress-Induced Gut Dysbiosis By Rajit Reddi

Abstract

Chronic psychological stress has increasingly been recognized as a major factor influencing the composition and metabolic activity of the gut microbiome. Through activation of the hypothalamic-pituitary-adrenal (HPA) axis and sympathetic nervous system, chronic stress alters levels of cortisol, catecholamines, inflammatory cytokines, and reactive oxygen species, all of which contribute to changes in microbial diversity and function. These alterations, often referred to as gut dysbiosis, can reduce beneficial bacterial populations, increase pathogenic organisms, impair gut barrier integrity, and disrupt microbial metabolite production. This meta-analysis examines the chemical and biochemical pathways through which chronic stress reshapes the gut microbiome, with a focus on stress hormones, neurotransmitters, short-chain fatty acids, oxidative stress, and gut permeability. Additionally, it evaluates biomarkers commonly associated with stress-induced dysbiosis and explores therapeutic strategies including probiotics, prebiotics, dietary interventions, and pharmacological compounds. Understanding the chemistry underlying stress-induced dysbiosis may lead to more personalized and effective treatments for both gastrointestinal and mental health disorders.

I. Introduction

A. Background and Rationale

1. Overview of the Gut Microbiome

The gut microbiome refers to the trillions of microorganisms, including bacteria, fungi, viruses, and archaea, that reside in the gastrointestinal tract. These microbes play critical roles in digestion, nutrient absorption, vitamin synthesis, immune regulation, and protection against pathogenic organisms. The majority of gut bacteria belong to the phyla Firmicutes and Bacteroidetes, although Proteobacteria, Actinobacteria, and Verrucomicrobia are also important contributors to gut function.

A healthy gut microbiome is characterized by high microbial diversity and a balanced ratio of beneficial to harmful organisms. Beneficial microbes, such as Lactobacillus, Bifidobacterium, and Faecalibacterium prausnitzii, support intestinal barrier integrity, reduce inflammation, and produce important metabolites like short-chain fatty acids (SCFAs). In contrast, dysbiosis occurs when this balance is disrupted, often leading to increased inflammation, impaired metabolism, and disease.

2. Overview of the Gut-Brain Axis and Its Role in Maintaining Gut Homeostasis

The gut-brain axis is a bidirectional communication network between the gastrointestinal tract and the central nervous system. This communication occurs through several pathways, including neural signaling through the vagus nerve, hormonal signaling through the HPA axis, immune system mediators, and microbial metabolites.

Under normal conditions, the gut-brain axis helps maintain homeostasis by regulating digestion, immune function, and mood. However, chronic stress can disrupt this system by increasing the release of cortisol and catecholamines. These stress-related compounds can directly affect microbial growth, intestinal permeability, and inflammation, ultimately altering the composition of the microbiome (Rusch et al., 2023; Verma et al., 2024).

3. Why Understanding the Chemical Basis of Stress-Induced Gut Dysbiosis is Important

Stress-induced gut dysbiosis has been linked to a wide range of physical and mental health disorders, including inflammatory bowel disease, irritable bowel syndrome, obesity, anxiety, depression, and neurodegenerative diseases. Although many studies have shown a connection between stress and gut microbial changes, fewer studies have focused specifically on the chemical mechanisms responsible for these changes.

Understanding how cortisol, neurotransmitters, inflammatory molecules, and microbial metabolites interact may provide opportunities for more targeted therapies. This is especially important given the growing interest in personalized medicine and microbiome-based interventions.

B. Objectives of the Meta-Analysis

The primary objective of this meta-analysis is to examine the biochemical pathways that connect chronic stress to gut dysbiosis. Specifically, this paper aims to:

1. Investigate the chemical and metabolic mechanisms driving gut microbiome changes under chronic stress.
2. Identify biomarkers associated with stress-induced dysbiosis.
3. Evaluate therapeutic interventions targeting chemical imbalances in the gut microbiome.

II. Methodology

A. Literature Search Strategy

Relevant literature was identified through searches of databases including Google Scholar, PubMed, ScienceDirect, and JSTOR. Keywords used in the search included “chronic stress,” “gut microbiome,” “gut dysbiosis,” “cortisol,” “HPA axis,” “short-chain fatty acids,” “oxidative stress,” “intestinal permeability,” “leaky gut,” “catecholamines,” and “neurotransmitters.”

Studies were included if they:

- Were peer-reviewed and published after 2015
- Focused on the chemical or metabolic aspects of stress-induced dysbiosis
- Included human or rodent models
- Reported measurable changes in microbial composition, metabolites, or biomarkers

Studies were excluded if they:

- Focused only on acute stress
- Did not include biochemical measurements
- Examined unrelated gastrointestinal diseases without discussing stress

B. Data Extraction and Statistical Analysis

Data extracted from the selected studies included:

- Stress-related biomarkers such as cortisol, norepinephrine, inflammatory cytokines, and reactive oxygen species
- Changes in microbial composition, including decreases in beneficial bacteria and increases in pathogenic organisms
- Alterations in microbial metabolites such as butyrate, acetate, propionate, serotonin, and GABA

III. Chemistry of Chronic Stress-Induced Gut Dysbiosis

A. Stress Hormones and Their Impact on the Gut Microbiome

Cortisol is the primary stress hormone released through activation of the HPA axis. During chronic stress, prolonged elevation of cortisol can weaken intestinal barrier function, suppress immune responses, and alter the availability of nutrients in the gut. Elevated cortisol has been associated with reduced microbial diversity and lower abundance of beneficial bacteria such as Lactobacillus and Bifidobacterium (Rusch et al., 2023; Ying et al., 2025). Catecholamines such as norepinephrine and dopamine are released during stress through activation of the sympathetic nervous system. These neurotransmitters can directly affect bacterial growth. For example, norepinephrine has been shown to stimulate the growth of pathogenic bacteria like Escherichia coli and Salmonella by increasing iron availability (Karl et al., 2018). Stress hormones may also increase bacterial virulence by promoting biofilm formation, adhesion to intestinal tissue, and toxin production. This can contribute to greater inflammation and worsen dysbiosis.

B. Short-Chain Fatty Acids and Metabolic Shifts

Butyrate is one of the most important SCFAs produced by gut bacteria through the fermentation of dietary fiber. It serves as the primary energy source for colonocytes and helps maintain gut barrier integrity. Chronic stress is often associated with reductions in butyrate-producing bacteria such as Faecalibacterium prausnitzii, Roseburia, and members of Lachnospiraceae. Lower butyrate levels can weaken the intestinal barrier and increase inflammation (Verma et al., 2024; Ying et al., 2025; Zhang et al., 2019). Recent evidence also suggests that SCFAs are not only passive markers of microbial activity, but active chemical mediators of gut-brain communication. Acute stress has been shown to reduce caecal SCFA concentrations, especially butyrate and acetate, while SCFAs can protect intestinal and blood-brain barrier models from LPS-induced disruption. This suggests that stress-related reductions in microbial metabolites may directly weaken both gut and brain barrier function (Rosell-Cardona et al., 2025). Butyrate may also influence stress-related mood symptoms through mechanisms involving oxidative stress and ferroptosis. Ferroptosis is an iron-dependent form of cell death associated with lipid peroxidation and redox imbalance. Because chronic stress increases oxidative stress, butyrate's ability to regulate inflammatory and redox-related pathways may help explain why loss of butyrate-producing bacteria is associated with depression-like behavior and gut-brain axis dysfunction (Wang et al., 2025).

Acetate and propionate also play important roles in metabolism and immune regulation. Changes in these SCFAs may alter appetite regulation, glucose metabolism, inflammation, and neurotransmitter synthesis. Some studies suggest that stress shifts the ratio of acetate, propionate, and butyrate, contributing to metabolic dysfunction. Importantly, lower microbial diversity and lower inferred capacity for butyrate production have been associated with greater cortisol reactivity to stress in healthy adults, suggesting that SCFA-producing microbiota may influence the body's endocrine response to stress (Karner et al., 2025).

A. Oxidative Stress and Gut Dysbiosis

Chronic stress increases the production of reactive oxygen species (ROS), which can damage intestinal cells and disrupt microbial balance. Excessive ROS may create an environment that favors pathogenic or inflammation-tolerant bacteria while reducing populations of beneficial anaerobes and antioxidant-producing microbes (Wu et al., 2025; Li et al., 2023; Sun et al., 2024). Oxidative stress and dysbiosis can form a self-reinforcing cycle. Dysbiosis may increase ROS production and inflammatory signaling, while oxidative stress can further damage the intestinal epithelium, alter microbial niches, and promote immune activation. This cycle is especially important because chronic inflammation and oxidative stress can impair epithelial barrier function and allow microbial products to enter systemic circulation (Li et al., 2023; Mostafavi Abdolmaleky et al., 2024; Herrera-Quintana et al., 2026).

Certain bacteria produce antioxidant compounds that help maintain redox balance in the gut. For example, Lactobacillus and Bifidobacterium species may help reduce oxidative stress by increasing glutathione levels, producing antioxidant enzymes, and supporting anti-inflammatory immune signaling. A healthy microbiome may therefore act as a biochemical buffer against oxidative damage, while stress-induced dysbiosis reduces this protective capacity. This redox imbalance is relevant to both gastrointestinal and neurological health. Because oxidative stress can influence immune signaling, epithelial permeability, and neuroinflammation, antioxidant-producing microbes and antioxidant-rich diets may represent important therapeutic targets in stress-induced dysbiosis (Kunst et al., 2023; Shandilya et al., 2022).

D. Gut Permeability and Chemical Mediators

The intestinal barrier is maintained by tight junction proteins such as occludin, claudin, and zonula occludens-1 (ZO-1). Chronic stress can reduce the expression of these proteins, increasing gut permeability. This condition, often called “leaky gut,” allows harmful molecules to enter the bloodstream (Kearns et al., 2025; Rusch et al., 2023). One particularly important mechanism involves glucocorticoid signaling. Chronic stress and elevated glucocorticoids have been associated with reduced claudin-1 expression, which weakens epithelial barrier integrity. This provides a direct molecular link between HPA-axis activation and increased intestinal permeability (Zheng et al., 2017).

Lipopolysaccharide (LPS), a component of gram-negative bacterial cell walls, can enter circulation when gut permeability increases. Elevated LPS levels activate immune responses and stimulate the release of inflammatory cytokines such as TNF-alpha, IL-6, and IL-1β. This contributes to systemic inflammation and may worsen both mental and gastrointestinal disorders.

Stress-induced barrier dysfunction can therefore amplify dysbiosis by allowing microbial products to trigger immune signaling outside the gut. In turn, inflammatory cytokines can further alter microbial composition, increase oxidative stress, and shift tryptophan metabolism toward the kynurenine pathway. This creates a feed-forward loop connecting stress hormones, barrier disruption, microbial metabolites, and neuroinflammation.

E. Neurotransmitters and the Gut-Brain Axis

Approximately 90% of serotonin is produced in the gut. Chronic stress can reduce serotonin production by altering microbial composition and tryptophan metabolism. Lower serotonin levels may contribute to anxiety, depression, and altered gut motility (Kearns et al., 2025; Bastiaanssen et al., 2019). Certain bacteria, including Lactobacillus and Bifidobacterium, can produce gamma-aminobutyric acid (GABA), an inhibitory neurotransmitter involved in reducing anxiety. Reduced levels of these bacteria may decrease GABA production during chronic stress. Stress-related changes in glutamate and dopamine may also affect gut-brain communication. Dysregulation of these neurotransmitters has been linked to mood disorders and altered gastrointestinal function.

F. Tryptophan Metabolism and the Kynurenine Pathway

A major chemical pathway that should be included in stress-induced gut dysbiosis is tryptophan metabolism. Tryptophan is an essential amino acid that can be metabolized into several biologically active products, including serotonin, kynurenine, indole derivatives, and downstream neuroactive metabolites. Under healthy conditions, the gut microbiome helps regulate how tryptophan is used by the host. However, chronic stress and inflammation can shift tryptophan metabolism away from serotonin production and toward the kynurenine pathway.

This shift is important because inflammatory cytokines can activate enzymes such as indoleamine 2,3-dioxygenase (IDO), which converts tryptophan into kynurenine. As more tryptophan is directed toward kynurenine metabolism, less may be available for serotonin synthesis. This provides a chemical explanation for how chronic stress-induced dysbiosis may contribute to mood-related symptoms such as anxiety and depression. Some kynurenine pathway products, including quinolinic acid, may also promote neuroinflammation and oxidative stress (Kearns et al., 2025).

The kynurenine pathway connects multiple parts of stress-induced dysbiosis: gut permeability, inflammation, neurotransmitter disruption, and neuroimmune signaling. If chronic stress weakens the gut barrier, bacterial products such as LPS can enter circulation and activate inflammatory pathways. These cytokine responses may further increase IDO activity, creating a cycle in which dysbiosis promotes inflammation, inflammation shifts tryptophan metabolism, and altered tryptophan metabolism affects brain function.

G. Metabolomic Profiling of Stress-Induced Dysbiosis

Metabolomics is especially useful for studying stress-induced gut dysbiosis because it focuses on the chemical products produced or modified by gut microbes. While 16S rRNA sequencing can identify which microbes are present, metabolomic techniques such as gas chromatography-mass spectrometry (GC-MS), liquid chromatography-mass spectrometry

(LC-MS), and nuclear magnetic resonance (NMR) spectroscopy can reveal what those microbes are doing chemically (Verma et al., 2024; Ying et al., 2025).

In the context of chronic stress, metabolomic profiling can be used to measure changes in short-chain fatty acids, bile acids, amino acids, neurotransmitter precursors, and inflammatory metabolites. For example, reduced butyrate may indicate a loss of beneficial fiber-fermenting bacteria, while altered tryptophan or kynurenine levels may indicate changes in inflammatory and neurotransmitter-related pathways. Measuring these compounds provides stronger evidence than microbial composition alone because it directly connects dysbiosis to biochemical function.

A stronger future direction for this field is personalized microbiome-based intervention. Rather than treating all stress-related dysbiosis the same way, researchers could use metabolomic profiles to determine whether a patient primarily shows reduced SCFA production, altered tryptophan metabolism, increased oxidative stress, or elevated inflammatory biomarkers. This would allow more targeted treatment using specific probiotics, prebiotics, dietary changes, or pharmacological compounds.

H. Evidence Map for Major Mechanisms

Mechanism	**Key Chemical Change**	**Microbial/Physiological Effect**	**Best Source Type**
HPA axis activation	Cortisol/corticosterone increase	Altered gut permeability, motility, nutrient availability, and microbial diversity	HPA axis review
Sympathetic activation	Norepinephrine/dopamine increase	Increased pathogen growth, adhesion, and virulence	Catecholamine/microbial endocrinology studies
SCFA disruption	Butyrate, acetate, and propionate shifts	Weakened epithelial barrier and increased inflammation	SCFA/metabolomics studies
Oxidative stress	ROS increase, antioxidant capacity decrease	Damage to gut epithelium and selection for inflammation-tolerant bacteria	Oxidative stress studies
Barrier dysfunction	Occludin, claudin, and ZO-1 decrease; LPS increase	Immune activation and systemic inflammation	Leaky gut/permeability studies

Tryptophan metabolism	Serotonin decrease; kynurenine/quinolinic acid increase	Neuroinflammation and mood-related symptoms	Kynurenine pathway reviews
Therapeutic intervention	SCFA restoration, cytokine reduction, improved microbial diversity	Improved barrier function and reduced stress-related symptoms	Probiotic/prebiotic/diet studies

IV. Therapeutic Interventions Targeting the Chemical Imbalance

A. Probiotics and Prebiotics

Probiotics are live microorganisms that may help restore microbial balance and improve biochemical function within the gut microbiome. Several strains of Lactobacillus and Bifidobacterium have been shown to increase short-chain fatty acid production, reduce inflammatory cytokine release, and strengthen intestinal barrier integrity. These bacteria may also influence neurotransmitter production, including gamma-aminobutyric acid (GABA) and serotonin precursors, further supporting gut-brain communication (Bastiaanssen et al., 2019; Verma et al., 2024).

One important mechanism through which probiotics exert beneficial effects is through restoration of butyrate-producing microbial communities. Increased butyrate production may improve colonocyte metabolism, reinforce tight junction proteins such as occludin and claudin, and reduce intestinal permeability. Some studies have also shown that probiotic supplementation decreases circulating inflammatory biomarkers including IL-6 and TNF-α, suggesting that microbiome-targeted interventions may help reduce systemic inflammation associated with chronic stress (Wu et al., 2025).

Prebiotics such as inulin, fructooligosaccharides, galactooligosaccharides, and resistant starch serve as fermentable substrates for beneficial gut bacteria. Fermentation of these compounds promotes production of SCFAs including butyrate, acetate, and propionate. Increased SCFA availability may improve gut barrier integrity, regulate immune signaling, and reduce oxidative stress. Resistant starch has been particularly associated with enrichment of Faecalibacterium prausnitzii and Roseburia, both important butyrate-producing taxa.

B. Dietary Interventions

Dietary interventions represent one of the most accessible approaches for improving stress-induced dysbiosis because diet directly shapes microbial composition and metabolic activity. Diets rich in dietary fiber, fermented foods, polyphenols, and omega-3 fatty acids may increase microbial diversity and support production of beneficial metabolites such as butyrate (Verma et al., 2024).

High-fiber diets promote fermentation by anaerobic bacteria and increase SCFA production, particularly butyrate. Increased butyrate availability supports colonocyte energy metabolism and enhances expression of tight junction proteins that help maintain intestinal

barrier integrity. Fermented foods such as yogurt, kefir, kimchi, miso, natto, and kombucha may also introduce beneficial microbial species and microbial metabolites that support gut homeostasis.

Omega-3 fatty acids may further improve gut health by reducing oxidative stress and suppressing inflammatory signaling pathways such as NF-κB activation. Some evidence suggests omega-3 supplementation increases populations of anti-inflammatory bacteria while reducing stress-associated inflammatory cytokines.

Traditional Eastern dietary patterns may provide additional protection against stress-induced dysbiosis due to their greater intake of fermented foods, legumes, seaweed, vegetables, and diverse plant fibers. In contrast, Western diets high in saturated fat, refined sugar, and ultra-processed foods have been associated with reduced microbial diversity, lower SCFA production, and increased inflammatory markers. These dietary differences may partially explain variations in microbiome composition and metabolic disease prevalence between populations.

C. Pharmacological and Natural Compounds

Several pharmacological and naturally derived compounds may help counteract the biochemical effects of chronic stress on the gut microbiome. Polyphenols, including flavonoids found in green tea, berries, cocoa, and turmeric, possess antioxidant and anti-inflammatory properties that may reduce ROS accumulation and improve microbial balance. Curcumin, a bioactive compound found in turmeric, has been shown to suppress inflammatory cytokine production and improve intestinal barrier function in some experimental models.

Natural compounds may also interact directly with microbial metabolism. Polyphenols can act as prebiotic-like substrates for certain beneficial bacteria, while microbial metabolism of these compounds produces additional bioactive metabolites with anti-inflammatory effects. This creates a reciprocal relationship in which the microbiome modifies dietary compounds and the resulting metabolites further influence microbiome composition.

Certain antidepressants and stress-regulating pharmacological compounds may indirectly influence the gut microbiome through modulation of neurotransmitter pathways and HPA axis activity. For example, selective serotonin reuptake inhibitors (SSRIs) may alter gut serotonin signaling and microbial composition. Corticosteroid regulators that reduce chronic cortisol elevation may also help restore intestinal barrier integrity and reduce dysbiosis-associated inflammation (Rusch et al., 2023).

Emerging therapeutic approaches include psychobiotics, engineered probiotics, and personalized microbiome interventions guided by metabolomic profiling. These strategies aim to target specific biochemical abnormalities, such as reduced SCFA production, altered tryptophan metabolism, or elevated inflammatory biomarkers, rather than treating dysbiosis as a single generalized condition.

V. Discussion and Future Directions

Overall, chronic stress alters the gut microbiome through a combination of hormonal, inflammatory, oxidative, and metabolic pathways. These changes reduce beneficial bacteria,

increase pathogenic organisms, weaken gut barrier integrity, and disrupt neurotransmitter production. Many studies rely on rodent models, which may not fully reflect the complexity of the human microbiome. Additionally, there is significant variation between studies in terms of stress models, dietary factors, and methods for measuring microbial changes. Future studies should focus on personalized microbiome interventions based on chemical profiling. Advances in metabolomics, machine learning, and microbiome sequencing may help identify which treatments work best for specific individuals. This could connect especially well with your broader interest in microbiome informatics and building personalized diet recommendations based on microbial metabolite production.

VI. Conclusion

Chronic stress has profound effects on the gut microbiome through its influence on hormones, neurotransmitters, inflammation, oxidative stress, and gut barrier function. These chemical changes contribute to dysbiosis and may increase the risk of both gastrointestinal and mental health disorders. Although more research is needed, therapeutic interventions targeting microbial metabolites and biochemical pathways may offer promising strategies for improving health outcomes.

Works Cited

Bastiaanssen, Thomaz F S, et al. “Making Sense of … the Microbiome in Psychiatry.” *International Journal of Neuropsychopharmacology*, vol. 22, no. 1, 7 Aug. 2018, pp. 37–52, academic.oup.com/ijnp/article/22/1/37/5067516, https://doi.org/10.1093/ijnp/pyy067. Accessed 10 Aug. 2019.

Chasiotis, Helen, and Scott P Kelly. “Effect of Cortisol on Permeability and Tight Junction Protein Transcript Abundance in Primary Cultured Gill Epithelia from Stenohaline Goldfish and Euryhaline Trout.” *General and Comparative Endocrinology*, vol. 172, no. 3, 30 Apr. 2011, pp. 494–504, https://doi.org/10.1016/j.ygcen.2011.04.023. Accessed 13 Oct. 2024.

Daniel Efrain Molotla-Torres, et al. “Role of Stress on Driving the Intestinal Paracellular Permeability.” *Current Issues in Molecular Biology*, vol. 45, no. 11, 18 Nov. 2023, pp. 9284–9305, https://doi.org/10.3390/cimb45110581.

Dziedziak, Marta, et al. “Gut Microbiota in Psychiatric and Neurological Disorders: Current Insights and Therapeutic Implications.” *Biomedicines*, vol. 13, no. 9, 29 Aug. 2025, pp. 2104–2104, www.mdpi.com/2227-9059/13/9/2104, https://doi.org/10.3390/biomedicines13092104.

Foster, Jane A., et al. “Stress & the Gut-Brain Axis: Regulation by the Microbiome.” *Neurobiology of Stress*, vol. 7, no. 7, Dec. 2017, pp. 124–136, www.sciencedirect.com/science/article/pii/S2352289516300509, https://doi.org/10.1016/j.ynstr.2017.03.001.

Hamid Mostafavi Abdolmaleky, and Jin-Rong Zhou. “Gut Microbiota Dysbiosis, Oxidative Stress, Inflammation, and Epigenetic Alterations in Metabolic Diseases.” *Antioxidants*, vol. 13, no. 8, 14 Aug. 2024, pp. 985–985, www.ncbi.nlm.nih.gov/pmc/articles/PMC11351922/, https://doi.org/10.3390/antiox13080985.

Herrera-Quintana, Lourdes, et al. “Interplay of Microbiome, Oxidative Stress and Inflammation in Health and Disease.” *Antioxidants*, vol. 15, no. 2, 8 Feb. 2026, p. 222, https://doi.org/10.3390/antiox15020222. Accessed 25 Feb. 2026.

Karl, J. Philip, et al. “Effects of Psychological, Environmental and Physical Stressors on the Gut Microbiota.” *Frontiers in Microbiology*, vol. 9, no. 9, 11 Sept. 2018, https://doi.org/10.3389/fmicb.2018.02013.

Karner, Thomas, et al. “Gut Microbial Diversity and Inferred Capacity to Produce Short-Chain Fatty Acids Are Associated with Acute Stress Reactivity in Healthy Adults.” *Neurobiology of Stress*, vol. 42, 24 Mar. 2026, pp. 100807–100807, https://doi.org/10.1016/j.ynstr.2026.100807. Accessed 22 May 2026.

---. “Gut Microbial Diversity and Inferred Capacity to Produce Short-Chain Fatty Acids Are Tied to Acute Stress Reactivity in Healthy Adults.” *BioRxiv (Cold Spring Harbor Laboratory)*, 25 June 2025, https://doi.org/10.1101/2025.06.24.661294. Accessed 22 May 2026.

Kearns R. “Gut-Brain Axis and Neuroinflammation: The Role of Gut Permeability and the Kynurenine Pathway in Neurological Disorders.” *Cellular and Molecular Neurobiology*, vol. 44, no. 1, 2024, pubmed.ncbi.nlm.nih.gov/39377830/, https://doi.org/10.1007/s10571-024-01496-z. Accessed 9 Oct. 2024.

Kelly, John R., et al. “Breaking down the Barriers: The Gut Microbiome, Intestinal Permeability and Stress-Related Psychiatric Disorders.” *Frontiers in Cellular Neuroscience*, vol. 9, 14 Oct. 2015, p. 392, www.frontiersin.org/articles/10.3389/fncel.2015.00392/full, https://doi.org/10.3389/fncel.2015.00392.

Kurhaluk, Natalia, et al. “Role of Gut Microbiota in Modulating Oxidative Stress Induced by Environmental Factors.” *Cellular Physiology and Biochemistry*, vol. 59, no. S2, 8 Aug. 2025, pp. 2–52, https://doi.org/10.33594/000000799.

La Torre, Danique, et al. “Psychosocial Stress-Induced Intestinal Permeability in Healthy Humans: What Is the Evidence?” *Neurobiology of Stress*, vol. 27, 1 Nov. 2023, p. 100579, www.sciencedirect.com/science/article/pii/S235228952300067X, https://doi.org/10.1016/j.ynstr.2023.100579.

Li, Lei, et al. “Oxidative Stress, Inflammation, Gut Dysbiosis: What Can Polyphenols Do in Inflammatory Bowel Disease?” *Antioxidants*, vol. 12, no. 4, 20 Apr. 2023, pp. 967–967, www.ncbi.nlm.nih.gov/pmc/articles/PMC10135842/, https://doi.org/10.3390/antiox12040967.

Li, Xuzhao, et al. “Relief from Chronic Unpredictable Mild Stress-Induced Constipation by Fluoxetine and Butyrate and Their Impact on Gut Microecology.” *Drug Design, Development and Therapy*, vol. 19, 2025, pp. 11675–11695, pubmed.ncbi.nlm.nih.gov/41479595/, https://doi.org/10.2147/DDDT.S552593.

Marwaha, Komal, et al. “Exploring the Complex Relationship between Psychosocial Stress and the Gut Microbiome: Implications for Inflammation and Immune Modulation.” *Journal of Applied Physiology*, vol. 138, no. 2, 15 Jan. 2025, https://doi.org/10.1152/japplphysiol.00652.2024.

Moloney, Rachel D., et al. “Stress and the Microbiota-Gut-Brain Axis in Visceral Pain: Relevance to Irritable Bowel Syndrome.” *CNS Neuroscience & Therapeutics*, vol. 22, no. 2, 10 Dec. 2015, pp. 102–117, https://doi.org/10.1111/cns.12490.

Petakh, Pavlo, et al. “Gut Microbiota and Immune Response under Chronic Social Stress in an LPS-Induced Inflammatory Model.” *Stress*, vol. 28, no. 1, 5 Dec. 2025, https://doi.org/10.1080/10253890.2025.2595777.

Rosell-Cardona, Cristina, et al. “Acute Stress-Induced Alterations in Short-Chain Fatty Acids: Implications for the Intestinal and Blood Brain Barriers.” *Brain, Behavior, & Immunity - Health*, 17 Apr. 2025, p. 100992, www.sciencedirect.com/science/article/pii/S266635462500050X, https://doi.org/10.1016/j.bbih.2025.100992.

Rusch, Jody A, et al. “Signalling Cognition: The Gut Microbiota and Hypothalamic-Pituitary-Adrenal Axis.” *Frontiers in Endocrinology*, vol. 14, 19 June

2023, pmc.ncbi.nlm.nih.gov/articles/PMC10316519/, https://doi.org/10.3389/fendo.2023.1130689.

Sharma, Seema, and Sampat Singh Tanwar. “Gut–Brain Axis in Neurological Disorders: Mechanistic Links and Regulatory Role of Chinese Herbal Medicines.” *Journal of Holistic Integrative Pharmacy*, vol. 7, no. 1, Mar. 2026, pp. 66–90, https://doi.org/10.1016/j.jhip.2026.02.004. Accessed 12 Apr. 2026.

Sun, Yiqi, et al. “The Role of Gut Microbiota in Intestinal Disease: From an Oxidative Stress Perspective.” *Frontiers in Microbiology*, vol. 15, 14 Feb. 2024, https://doi.org/10.3389/fmicb.2024.1328324.

Tetel, M. J., et al. “Steroids, Stress and the Gut Microbiome-Brain Axis.” *Journal of Neuroendocrinology*, vol. 30, no. 2, Feb. 2018, p. e12548, https://doi.org/10.1111/jne.12548.

van de Wouw, Marcel, et al. “Short-Chain Fatty Acids: Microbial Metabolites That Alleviate Stress-Induced Brain-Gut Axis Alterations.” *The Journal of Physiology*, vol. 596, no. 20, 28 Aug. 2018, pp. 4923–4944, physoc.onlinelibrary.wiley.com/doi/abs/10.1113/JP276431, https://doi.org/10.1113/jp276431.

Verma, Ankita, et al. “Gut-Brain Axis: Role of Microbiome, Metabolomics, Hormones, and Stress in Mental Health Disorders.” *Cells*, vol. 13, no. 17, 27 Aug. 2024, pp. 1436–1436, www.mdpi.com/2073-4409/13/17/1436, https://doi.org/10.3390/cells13171436.

Wang, Zhen, et al. “The Gut Microbiota Metabolite Butyrate Modulates Acute Stress-Induced Ferroptosis in the Prefrontal Cortex via the Gut–Brain Axis.” *International Journal of Molecular Sciences*, vol. 26, no. 4, 17 Feb. 2025, pp. 1698–1698, www.proquest.com/docview/3171027091?sourcetype=Scholarly%20Journals, https://doi.org/10.3390/ijms26041698.

Wu, Changmei, et al. “Gut Microbiota: A New Target for the Prevention and Treatment of Insomnia Using Chinese Herbal Medicines and Their Active Components.” *Frontiers in Pharmacology*, vol. 16, 6 May 2025, https://doi.org/10.3389/fphar.2025.1572007. Accessed 6 May 2026.

Ying, Li, et al. “Chronic Stress Is Associated with Altered Gut Microbiota Profile and Relevant Metabolites in Adolescents.” *BMC Microbiology*, vol. 25, no. 1, 8 July 2025, https://doi.org/10.1186/s12866-025-04094-1.

Zhang, Jindong, et al. “Beneficial Effect of Butyrate‐Producing Lachnospiraceae on Stress‐Induced Visceral Hypersensitivity in Rats.” *Journal of Gastroenterology and Hepatology*, vol. 34, no. 8, 1 Aug. 2019, pp. 1368–1376, www.ncbi.nlm.nih.gov/pmc/articles/PMC7379616/, https://doi.org/10.1111/jgh.14536. Accessed 25 Mar. 2021.

Zheng, Gen, et al. “Chronic Stress and Intestinal Barrier Dysfunction: Glucocorticoid Receptor and Transcription Repressor HES1 Regulate Tight Junction Protein Claudin-1 Promoter.” *Scientific Reports*, vol. 7, no. 1, 3 July 2017, p. 4502,

www.nature.com/articles/s41598-017-04755-w, https://doi.org/10.1038/s41598-017-04755-w.

Zong, Ye, et al. “Chronic Stress and Intestinal Permeability: Lubiprostone Regulates Glucocorticoid Receptor‐Mediated Changes in Colon Epithelial Tight Junction Proteins, Barrier Function, and Visceral Pain in the Rodent and Human.” *Neurogastroenterology & Motility*, vol. 31, no. 2, 4 Oct. 2018, p. e13477, https://doi.org/10.1111/nmo.13477. Accessed 18 Mar. 2021.

A Narrative Study on Legal Consciousness of Chinese Domestic Workers – From the Perspective of Workers' Lived Experiences By Wenxin Yu

Introduction and Literature Review

Underrepresented domestic workers and legal consciousness in the gig economy

As hundreds of thousands of domestic workers have entered platform-based labour markets in China (Li, 2017), current academic literature has not sufficiently examined how these workers understand and make sense of law through their everyday experiences. Domestic workers, who provide care and household services in private homes, have long drawn scholarly attention because their labour lies at the intersection of gender, migration, household status (hereafter *hukou*), and labour precarity. In China, domestic workers are predominantly migrant women in urban areas who were historically excluded from education and stable formal employment after the transition away from the planned economy (Luo, 2019). Existing scholarship has therefore focused on how their social identities contribute to enduring inequality, including precarious working conditions (Hu, 2010), low and unstable income (Wang & Wu, 2017), migrant status (Li, 2017), gendered violence (Chen, 2011; Tong, 2018), and struggles over legal rights (Peng, 2019; Luo, 2023).

Within legal studies, Chinese domestic workers are often understood as a socially marginalised population with limited access to formal legal knowledge and weak institutional protection. Because many are recruited and employed through informal channels, they are especially vulnerable to rights violations and work-related difficulties (Peng, 2014; Su & McDonnell, 2022). Prior studies have shown that Chinese migrant workers are often hesitant or resistant to taking legal action when their labour rights are violated, leading some scholars to argue that state institutions should improve regulatory practice and raise workers' legal awareness (Chen, 2011). For example, Wong (2011) suggests that domestic workers' difficulty in attaining legal rights is partly shaped by insufficient rights consciousness and a lack of personal strategies for pursuing protection. Their limited understanding of rights is tied not only to long-standing structural barriers, but also to the instability and informality of their work (Wong, 2011; Lu, 2017). Similarly, Ellerman argues that Chinese domestic workers' silence in the face of legal violations reflects both limited legal consciousness and the multiple subordinations attached to being domestic workers, low-income migrants in urban areas, and women (2017). Together, these studies suggest that legal consciousness is not an abstract body of knowledge, but it is shaped by social position and everyday experience. However, it remains unclear whether earlier understandings of domestic workers' legal consciousness still apply in the context of the contemporary gig economy.

The Shifting Conception of Legal Consciousness

Legal consciousness refers to how ordinary people perceive, interpret, and engage with law in everyday life. Rather than focusing only on legal institutions, this concept is used to understand the social power of legality in daily experience (Headworth, 2024). Existing scholarship generally falls into two broad approaches. One examines how state law governs

behaviour from the top down, while the other focuses on how individuals interpret, negotiate, and mobilise legality from the bottom up in relation to their own lives (Marshall, 2006; Hull, 2016). Some scholars argue that legal consciousness research should reveal how state law operates in action by shaping people's goals, choices, and understandings of social problems (Marshall & Barclay, 2003; Miežanskienė, 2020). More recent scholarship, however, emphasises the relational and interpretive nature of legal consciousness, arguing that people actively make sense of law through the conditions of their lived experience (Nguyen, 2017; Saguy & Stuart, 2008; Lavelanet, 2024). In this view, legal consciousness is not limited to knowing official rules; it involves how individuals interpret law, assign meaning to it, and respond to it through concrete events and relationships. Chua and Engel (2019), for example, argue that research on legal consciousness should focus on how people understand and respond to life experiences in relation to the law, especially through shifting perceptions shaped by specific incidents over time.

This shift is important because it expands the study of law beyond formal legality. Researchers increasingly examine how ordinary people interact not only with state-made law but also with rules, norms, and expectations that function as law in practice and shape ordinary people's legal consciousness. Lavelanet (2024) suggests that legal consciousness must be understood relationally, through the connection between what people live and what they experience as law. In this sense, legal consciousness is situational: what people understand as "law" may depend on the social structures, institutional absences, and power relations shaping daily life. What people experience as law may therefore include not only formal legal protections, but also the rules they come to see as normal, binding, or unavoidable. Young (2014), for instance, shows how people may become obedient to private "house law" when formal legal systems do not adequately protect them, and how such informal rules can be absorbed into broader understandings of legality. Neilson (2000), therefore, calls for methodological attention to the contexts in which different groups interpret and experience what they understand as law.

This perspective is especially relevant to domestic workers. Studies across different contexts have shown that domestic workers' legal awareness is shaped not only by official rights discourse but also by workplace subordination, social norms, and informal rules governing everyday conduct (Hayashi, 2010; Güdük & Desmet, 2011; Rakotomalala, 2025). However, relevant research in the Chinese context remains limited, particularly regarding workers' own storied understandings of legality. As shown, existing scholarship often identifies structural barriers, but less often captures how domestic workers themselves narrate and interpret experiences of protection, vulnerability, obligation, and injustice. This limitation becomes even more significant in the context of platform-mediated labour, where workers' encounters with law may be increasingly indirect and filtered through algorithms, ratings, private employers, and platform rules. To understand domestic workers' legal consciousness in contemporary China, it is therefore necessary not only to identify structural constraints but also to explore how workers make meaning of these experiences in their own accounts.

The Creation of "Gig Workers" and Their Legal Consciousness

Research on workers' rights and legal consciousness in China has long been tied to wider historical and socioeconomic transformations, including changes in state policies and labour law (Gallagher, 2006). This is also true for domestic workers (Wong, 2011). One major structural factor is the *hukou* system, which has historically regulated migration and determined rural workers' access to urban employment and welfare (Windrow & Guha, 2005). After the collapse of the planned economy and the shift from state-owned enterprises to private industry, millions of workers entered informal labour sectors. By the 1990s, an estimated 45 million workers had been laid off or pushed into precarious work, helping form the labour pool that later sustained China's expanding capitalist economy (Xu, 2022; Liang et al., 2016). Many of these workers, especially rural migrants in cities, lost the protections once provided by traditional work units and were instead absorbed into unstable forms of informal employment (Xue et al., 2014).

These structural conditions have persisted into the era of the gig economy. With the rise of large digital platforms and the decline of the formal sector around 2015, China's gig economy expanded rapidly, reaching approximately 200 million workers by 2020 (Xiao & Liu, 2023). However, this growth is not matched by equivalent labour protections, highlighting the gap between rapid economic development and labour standards, especially workers' welfare and social security (Cooke et al., 2019). Gig workers in China face "multiple decent work deficits," including low pay, weak representation, overtime work, and exposure to health and safety risks (Xu & Liu, 2021). These vulnerabilities are not evenly distributed. More than half of gig workers are non-local *hukou* residents (Xu & Liu, 2021), and many enter platform labour because they have limited access to stable urban employment due to their education level and migration status (Yu, 2018). The Precarity Framework (PF) further explains how low levels of education can exclude individuals from the formal labour market and narrow their employment choices (Li et al., 2025). In this sense, the gig economy's "low entry threshold" makes it particularly attractive to workers who need immediate income but do not qualify for formal urban jobs, especially rural migrants (Yu, 2018). However, this accessibility is inseparable from constraint. Workers often accept platform labour not because it offers genuine freedom, but because they lack viable alternatives (Yu, 2018; Li et al., 2025).

Research further shows that gig workers often lack formal legal protection because their labour is classified as informal or as independent contracting. Platforms typically do not provide comprehensive social insurance or compensation, leading to low social security enrollment rates among workers (Wu et al., 2025; Zhang & Liu, 2025). Voluntary social security enrollment has been implemented to alleviate. However, Zhang and Liu argue that voluntarism can only serve as a "supplement" to mandatory coverage and, in practice, reaches only a limited portion of workers (2025). Similar patterns have been documented elsewhere: in Hangzhou, 63% of platform workers lacked social security coverage (Chen et al., 2022), a condition that conflicts with international labour standards concerning social protection and fundamental rights at work (Cooke et al., 2019). At the same time, gig workers do not necessarily enjoy high levels of autonomy. Their earnings and hours depend heavily on platform algorithms and opaque work-allocation systems. Because workers often do not know how tasks are distributed, rated, or

penalised, platforms exercise control through one-sided rules that remain largely hidden from those subject to them. Refusing tasks may result in punishment or reduced access to future opportunities. This algorithmic management deepens uncertainty and pushes workers toward physical and psychological exhaustion in order to avoid sanctions or blacklisting (Xie et al., 2025).

These findings suggest that the Chinese gig economy is structurally disembedded from formal labour protections (Chen et al., 2020). However, most existing research on platform control and worker precarity focuses on more visible on-demand sectors such as ride-hailing and food delivery. Much less attention has been paid to the domestic service sector, even though it is predominantly composed of female migrant workers. This further investigation matters because domestic work is not carried out in public spaces, but in private homes and intimate labour relationships. As a result, domestic workers may encounter not only platform rules, but also household expectations, employer discretion, and informal norms that regulate conduct in ways that may feel "law-like". For a study of legal consciousness, this means domestic workers may be especially likely to interpret legality through lived and narrated experience rather than through direct encounters with formal law.

Existing studies on Chinese gig workers help elucidate this issue, but remain limited. Scholars have shown that, although gig workers are often aware of inadequate labour protection, they rarely challenge the broader structure of platform work. High economic dependency and intense competition push workers to prioritise immediate earnings and job retention over collective claims-making (Xu & Liu, 2021). Workers may regard platforms as too powerful to confront and may see themselves as easily replaceable. As a result, platform rules and regulations become central to their everyday understanding of work and constraints. Qian et al. (2024), for example, examine how food delivery workers perceive algorithmic management and platform rules, finding growing mistrust toward the platforms and resistance to their daily encounter with "management," "rules," and "policies." Ou (2025) similarly argues that gig workers' legal awareness is often grounded in "rules" rather than "individual rights." Instead of challenging the system directly, workers commonly engage in "everyday resistance" by exploiting algorithmic loopholes to protect their own interests. However, these studies focus primarily on workers' perceptions of platforms rather than on how workers narrate and interpret the meaning of legality across the broader arc of their lives and work experiences.

Currently, only one study directly addresses Chinese gig workers' legal consciousness, focusing on food delivery workers. Lin finds that workers perceive legality as "something to be avoided…dangerous to invoke" because of their ambiguous legal status and constrained socioeconomic position (2025: 22). This is a significant contribution because it shows that legal consciousness among gig workers is relational, situational, and shaped by daily encounters rather than by abstract legal doctrine alone. However, Lin's study remains limited in two important ways. First, it focuses on food delivery workers, a predominantly male group (Zheng et al., 2026). Second, it does not fully explore how gendered and sector-specific lived experiences shape legal consciousness as workers interpret events over time. This matters because female gig

workers face distinctive forms of precarity, including widening wage gaps (Yeasmin, 2024a; Han et al., 2024), increasing unfair treatment (Zhou, 2024), and gendered precarity reflected in their everyday experiences (Kwan, 2022; Liu & Huang, 2025). Domestic workers occupy a particularly important position at the intersection of gender, migration, and gig labour, yet the topic of domestic workers in the platform era remains underexplored (Yin, 2022; Chang & Huang, 2023).

Taken together, the literature suggests that domestic workers' legal consciousness is likely shaped not only by formal legal exclusion but also by platform governance, informal workplace rules, gendered labour relations, and migrant precarity. However, current scholarship has not sufficiently explained how Chinese domestic workers themselves make sense of these conditions through the stories they tell about work, conflict, vulnerability, protection, and survival. Most existing studies either focus on domestic workers outside the platform context or examine male-dominated gig sectors such as food delivery. As a result, the lived and narrated dimensions of legal consciousness among female migrant domestic workers remain underrepresented in current research. Therefore, this study seeks to explore the following research questions:

- How do Chinese domestic workers narrate their lived experiences in platform-based domestic labour?
- How do those narratives reveal their legal consciousness in relation to rules, rights, and regulations?

Methodology and Method

Data Collection

This study explores how Chinese female domestic workers in the gig economy understand law, rules, and regulations through lived experience. Because the research focuses on meaning-making rather than measuring attitudes or behaviours, narrative inquiry was used as the methodological approach.

Narrative inquiry is appropriate because it examines how individuals interpret events and construct meaning through stories over time. It aligns with legal consciousness research, which understands legality as grounded in everyday experience rather than only formal legal knowledge. Semi-structured narrative interviews were conducted to allow participants to describe work experiences in their own terms while enabling follow-up questions about platform rules, injury, social security, and unfair treatment (Peterson-Badali & Abramovitch, 1992; Barton, 2004; Berrey et al., 2012; Guzik, 2018; Horák et al., 2021).

This study used narrative interviews to allow participants to recount their lived experiences as domestic gig workers navigating platform-based labour. Although the research focus guided the interviews, they were conducted flexibly and conversationally, allowing participants to narrate meaningful events in their own terms. This approach reduced excessive researcher control while still allowing follow-up questions to clarify details and explore issues relevant to the study, such as platform rules, injury, social security, and experiences of unfair

treatment (Muller & Segal, 2015; Bihu, 2020). Rather than treating legal consciousness as a set of predefined opinions, this method enabled examination of how participants' understandings of the law emerged through stories about work, conflict, and everyday experience.

Participant Recruitment and Sampling Strategy

Purposive sampling was used to recruit four female domestic workers who worked through domestic service platforms, held rural *hukou*, and lived in urban areas. These criteria reflected the study's focus on the intersection of gender, migration, and platform labour. The small sample size was appropriate to the goals of narrative inquiry, which prioritises depth, detail, and interpretive richness over statistical generalizability (Adhikari, 2021). Narrative research seeks to understand the complexity and uniqueness of participants' lived experiences, and this requires sufficient time and attention to each participant's account. A small participant pool also supported the development of trust between researcher and participant, which is especially important when discussing sensitive issues such as work-related vulnerability, platform discipline, and experiences of rights violations (Goldstein, 2017). Ethical protocol is attached in the appendix.

Interview Protocol

All interviews were conducted remotely via Tencent Meeting. Although remote interviewing may reduce access to some real-time emotional and embodied cues due to the physical distance between the researcher and the participant (Engwand, Goldspink & Wood, 2022), it was appropriate for this study for two reasons. First, a non-face-to-face setting may help participants feel safer and more comfortable when discussing sensitive personal experiences (Jenner & Myers, 2019). Second, online interviewing allowed participants to join from locations that were more convenient and familiar to them, which may have supported greater openness in storytelling (Johnson, Scheitle & Ecklund, 2021).

All interviews were conducted in Chinese, the first language of both the researcher and the participants, because participants had limited English proficiency. Conducting interviews in participants' first language helped ensure that they could express experiences, emotions, and interpretations more fully and precisely, especially when recounting complex or sensitive events (Cortazzi et al., 2011). Since this study focused on the meanings participants attached to their experiences, linguistic comfort and expressive accuracy were essential.

Data Analysis

Audio recordings were transcribed for analysis, and the researcher manually checked the accuracy of each transcript. Because this study focused on how legal consciousness was expressed through lived experience, the analysis aimed to preserve the narrative quality of participants' accounts while also identifying cross-narration patterns relevant to the research question.

This study employed a story constellation analytic approach. Story constellation is useful for examining individual stories in relation to wider social and structural conditions, which allows the researcher to connect personal experience with broader contexts (Hancock et al., 1998). This approach was appropriate because the study did not treat participants' experiences as

isolated incidents. Instead, it sought to understand how individual narratives reflected the broader conditions shaping domestic workers' legal consciousness in the gig economy.

The analysis followed Craig's (2007) multidimensional framework of story constellation, which builds on Clandinin and Connelly's (2000) conceptualisation of narrative analysis through three interrelated stages: broadening, burrowing, and restorying.

Broadening: the researcher identified the wider social, institutional, and structural factors that appeared in participants' narratives. These included factors such as *hukou* status, gender, platform governance, employer expectations, and access to protection or compensation. The purpose of this stage was to situate each participant's story within the broader social world that shapes her work experiences.

Burrowing: the analysis moved inward to examine how each participant personally experienced and interpreted those conditions. At this stage, the researcher focused on the meanings participants attached to specific incidents, the language they used to describe rules and rights, and the ways they connected personal background to their understanding of work and legality. This stage was especially important because legal consciousness in this study was understood not as formal legal knowledge alone, but as participants' lived interpretations of regulation, obligation, and vulnerability.

Restorying: the researcher reconstructed participants' narratives analytically by tracing how beliefs, identities, and experiences were interconnected across their stories. The purpose of restorying was not to replace participants' voices but to present their narratives coherently, highlighting both individual meaning-making and shared patterns across cases. This final stage made it possible to show how legal consciousness was shaped through recurring experiences of platform control, informal work rules, dependence, and constraint.

Each transcript was read multiple times, coded for recurring themes, and compared across participants to identify shared patterns while preserving individual differences. This process allowed the study to examine how legal consciousness emerged through lived experience rather than formal legal knowledge alone.

Findings and Discussion

This study sought to understand how Chinese domestic workers in the gig economy narrate their lived experiences and how these narratives reveal their legal consciousness regarding rules, rights, and regulations. The charts below show individual narrations with exemplary quotes:

THEMES	SUB-THEME	CONTEXT	EXAMPLE NARRATIVES
PLATFORM PENALTIES AS SOVEREIGN MANDATES	The Fining System as an Unquestionable Legal Code	• The participant views the platform's 50 RMB fine for order cancellations—even those caused by clients—as an absolute rule rather than a potentially illegal labour deduction.	*""Generally, it is guest reasons [for cancellation] ... we are fined 50 yuan per order." (1a)*
	Accepting "Clearing Out" as a Legitimate Termination	• Because she believes the lack of a "labour contract" voids all protections, she accepts the platform's power to "clear her out" for taking leave as a natural consequence.	*"If you take too much leave, they will just gradually clear us out... we didn't sign a labour contract." (1* *"Producing negative reviews... will affect our order receiving... [we] need to go back for retraining." (1* "Do you have ways to report these bad comments? *It's hard to withdraw a bad comment*" (1b)
MISUNDERSTANDING EMPLOYMENT STATUS	**The "Individual Household" Illusion**	• She reflects on her status as an "individual business owner" as a standard market reality, unaware that this classification is often a legal strategy used by platforms to avoid social security obligations.	*"Now they changed it to a point system... you not only have to do cleaning, but you also have to do sales... get customers to buy packages" (1c)*
LEGAL ENGAGEMENT AND PRACTICAL REDRESS MECHANISMS	Avoidance of formal Legal Action	• In her reflection on platform inconveniences. She chooses to "exit" the platform rather than directly challenging it.	*"After the rules changed, we cannot take a leave... the salary is around five to six thousand yuan a month, which is pretty low for living in Guangzhou. So, I just left." (1d)*
MISUNDERSTANDING EMPLOYMENT STATUS	**The "Individual Household" Illusion**	• She reflects on her status as an "individual business owner" (*getihu*) as a standard market reality, unaware that this classification is often a legal strategy used by platforms to avoid social security obligations.	*"I understood a bit... we aren't signing a labour contract... It's like we are individual business owners (getihu)." (1f)*

THEMES	SUB-THEME	CONTEXT	EXAMPLE NARRATIVES
THE DISSONANCE BETWEEN OFFICIAL PROTECTION AND LIVED REALITY	Internalization of Non-Labor Status	Participant views her "third-party status" as the reason for her low economic return.	*"Because we are third-party cooperation, what gets to our hands is just a tiny bit." (2a*
INDIVIDUAL SELF-RELIANCE	**Rejecting Systemic Support for Self-Sufficiency**	Rather than seeking to reform platform rules during the "off-season," she believes a worker's only path is to "rely on yourself," even if they were working for a formal company.	*"Even if you work in a company, you still have to rely on yourself."(2b)*
INDIVIDUAL SELF-RELIANCE	**Internalizing the "Condition" Barrier**	She accepts a denied injury claim based on the platform's vague "conditions," lacking the legal knowledge to recognize that workplace injuries are protected by statutory laws regardless of platform policy.	*"I had a foot injury... but the company didn't compensate. Their reason was I didn't 'meet the conditions' ... I spent over 8,000 yuan out-of-pocket." (2c)*
INTERNALIZED POWERLESSNESS AND "SPEECHLESSNESS"	Resignation toward Systemic Unfairness	Participants reflect on the speechlessness of the algorithmic rules, feeling powerless to argue against platform decisions	*"If the rules are unreasonable, we can't say anything... we are speechless. The aunties just have to work like this." (2d)*

THEMES	SUB-THEME	CONTEXT	EXAMPLE NARRATIVES
MISUNDERSTANDING THE SOURCE OF RIGHTS	**Freedom as an Exemption from Legal Protection**	She mistakenly believes that obtaining social security or a labour contract would *legally require* her to give up her flexibility, viewing platform precarity as a mandatory "trade-off."	*"Currently there isn't a company that provides [insurance] and gives you this kind of freedom. You have to make a trade-off." (3a*
LAW AS A DISTANT, NON-APPLICABLE TOOL	**High Monetary Threshold for Justice**	Her legal consciousness defines "law" only as a tool for "big money" disputes. She chooses to "absorb" daily platform injustices to avoid the "trouble" of a system she perceives as exhausting.	*"Only for big things involving a lot of money would people go to court... for small problems... we don't want to cause so much trouble." (3b*
LEGAL ENGAGEMENT AND PRACTICAL REDRESS MECHANISMS	**Helplessness Against Algorithmic Bias**	She views an unfair customer rating as a problem with no external legal or organizational solution.	*"The platform judges the bad review... I don't know what to do" (3c)*

THEMES	SUB-THEME	CONTEXT	EXAMPLE NARRATIVES
SELF-IMPROVEMENT AS A SOLUTION TO STRUCTURAL INEQUALITY	**The "Inspirational Aunties Framework**	She restoried her physical pain and lack of protection into a narrative of personal resilience, believing that her professional well-being depends on her "positive energy" rather than legal rights.	*"My right hand is completely numb... but the next day you work and it's fine... I think I am an inspirational (lizhi) auntie." (4a)*
SELF-IMPROVEMENT AS A SOLUTION TO STRUCTURAL INEQUALITY	**Overcoming Psychological Barriers through Individual Efforts**	She reflects that workers must "force themselves" to overcome the "low class" feeling of the work, framing success as a mental change rather than a challenge to platform conditions.	*"People will feel... 'lower than others'... but you must take that step to solve problems... you aren't even afraid of being poor, so why care about 'face'?" (4b)*
RELATIONAL SENTIMENT AS A SUBSTITUTE FOR LEGAL FRAMEWORKS	**Emotional Legitimacy over Contractual Rights**	Her legal consciousness is entirely bypassed by a focus on "family feeling." She values being "appreciated" over having a formal contract.	*"They are like family... As they say, not relatives, but better than relatives." (4c)*

First Constellation: State Law as Distant and Weakly Actionable

The first constellation that emerged from the participants' narratives concerns their relationship to formal state law. In their view, state law appeared distant, partial, and weakly actionable in relation to their actual work experiences, as participants did not seem to understand the precise role it plays in their daily lives.

State law appeared distant in the way participants spoke about contracts and legal classification. One participant explained that workers were not "*signing a labour contract*", and another accepted her status as a "*third-party cooperator*", treating it as a normal feature of the job rather than as a questionable legal arrangement. They do not perceive it as a violation of their legal rights or employment laws. In these accounts, formal labour law was not narrated as something workers could readily claim. Instead, it remained outside the practical logic through which they made sense of work.

The distance from state law is also evident in participants' lack of legal knowledge. For example, participant 2's account adds details for further understanding. When asked about her perceptions of unreasonable regulations, she referred to "*high taxes*" and argued that platforms should not consider her low earnings "*within the scope of tax deductions"*. Her narrative blurred the distinction between government deductions and platform-imposed financial loss. This suggests that state regulation was experienced indirectly rather than as a coherent legal framework. As a result, the role of state-made law becomes distant and inaccessible in participants' daily experiences. Participants were aware of state law, but experienced it as distant and weakly actionable.

Second Constellation: Platform Rules as Lived Law

A careful reading shows that, in participants' experiences, state law seemed distant, whereas platform rules seemed immediate. The second constellation that emerged from the data, therefore, concerns platform rules as the most visible and binding form of law in workers' daily lives. These narratives reveal that participants often understood regulation less through formal rights discourse and more through the practical authority of deductions, rankings, penalties, and platform judgments.

The 'lived laws' on the platform affect participants' responses to emergent issues. One participant explained that when the platform deducted wages, "*we just have to work like this*", adding that even when workers asked for an explanation, "*they have their own sort of reason*". In this case, platform rules remain unchallengeable. A third participant recounted paying 8000 yuan out of pocket for a foot injury because it "*did not meet platform conditions*". Nevertheless, she did not challenge the platform further or seek legal aid due to the lack of understanding of her legal rights. The 'lived law' holds power over participants' experience.

Instead of questioning the fairness of algorithmic management, workers either accepted the "clearing out" process and internalised the issue or left the platform silently (2025). Workers experienced platform governance as practical legality (Young, 2014). The rules that mattered most in everyday life were the ones that determined income, punishment, visibility, and continued access to work. This supports scholarship conceptualising legal consciousness as extending beyond formal state law to the rules and norms people experience as binding in practice (Young, 2014; Lavelanet, 2024). It also connects with Lin's argument that platform workers may perceive legality as something dangerous or difficult to invoke, and with Ou's suggestion that platform decisions function as "law in practice".

This pattern appears in domestic platform work and is tied to platform mechanisms used to avoid punishment. One participant explained that if workers took too much personal leave, "*they will just 'clean' us out*", meaning the algorithm would assign them fewer jobs and effectively push them off the platform—the lack of visibility on the platform results in lower incomes. Rather than questioning whether the review or punishment was fair, she focused on how to survive within the system instead of seeking legal assistance. This change in the platform system reflects previous scholars' findings that although gig workers are perceived as "independent workers", the platform still exerts control like an employer, usually through the point system and performance indicators (Wang & Cooke, 2021). Repeated encounters with platform authority shaped legal consciousness, making platform rules more immediate and meaningful than formal state law.

Third Constellation: Work Negotiation, Trade-offs, and Flexibility

The third constellation that emerged from the narratives concerns how workers themselves negotiated precarious labour conditions through trade-offs over flexibility. This section shows that the current state of legal consciousness shapes individuals' actions, as evidenced by how workers in this study interpreted the practical bargains required to sustain work and family life.

Participants negotiated platform conditions through necessary trade-offs. Rather, they narrated a series of constrained negotiations in which insecurity was often accepted because flexibility remained necessary.

This pattern was especially visible in Participant 3's story. She explained that she chose her current platform largely because it offered the flexibility she needed to balance paid domestic labour with care responsibilities for children at home. Her narrative suggests that platform work was meaningful not because it provided full protection or fairness, but because it allowed her to manage conflicting demands that formal jobs might not accommodate. In this case, flexibility was not the opposite of precarity. It was one reason precarious labour became acceptable to participants.

This supports prior scholarship that gig workers trade security for flexibility (Aloisi, 2016). Nevertheless, it also complicates that argument. In the narratives of this study, flexibility was not described solely as free worker autonomy. It was a constrained accommodation to migration status, family responsibility, and limited alternatives in individual personal lives. Participants often accepted the absence of contracts, paid leave, and social security because the platform's flexibility allowed them to continue earning while managing reproductive labour and economic pressures.

This negotiation shaped legal consciousness in important ways. When flexibility was experienced as necessary for survival, the lack of formal legal protection could come to seem like part of the arrangement rather than a violation worth challenging. Workers did not always narrate this as fairness, but often as reality. As one participant's account suggests, workers recognised that being on the platform meant giving up traditional employee protections.

However, they accepted this because the work remained compatible with their lives in ways that other jobs were not. In this sense, legality was interpreted through compromises necessary to keep working.

Fourth Constellation: Negotiation and Self-responsibilization

The fourth constellation concerns how participants responded to unfairness, exclusion, or harm. In relation to the research question, the findings show that workers' narratives often revealed legal consciousness through the coping strategies they adopted: silence, endurance, exit, and self-responsibilization. Workers prioritised survival over formal legal mobilisation.

Workers often viewed voice as ineffective and exit as more realistic (Hirschman, 1970). One participant explicitly stated, "*Individuals cannot fight with these decisions*". This helps explain why legal consciousness in the study did not often appear as rights-based mobilisation. Workers recognised unfairness, but they narrated the costs of challenging it as too high relative to the likely outcome.

Participant 2's narrative shows a similar pattern of endurance, but with a stronger tendency toward internalised resignation. She framed work largely through the language of self-reliance, saying that "*even if you work in a company, you still have to rely on yourself*". Her account suggests that insecurity was not only experienced as externally imposed, but also internalised as part of the worker's own burden to manage.

Participant 4 provides the clearest example of self-responsibilisation. When describing how her right hand became "*completely numb*" from labour, she dismissed the need for further concern and said, "*the next day you work, and it is fine…*" In her narrative, the central problem was not the absence of labour protection but the worker's own attitude toward hardship. This reflects what scholars describe as neoliberal self-responsibilization and the creation of personalised work identities under precarious conditions (Ou, 2025; Petriglieri et al., 2019).

Legal consciousness appeared in the ways workers justified endurance, narrated resignation, or converted structural unfairness into a personal challenge. In this study, participants' legal consciousness was revealed not only in what they said about rules and regulations, but also in how they positioned themselves in relation to those forces. Silence, endurance, and self-responsibilization were signs of how legality had already been absorbed into everyday strategies of survival.

Discussion: Constellation Synthesis

Taken together, these four constellations answer the research question by showing that Chinese domestic workers in the gig economy narrate their lived experiences through a layered and unequal understanding of legality. Across the narratives, state law appeared as formally present but practically distant: participants rarely described it as an accessible or effective resource for addressing routine work problems. By contrast, platform rules were experienced as immediate, enforceable, and consequential in daily life. As a result, legal consciousness was shaped less through direct engagement with formal rights and more through workers' repeated encounters with deductions, rankings, leave penalties, injury exclusion, and platform judgment.

The four constellations also show that participants' legal consciousness was not formed only in the present moment of platform work. Although workers did not always explicitly name hukou, migrant status, gender, or occupational marginality, these structural conditions remained embedded in their backgrounds. This historically secluded situation reflects the "anthropological disembeddedness" described by scholars, in which individuals become atomised and disconnected from social institutions, leading to greater acceptance of exploitation (Chen et al., 2020). Disembedded from the formal employment scheme, domestic platform workers remain in a 'second-tier category' where social rights of a citizen are denied (Demurger et al., 2009). Their narratives suggest that entry into platform domestic labour was already shaped by unequal access to formal employment, limited social protection, and the need to balance paid labour with care responsibilities. In this sense, workers' legal consciousness emerged not simply from isolated workplace disputes, but from a broader life context in which legal protection had long been partial, unstable, or difficult to claim.

This finding helps explain why participants often narrated platform rules as their most meaningful 'lived laws.' When state-law protection is experienced as distant or fragmented, platform governance becomes the most visible form of authority in everyday work. As scholars such as Lin (2025) and Gallagher (2006) argue, workers consider wages and remuneration the sole criteria for determining legitimacy. Workers, therefore, came to interpret legality in terms of

what immediately affected their earnings, visibility, and continued access to jobs. Their narratives suggest that legal consciousness in the gig economy is shaped not by the absence of law altogether, but by an uneven hierarchy in which formal state law is weakly actionable while platform regulation is directly lived.

At the same time, the constellations show that workers did not respond uniformly to this legal environment. Instead, they negotiated it through constrained trade-offs, especially around flexibility, caregiving, and income survival. This is important because it complicates any reading of participants as simply passive or legally unaware. Their narratives show that in the lack of formal legal knowledge, workers transform platform unfairnesses into personal choices or inadequacy, revealing forms of practical reasoning: accepting insecurity to preserve flexibility, remaining silent because the challenge seemed ineffective, leaving the platform when voice was useless, or interpreting hardship as an individual responsibility to overcome. These responses suggest that legal consciousness was expressed not only through explicit talk about rights, but also through the strategies workers used to endure and make sense of unequal regulation.

These findings extend the literature in four ways. First, they support a scholarship that understands legal consciousness as relational and grounded in lived experience rather than in formal legal knowledge alone. Second, they deepen studies of platform labour by showing that for domestic workers, platform rule and state law are not simply separate systems; they are experienced through an uneven relationship in which platform authority dominates everyday life while formal protection remains inaccessible. Third, they provide evidence of how domestic workers, when formal laws are distant, rationalise unjust platform regulations through shifting them to personal choices and responsibilities. Fourth, they demonstrate the value of narrative inquiry in this context. By examining how workers tell stories about migration, care, injury, discipline, and survival, this study reveals how legal consciousness is constructed over time through the accumulation of lived experiences rather than through abstract legal reasoning alone.

Conclusion

This study shows that female domestic gig workers in China experience legality through platform rules more than formal law. State law appeared distant and weakly actionable, while platform decisions became their ‘lived law’ which shaped income, punishment, and continued access to work. Workers often interpreted insecurity as a necessary trade-off for flexibility and survival, leading to silence, endurance, and self-responsibilization rather than formal legal mobilisation.

Several limitations should be noted. First, participants were recruited through platforms in Guangzhou, which may not apply to other urban cities due to regulatory differences. Second, interviews are conducted via online meetings, which inherently limit access to certain real-time physical cues that a researcher would normally capture in person. Third, this study did not consider different family backgrounds or earlier traumatic events that may potentially shape individuals’ responses to challenges. Finally, as narrative inquiry often prioritises interpretive depth over statistical generalisability, it lacks generalisability.

However, this article holds several implications for future research. Further research could focus on identifying the specific reasons behind workers' resistance to seeking legal aid by assigning a point system to factors such as "lack of access" or "stigmas," and quantitatively measuring workers' helplessness. It also suggests the need for corporate compliance with labour protection. More binding policies should be implemented to ensure workers' enrollment in social security and other social schemes, especially for migrant workers who cannot afford urban social security on their own.

Works Cited

Adhikari, S. P. (2021). Revealing the story of an individual through narrative inquiry: A methodological review. *Interdisciplinary Research in Education,* 6(1), 71–80.

Barton, E. (2004). The construction of legal consciousness in discourse: Rule and relational orientations toward the law in a disability support group. *Journal of Pragmatics*, 36(4), 603–632. https://doi.org/10.1016/S0378-2166(03)00098-5

Berrey, E., Hoffman, S., & Nielsen, L. B. (2012). Situated justice: A contextual analysis of fairness and inequality in employment discrimination litigation. *Law & Society Review*, 46(1), 1–36. https://doi.org/10.1111/j.1540-5893.2012.00471.x

Bihu, R. (2020). Using unstructured interviews in educational and social science research: The process, opportunity and difficulty. *Global Scientific Journals*, GSJ, 8(10).

Chang, C., & Huang, W. (2023). The Rising Gig Economy in China: Implications for the Protection of Migrant Workers. In *Protecting the Future of Work: New Institutional Arrangements for Safeguarding Labour Standards* (pp. 97–110). emerald insight. https://doi.org/10.1108/978-1-80071-248-520221013

Chen, T., Song, W., Song, J., Ren, Y., Dong, Y., Yang, J., & Zhang, S. (2022). Measuring well-being of migrant gig workers: Exemplified by Hangzhou City in China. *Behavioral Sciences*, 12(10), 365. https://doi.org/10.3390/bs12100365

Clandinin, D. J., & Connelly, F. M. (2000). *Narrative Inquiry: Experience and story in qualitative research*. San Francisco: Jossey-Bass.

Cooke, F. L., Xu, J., & Bian, H. (2019). The prospect of decent work, decent industrial relations and decent social relations in China: Towards a multi-level and multi-disciplinary approach. *The International Journal of Human Resource Management,* 30(1), 122–155. https://doi.org/10.1080/09585192.2018.1521461

Craig, C. J. (2007). Story constellations: A narrative approach to contextualising teachers' knowledge of school reform. *Teaching and Teacher Education*, 23(2), 173–188. https://doi.org/10.1016/j.tate.2006.04.014

Ellerman, M. L. (2017). The power of everyday subordination: Exploring the silencing and disempowerment of Chinese migrant domestic workers. *Critical Asian Studies*, 49(2), 187–206.

Engward, H., Goldspink, S., Iancu, M., Kersey, T., & Wood, A. (2022). Togetherness in separation: Practical considerations for doing remote qualitative interviews ethically. *International Journal of Qualitative Methods*, 21, 16094069211073212.

Gallagher, M. E. (2006). Mobilising the law in China: "Informed disenchantment" and the development of legal consciousness. *Law & Society Review*, 40(4), 783–816.

Goldstein, S. E. (2017). Reflexivity in narrative research: Accessing meaning through the participant-researcher relationship. *Qualitative Psychology*, 4(2), 149.

Golzar, J., Noor, S., & Tajik, O. (2022). Convenience sampling. *International Journal of Education & Language Studies*, 1(2), 72–77.

Güdük, A., & Desmet, E. (2022). Legal consciousness and migration: Towards a research agenda. *International Journal of Law in Context*, 18(2), 213–228.

Guzik, K. (2008). The agencies of abuse: Intimate abusers' experience of presumptive arrest and prosecution. *Law & Society Review*, 42(1), 111–144. https://doi.org/10.1111/j.1540-5893.2008.00336.x

Hancock, B., Ockleford, E., & Windridge, K. (2001). An introduction to qualitative research. *London: Trent Focus Group*.

Han, J., Li, M., Li, S., & Hu, Y. (2024). The widening gender wage gap in the gig economy in China: The impact of digitalisation. *Humanities and Social Sciences Communications*, 11(1), 1639.

Hayashi, R. R. (2010). Empowering domestic workers through law and organising initiatives. *Seattle Journal for Social Justice*, 9, 487.

Headworth, S. (2024). Liability chains and access to justice. *Punishment & Society*, 26(4), 672–692.

Horák, F., Lacko, D., & Klocek, A. (2021). Legal consciousness: A systematic review of its conceptualisation and measurement methods. *Anuario de Psicología Jurídica*, 31, 9–34. https://doi.org/10.5093/apj2021a2

Hu, X. (2010). Paid domestic labour as precarious work in China. *Summit (Simon Fraser University*). http://summit.sfu.ca/item/9971

Hull, K. E. (2016). Legal consciousness in marginalised groups: The case of LGBT people. *Law & Social Inquiry*, 41(3), 551–572.

Jenner, B. M., & Myers, K. C. (2019). Intimacy, rapport, and exceptional disclosure: A comparison of in-person and mediated interview contexts. *International Journal of Social Research Methodology*, 22(2), 165–177.

Johnson, D. R., Scheitle, C. P., & Ecklund, E. H. (2021). Beyond the in-person interview? How interview quality varies across in-person, telephone, and Skype interviews. *Social Science Computer Review*, 39(6), 1142–1158.

Li, M., Hu, X., Jin, K., & Han, J. (2025). Exploring factors influencing entry into the gig economy: A study of Chinese workers. *Acta Psychologica*, 259, 105301. https://doi.org/10.1016/j.actpsy.2025.105301

Liang, Z., Appleton, S., & Song, L. (2016). Informal employment in China: Trends, patterns and determinants of entry (No. 10139). *IZA Discussion Papers.*

Luo, C. (2019). Domestic workers in China: A comparative study utilising the capability approach (Doctoral dissertation, University of Canterbury).

Luo, C. (2023). The capability approach: A suitable theoretical framework to improve legal protection of domestic workers in China. *International Journal of Law and Politics Studies,* 5(1), 63–69.

Lavelanet, A. F. (2024). Advancing our understanding of legal consciousness: Unpacking the lived experiences of abortion seekers (Doctoral dissertation, University of Birmingham).

Lu, L. (2017). Chinese domestic workers as inferior “other”: Why are they particularly vulnerable yet neglected? (Doctoral dissertation, University of Ottawa).

Marshall, A. M. (2006). Communities and culture: Enriching legal consciousness and legal culture. *Law & Social Inquiry*, 31(1), 229–249.

Marshall, A. M., & Barclay, S. (2003). In their own words: How ordinary people construct the legal world. *Law & Social Inquiry*, 28(3), 617–628.

Miežanskienė, R. (2020). Exploring tendencies in research on migrants’ legal consciousness and uncovering factors for socio-legal integration. *European Integration Studies*, 14(1), 26–38.

Mueller, A. E., & Segal, D. L. (2015). Structured versus semistructured versus unstructured interviews. *The Encyclopedia of Clinical Psychology,* 1–7.

Neilson, L. (2000). Partner abuse, children and statutory change: Cautionary comments on women's access to justice. *Windsor Yearbook of Access to Justice*, 18, 115.

Nguyen, T. P. (2017). Legal consciousness and workers’ resistance in Đồng Nai Province, Vietnam. *Asian Journal of Comparative Law*, 12(2), 311–333.

Ou, L. (2025). What do platform workers think about the law? The ambiguous legal status and legal consciousness of On-Demand food delivery riders in China. *Industrial Law Journal*. https://doi.org/10.1093/indlaw/dwaf025

Peng, Q. (2014). Multi-layered gaps between ILO conventions and the Chinese legal protection for migrant women workers as domestic helpers in China.

Peng, Q. (2019). No country for the low-end? An intersectional analysis of the status of migrant women workers as domestic helpers and relevant laws in China. *International Journal on Minority and Group Rights*, 26(4), 575–599.

Peterson-Badali, M., & Abramovitch, R. (1992). Children’s knowledge of the legal system: Are they competent to instruct legal counsel? *Canadian Journal of Criminology*, 34(2), 139–160.

Rakotomalala, L. V. (2025). Caring or complaining? Emotional management and performative legal consciousness among undocumented domestic workers in Belgium. *Social & Legal Studies.*

Saguy, A. C., & Stuart, F. (2008). Culture and law: Beyond a paradigm of cause and effect. *The Annals of the American Academy of Political and Social Science,* 619(1), 149–164.

Su, Z., McDonnell, D., Cheshmehzangi, A., Ahmad, J., Chen, H., Šegalo, S., & Cai, Y. (2022). What “family affair?” Domestic violence awareness in China. *Frontiers in Public Health*, 10, 795841.

Tong, X. (2018). Gendered labour regimes: On the organising of domestic workers in urban China. *Asian Journal of German and European Studies*, 3(1), 14.

Truong, K. P., Westerhof, G. J., Lamers, S. M., & de Jong, F. (2014). Towards modelling expressed emotions in oral history interviews: Using verbal and nonverbal signals to track personal narratives. *Literary and Linguistic Computing,* 29(4), 621–636.

Wang, J., & Wu, B. (2017). Domestic helpers as frontline workers in China's home-based elder care: A systematic review. *Journal of Women & Ageing,* 29(4), 294–305.
Windrow, H., & Guha, A. (2005). The hukou system, migrant workers, and state power in the People's Republic of China. *Northwestern Journal of International Human Rights*, 3, 1.
Wong, L. (2011). Chinese migrant workers: Rights attainment deficits, rights consciousness and personal strategies. *The China Quarterly*, 208, 870–892.
Wu, X., Wang, J., & Jiang, J. (2025). Insurance and the gig economy: Analysing Chinese platform workers' participation in work-related injury insurance. *Acta Psychologica*, 256, 105016.
Xiao, W., & Liu, Y. (2023). Gig economy in China? A scoping review. *Journal of Theory and Practice of Engineering Science*, 3(8), 29–36.
Xie, P., Chen, H., Cheng, Z., Fan, W., & Ge, S. (2025). How do courts respond to the platform regulation policy in protecting gig workers? Evidence from court decisions in China. *Asia Pacific Journal of Human Resources*, 63(2), e70014.
Xu, Y., & Liu, D. (2021). Decent work for the digital platform workers: A preliminary survey in Beijing. *Digital Law Journal*, 2(1), 48–63.
Xu, Z. (2022). The gig economy in China. In Routledge eBooks (pp. 392–400).
Xue, J., Gao, W., & Guo, L. (2014). Informal employment and its effect on the income distribution in urban China. *China Economic Review*, 31, 84–93.
Yeasmin, S. (2024). Uncovering gender pay disparities: An examination of Chinese gig workers' earnings and assets (Doctoral dissertation, Brac University).
Yin, S. (2024). Situating platform gig economy in the formal subsumption of reproductive labour: Transnational migrant domestic workers and the continuum of exploitation and precarity. *Capital & Class*, 48(1), 119–133.
Young, K. M. (2014). Everyone knows the game: Legal consciousness in the Hawaiian cockfight. *Law & Society Review*, 48(3), 499–530.
Yu, H. (2018). China: The destiny of web platform workers in China: Employees, nothing or a "third option"? *Japan Labour Issues*, 2(8), 92–100.
Zhang, C., Jian, Y., Zhou, Z., & Zhou, B. (2024). Can the digital economy close the gender wage gap? Evidence from China. *International Review of Economics & Finance*, 96, 103604.
Zhang, H., & Liu, K. (2025). Can a voluntarist approach to social security extend protection to gig workers? Evidence from the platform-based food-delivery sector in China. *Journal of Industrial Relations*, 67(1), 101–129.
Zheng, Q., Qiu, Z., & Yang, W. (2026). Can women escape from precarity in the gig economy: Evidence from labour market, labour process, and labour reproduction in China. *Sage Open*, 16(1), 21582440251412488.

Appendix A: Research Ethics Protocol

Ethical Procedures

Before each interview, participants were informed of the study's purpose and provided with a Participant Information Sheet and Consent Form. They were told about the research

objective, how their data would be used, the anonymity measures that would be applied, and their right to refuse to answer any question or withdraw from the study at any time without consequence. Each interview lasted approximately forty to sixty minutes and was audio-recorded with permission.

To protect confidentiality, all recordings and transcripts were stored on a passcode-protected device accessible only to the researcher. Identifying information was removed during transcription and analysis. Because the study involved workers discussing potentially vulnerable aspects of their employment, maintaining anonymity and voluntary participation were especially important.

Game Theory In Football Market Transfers By Atharva Kumar

Abstract

Football transfer markets are increasingly characterised by strategic competition, financial constraints, and rapid decision-making. Traditional economic models explain many aspects of player valuation but often fail to capture the interdependent decision-making that occurs when multiple clubs compete for the same player. This paper applies principles of game theory, particularly Nash Equilibrium and strategic payoff analysis, to examine how football clubs make transfer decisions involving release clauses. A conceptual game-theoretic model is developed to represent the strategic interaction between two competing clubs choosing whether to trigger a player's release clause or negotiate a transfer fee. The model incorporates variables including player market value, release clause amount, negotiated price, and the strategic cost of losing the player to a rival. To demonstrate the model's practical application, two historical transfer case studies, Neymar Jr.'s transfer from FC Barcelona to Paris Saint-Germain (2017) and Erling Haaland's transfer from Borussia Dortmund to Manchester City (2022), are analysed. The findings illustrate that clubs frequently prioritise long-term sporting utility, competitive positioning, and commercial value over immediate financial returns. The study also demonstrates how changes in rivalry intensity, financial asymmetry, and player importance influence optimal strategic decisions. The paper contributes to the growing literature on sports economics by extending game-theoretic analysis to football release-clause negotiations and proposes a framework that may assist future research into strategic behaviour in football transfer markets.

Keywords: Game Theory; Football Transfer Market; Nash Equilibrium; Release Clauses; Strategic Decision-Making; Sporting Utility; Football Economics; Case Study Analysis.

Introduction

What is Game Theory?

Game theory is the mathematical study of strategic decision-making. It is used to analyse situations in which the outcome for each participant depends not only on their own choices but also on the choices made by others. As a branch of applied mathematics and economics, game theory has wide-ranging applications in everyday life, helping explain how individuals and organisations make decisions when their interests are interconnected.

Importance in Football

Unlike situations in which an individual makes decisions independently based solely on personal benefits, football clubs operate in a highly competitive environment where every decision is influenced by the actions of rival clubs. During transfer negotiations, clubs must make quick strategic decisions while managing significant financial investments. Understanding these decisions, the incentives behind them, and the resulting outcomes requires the application of game theory.

How to Model a Game

To simplify the complexity of real-world football transfer markets, the situation can be reduced to three fundamental components:

1. **Players** – Who is making the decisions? In this study, the decision-makers are the football clubs involved in the transfer, referred to as Club A and Club B.
2. **Actions** – What strategic choices are available to the players?
3. **Payoffs** – What are the incentives behind each action, and what benefits or costs do the clubs receive as a result?

Throughout this paper, it is assumed that all players are rational. In other words, each club chooses the strategy that it believes will maximise its expected payoff.

What is Nash Equilibrium?

Named after the renowned mathematician John Nash, Nash Equilibrium is one of the central concepts in game theory and is widely used to analyse strategic interactions. It describes a situation in which every player has selected the best possible strategy in response to the strategies chosen by all other players. Once a game reaches Nash Equilibrium, no player has an incentive to change their decision unilaterally because doing so would not improve their outcome. In the context of this paper, identifying the Nash Equilibrium allows us to predict the likely actions of competing clubs and determine the most probable strategic outcome of a transfer negotiation.

Types of Games

Game theory generally distinguishes between two broad categories of games: **zero-sum games** and **non-zero-sum games**.

A **zero-sum game** is a strictly competitive situation in which one player's gain is exactly equal to the other player's loss, meaning that the total payoff always equals zero. An example would be two direct rival clubs competing for the same domestic title, where one club's success directly reduces the other's chances of winning.

A **non-zero-sum game**, on the other hand, allows the total payoff to either increase or decrease depending on the strategies adopted by the players. For example, if both clubs aggressively bid for a player, they may drive the transfer fee higher, resulting in negative outcomes for both parties (a negative-sum game). Conversely, if both clubs negotiate effectively, they may reduce costs and achieve mutually beneficial outcomes (a positive-sum game). Throughout this paper, we examine how differences between clubs can transform a transfer scenario into a positive-sum, zero-sum, or negative-sum game.

How Do Football Clubs Operate in the Transfer Market?

Football transfer markets differ significantly from conventional markets. While they involve buying and selling, the asset being transferred is a professional football player who possesses personal preferences and decision-making power. Furthermore, transfer markets operate within a highly regulated framework that includes financial fair play regulations, spending limits, and wage controls.

One distinctive feature of football transfers is the **release clause**, a contractual provision that allows a buying club to bypass negotiations entirely by paying a predetermined transfer fee.

Once the release clause is activated, the selling club is legally obligated to permit negotiations with the player. This paper focuses specifically on modelling strategic interactions involving release clauses through the application of game theory.

The Clubs' Main Objectives

Unlike conventional businesses whose primary objective is to maximise profit, football clubs generally seek to maximise **sporting utility**, with profitability serving as a secondary objective. Sporting utility refers to achieving sustained success on the field through winning matches, qualifying for prestigious competitions, improving squad quality, and enhancing the club's long-term reputation. A stronger reputation also makes the club more attractive to investors and elite players, thereby increasing its future chances of success.

Because football is inherently competitive, preventing a direct rival from strengthening its squad is often a major strategic objective. A club may therefore pursue a player not only because of the player's individual quality but also to prevent a rival club from acquiring them. At the same time, clubs must balance these sporting ambitions against financial and regulatory constraints, including transfer budgets, wage structures, youth development, and governing regulations. Consequently, every transfer decision represents a balance between sporting objectives and financial responsibility.

Incentives of the Players

Players are not passive assets within the transfer market; they possess their own objectives and make decisions based on a variety of personal and professional considerations. Financial incentives play a significant role in determining a player's preferred destination. These may include higher wages, signing bonuses, performance-related incentives, endorsement opportunities, and the potential to maximise long-term earnings.

Sporting considerations are equally important. Players may wish to compete in prestigious competitions such as the UEFA Champions League, secure a guaranteed starting position to maintain their place in their national team, or continue their professional development under a highly respected manager. In situations where multiple clubs trigger a player's release clause simultaneously, the final decision ultimately rests with the player, who evaluates these financial and sporting factors before choosing between competing offers.

The area of research will be how to use game theoretic models in transfer markets to maximise the utility of the club. Football is an industry that is one of the most popular in the world, and based on the trajectory, its not slowing down anytime soon. In professional football, a very important part of it is the transfer market. Within this, there are many factors and clubs try to find ways to maximise their utility by taking into factors various things. As a result, research on this topic has increased due to the steady rise in its popularity. Understanding this topic is important as it provides insight into the thought processes behind a transfer that a person may normally overlook.

The global football industry has transformed the transfer market into a highly complex economic game, where elite clubs constantly seek strategic ways to maximize their utility. One of the key aspects to understand is the factors affecting a player's market value. Several studies

have examined the possible determinants of individual player valuation and why some clubs may be willing to pay more or less for a player than their rivals. While Carmichael and Thomas (1993) state that selling clubs possess a higher degree of bargaining power, the literature argues that market values are determined by a wide range of interconnected factors. For example, variables such as stadium attendance (Carmichael & Thomas, 1993), historic league rankings (Dobson et al., 2000), and financial metrics (Frick, 2007) can be used to measure player values. These determinants can be divided into (1) player characteristics, including demographics and physical attributes, (2) player performance, in terms of on-pitch performance, and (3) player popularity through external sources such as social media or news (Lennart Ante, 2019).

How Do The Clubs Make Decisions?

Traditional research usually viewed transfers as standard market transactions. However, recent studies have increasingly labelled the market as a non-cooperative game. As Zhou (2025) notes, "the buying club also needs to consider the existence of other competitors during the bidding process and thus may face considerable pressure." This emphasises that transfer choices are completely interdependent, and a club's final payoff depends on what its closest rivals do.

Gap In Research

However, despite the existing research, there is limited focus on how release clauses, a legal loophole, function in these kinds of scenarios. Most studies also focus on large, interconnected transfers, while limited research has been conducted on scenarios involving strictly two clubs using game-theoretic models to maximise utility.

Therefore, this study aims to address this gap by examining how game theory can be used to maximise the utility of one club when competing directly against another club. The variables in this study include the rivalry between the teams (domestic or international), the relative financial power of the clubs, and the extent to which each club needs the player based on the position. For example, one team may require a striker more urgently than another.

Methodology

This study adopts a qualitative conceptual research design based on game-theoretic modelling and comparative case study analysis. Rather than collecting primary data through surveys or interviews, the research develops a theoretical framework using established concepts from game theory and applies the framework to real football transfer cases.

The research was conducted in three stages.

First, a comprehensive review of academic literature relating to football economics, player valuation, transfer markets, bargaining theory, and game theory was undertaken. Existing studies were used to identify the major variables influencing transfer decisions, including market value, release clauses, bargaining power, financial resources, and sporting utility.

Second, these findings were synthesised to construct a two-player non-cooperative game representing competition between two clubs pursuing the same player. The model defines the players, available strategies, payoff variables, and assumptions regarding rational

decision-making. Nash Equilibrium is then used to evaluate the strategic outcomes under different market conditions.

Finally, the model was applied to two real-world football transfers: Neymar Jr.'s transfer to Paris Saint-Germain in 2017 and Erling Haaland's transfer to Manchester City in 2022. These cases were selected because they represent contrasting strategic environments. The Neymar transfer demonstrates an asymmetric market with minimal competition due to Paris Saint-Germain's exceptional financial strength, while the Haaland transfer illustrates direct strategic competition between multiple elite clubs. Publicly available information, including historical market values, transfer fees, release clauses, and documented club circumstances, was used to estimate the variables within the payoff matrices.

The analysis compares the predicted outcomes generated by the game-theoretic model with the actual decisions made by the clubs. This allows evaluation of how effectively game theory explains strategic behaviour in football transfer markets and identifies factors beyond purely financial considerations that influence decision-making.

Literature Review

Football transfer markets are considerably more complex than conventional markets. They represent highly strategic environments in which a club's success depends not only on its own decisions but also on accurately anticipating the actions of competing clubs.
The modern transfer system was formally established in 2002 when the European Commission and FIFA introduced reforms governing player contracts and international transfers across Europe (European Commission, 2002). Under these regulations, players remain contractually bound to their clubs until the expiry of their contracts unless an agreement is reached between the clubs or specific contractual provisions apply (Hoey et al., 2021).

One such provision is the release clause, which enables a buying club to bypass lengthy negotiations by paying a predetermined transfer fee. As Hoey et al. (2021) explain, release clauses allow rival clubs to acquire a player without requiring the selling club's consent, provided that the contractual amount is paid in full. Since release clauses require clubs to make immediate, high-stakes decisions under fixed pricing conditions, traditional supply-and-demand models alone are insufficient for explaining these interactions.

Instead, these strategic situations are more appropriately analysed using game theory, which provides a framework for understanding how clubs anticipate rival behaviour, maximise sporting utility, and optimise transfer decisions under competitive conditions.
To develop a game-theoretic model of football transfer markets, it is first necessary to define the key concepts used throughout this study.

Players:-

The players within this model include Club A, the selling club that owns the player under contract, and Clubs B and C, the competing clubs seeking to acquire the player. Unlike traditional markets, however, the footballer is also an independent decision-maker with personal

objectives and preferences. In addition, player agents represent another important stakeholder, as they seek to maximise both their own interests and those of the players they represent.

Actions (Strategies):-

The competing clubs each have two strategic options, chosen without complete knowledge of the opposing club's decision:

Trigger (T): Pay the player's release clause immediately, bypassing negotiations and attempting to secure the transfer directly.

Negotiate (N): Decline to activate the release clause and instead attempt to negotiate a lower transfer fee with the selling club.

Payoffs:-

In game theory, payoffs represent the expected utility gained from a particular strategic decision. Within football transfer markets, however, utility extends beyond simple financial profit.

The payoff associated with signing a player may include increased sporting success, qualification for major competitions, squad improvement, commercial growth, fan engagement, and long-term club development. As Feess et al. (2003) explain, football transfer payoffs can be understood as the allocation of surplus created through bargaining, rather than as purely monetary outcomes.

Consequently, this study adopts a broader interpretation of utility that incorporates both sporting and commercial benefits.

Market Type:-

Football transfer markets involving release clauses are best characterised as non-cooperative games. In many situations, one club's success in securing a player directly reduces another club's opportunity to strengthen its squad, creating highly competitive strategic interactions. Nevertheless, these markets are not always strictly zero-sum because multiple stakeholders—including the buying club, selling club, player, and agent—may all derive different forms of utility from the same transfer.

CASE STUDY 1 – NEYMAR JR. (2017)

Neymar Jr.'s transfer to PSG from FC Barcelona in 2017 was the most expensive transfer in the world at that time. A whopping €222 million was paid to secure the 25-year-old Brazilian's signature on a five-year contract. We will be using this case study as a model to determine the application of game theory. This case study was chosen because it represents a scenario where there are no rivals and demonstrates the asymmetric financial power model with no competing club.

Fitting this data into the variables:

- **V** – €100m (Market value of Neymar Jr. at the time of the transfer)
- **R** – €222m (Release clause set by FC Barcelona)
- **P** – N/A (FC Barcelona flatly refused negotiations for Neymar Jr. as they wanted to keep him.)

To calculate the defeat cost, we need to answer the following three questions:

- **How urgently did the buying club need that position and signing?**
  There was a desperate need to sign Neymar Jr., as there was a significant need for a left winger. PSG also needed a global superstar to become the face of the club following Zlatan Ibrahimović's departure.
- **Were they direct rivals competing for the same trophies?**
  Although they were not domestic rivals, they were ultimately competing for prestigious European trophies such as the UEFA Champions League.
- **What was the reputational cost of losing out?**
  It was huge, as the Qatari investors had recently acquired the club and needed to make statement signings. Additionally, the previous year, PSG had crashed out of the Champions League after losing 6–1 in FC Barcelona's historic comeback, spearheaded by Neymar himself, who scored two goals and provided one assist in that match. This transfer was a statement signing to announce PSG as a global footballing superpower by signing the joint most expensive player at that time.

After considering all these factors, the defeat cost of losing Neymar to a rival would be around €70m. However, the cost of simply not signing him would be lower, approximately €40m.

Given the absence of a financially credible rival, this bilateral game collapses into a unilateral game, where the actions of **Club B** become irrelevant simply because there is no Club B. This reduces the matrix to only two relevant outcomes:

1. PSG triggers (T)
2. PSG does not trigger

   Therefore, the relevant payoff in this case depends solely on PSG's actions.

To calculate the defeat cost to FC Barcelona (**L**), this will also be estimated by answering the following questions:

- **Market value at that time** – €100m
- **Importance to the club in the footballing sense** – Huge. Neymar was an irreplaceable part of the famous MSN trio, which had recently secured a historic treble. He was also the main creative force alongside Messi.
- **Importance to the club outside the football ground** – Neymar was an up-and-coming superstar who was expected to become the next global football icon. He brought significant popularity to the club, as evidenced by the €150m increase in revenue experienced by PSG after signing Neymar Jr.

After considering these factors, the defeat cost to FC Barcelona of losing Neymar would be approximately €160m.

Putting all these variables into the payoff matrix

	PSG Triggers	**PSG Doesn't Trigger**

PSG Payoff	€100m − €222m = **−€122m**	**−€40m**
FC Barcelona Payoff	€222m − €160m = **€62m**	**€160m**

What this case study tells us

In pure sporting utility terms, the dominant strategy for PSG was not to trigger the release clause, as **−€40m > −€122m**. However, PSG still chose to trigger it. This suggests that the true value (**V**) of Neymar Jr. was greater than €100m, as the original definition of **V** does not account for the non-sporting utility gained from signing Neymar, such as brand value, commercial revenue, and global recognition. These additional benefits increase the true value of **V**, making **(Trigger, Accept)** the dominant strategy.

Similarly, it would have made more sense for FC Barcelona to keep Neymar Jr. However, the release clause legally removed this choice, essentially forcing the club to accept the transfer. This represents a non-zero-sum asymmetric game, as PSG's loss is not exactly equal to FC Barcelona's gain.

CASE STUDY 2 – ERLING HAALAND TO MANCHESTER CITY

Erling Haaland's transfer from Borussia Dortmund to Manchester City in 2022 was a landmark transfer. Haaland signed for City on a five-year contract. At the time, he was one of the young, up-and-coming superstars of world football, and Manchester City, after parting ways with many of their key players, wanted to kickstart a new era with Haaland leading the line.

Fitting the data into the variables:

- **V** – €150m (Market value of Haaland at the time of the transfer)
- **R** – €60m (Release clause set by Borussia Dortmund)
- **P** – N/A (There were no negotiations, as the release clause was significantly lower than the market value.)

To calculate the defeat cost, we again need to answer the following three questions:

- **How urgently did the buying club need that position and signing?**
  Very urgently, as Sergio Agüero, their legendary striker, had just left the club, and they had no striker capable of adequately filling that role.
- **Were they direct rivals competing for the same trophies?**
  Similar to the previous case, they were not domestic rivals but were competing for the same European trophies.
- **What was the reputational cost of losing out?**
  It was significant because, a year earlier, Manchester City had failed to sign Harry Kane, which would have once again forced the club to operate with a false nine system. Furthermore, losing Haaland would have significantly strengthened Real Madrid, who were viewed as Manchester City's main European rivals.

After considering all these factors, the defeat cost of losing Haaland to Real Madrid would be around €80m, primarily due to how low the release clause was relative to his actual market value.

Unlike the previous case, there were several clubs interested in signing Haaland, including Manchester City, Real Madrid, and Bayern Munich. In this model, **Club A** is Manchester City, who eventually signed Haaland, while **Club B** is Real Madrid, who were considered the closest rivals.

For Manchester City, as mentioned earlier, the defeat cost would be **€80m**.
For Real Madrid, the defeat cost would be lower, as Karim Benzema, their starting striker, was still at the club, and they were also simultaneously pursuing Kylian Mbappé. Therefore, the defeat cost would be approximately **€25m**.

Putting these values into the matrix

MCI / RMA	**Trigger (T)**	**Negotiate (N)**
Trigger (T)	€5, €32.5	€90, −€25
Negotiate (N)	−€80, €90	*N/A*

*The (N, N) scenario is structurally unavailable in this case, as Dortmund, aware that the release clause was set significantly below the market value, had no incentive to accept a negotiated fee lower than **R**. Therefore, **P = N/A**, and the (N, N) cell is not possible.*
For the case **(Trigger, Trigger)**, the payoffs would be:

- **Manchester City** – (€150 − €60 − €80)/2 = **€5**
- **Real Madrid** – (€150 − €60 − €25)/2 = **€32.5**

For the case **(Trigger, Negotiate)**, the payoffs would be:

- **Manchester City** – €150 − €60 = **€90**
- **Real Madrid** – **−€25**

For the case **(Negotiate, Trigger)**, the payoffs would be:

- **Manchester City** – **−€80**
- **Real Madrid** – €150 − €60 = **€90**

What this means?

This shows that **Trigger** was the dominant strategy for both clubs. However, Real Madrid ultimately chose **Negotiate** by prioritising the Mbappé deal, a rational but ultimately incorrect expectation that left them with neither player.

The Nash Equilibrium of this game was **(T, N)**—Manchester City triggers while Real Madrid negotiates. Although the predicted equilibrium was **(T, T)**, Real Madrid's simultaneous pursuit of Mbappé changed their effective payoff structure, which ultimately benefited Manchester City.

This illustrates a key extension of the model: when a club simultaneously pursues multiple transfer targets, its effective **D** value for any single player is reduced, making **Negotiate** appear to be a rational strategy even when **Trigger** is objectively dominant. Madrid's reduced value of **D**, which was €25m, combined with their expectation of signing Mbappé, pushed them towards **Negotiate**, ultimately resulting in the worst possible outcome of signing neither player.

Now analysing Dortmund's defeat cost (L)
This will again be estimated by answering the following three questions:

1. **Market value at that time** – €150m.
2. **Importance to the club in the footballing sense** – Huge. Erling Haaland was an exciting young talent who contributed significantly to Borussia Dortmund's performances in both the Bundesliga and the UEFA Champions League. However, Dortmund were aware that Haaland's relatively low release clause meant that he would eventually leave for a bigger club when the opportunity arose.
3. **Importance to the club outside the football ground** – Significant. Signing Haaland enhanced Borussia Dortmund's reputation as a club capable of attracting elite young talent, despite competition from several other major clubs.

While signing Haaland, Borussia Dortmund knew that they would eventually have to let him leave, as evidenced by his relatively low release clause. However, this was not a major concern because Dortmund's entire sporting model is built around signing talented young players at relatively low prices and later selling them for substantial profits. They signed Haaland for approximately €20m and later sold him for €60m, making it a planned and profitable transfer. After considering all these factors, Borussia Dortmund's defeat cost (**L**) would be approximately **€25m**, as the club had already anticipated that Haaland would eventually leave.

Dortmund's payoff

	MCI Triggers	**MCI Negotiates**
BVB's Payoff	R − L = €60 − €25 = **€35**	**0**

Therefore, this also represents a non-zero-sum game with different **D** values.

Discussion

The objective of this study was to investigate whether game theory can effectively explain strategic decision-making in football transfer markets, particularly when release clauses are involved. The results from both case studies demonstrate that transfer decisions extend beyond simple financial calculations and are largely influenced by strategic interactions between competing clubs.

The baseline model successfully illustrates how competing clubs adjust their strategies according to expected actions of their rivals. In situations where two clubs possess similar financial strength and equal sporting need, triggering a release clause often emerges as the dominant strategy because delaying negotiations increases the risk of losing the player. This outcome closely aligns with the principles of Nash Equilibrium, where each club chooses the strategy that maximises its expected utility given the anticipated behaviour of the opposing club.

The Neymar transfer illustrates the limitations of purely monetary valuation. Although the immediate payoff appears negative when considering only transfer cost and market value, Paris Saint-Germain still pursued the transfer because the player's commercial value, global

brand appeal, sponsorship opportunities, and long-term sporting impact substantially increased the club's overall utility. This suggests that football clubs optimise broader organisational objectives rather than purely financial returns.

The Haaland case demonstrates how competition between multiple clubs affects strategic choices. While the model predicts that triggering the release clause is the dominant strategy for both Manchester City and Real Madrid, Real Madrid's simultaneous pursuit of another elite player altered its perceived opportunity cost. Consequently, strategic priorities outside the model influenced the final outcome. This highlights that real-world football transfers often involve multiple interconnected games occurring simultaneously.

Overall, the findings indicate that game theory provides a strong framework for explaining football transfer behaviour, but practical decision-making also incorporates uncertainty, managerial preferences, commercial considerations, player ambitions, and incomplete information that cannot always be captured within simplified mathematical models.

Recommendations

Based on the findings of this study, several recommendations can be made.

First, football clubs should incorporate formal game-theoretic analysis into their transfer planning to better anticipate the likely actions of competing clubs before entering negotiations.

Second, transfer decision-making should evaluate player utility using broader performance indicators that include commercial value, branding potential, fan engagement, and long-term sporting contribution rather than relying solely on estimated market value.

Third, clubs should develop scenario-based strategic models for different competitive environments, including domestic rivalries, European competition, and financially asymmetric markets, as each environment produces different equilibrium outcomes.

Finally, clubs should continuously monitor competing clubs' financial situations, squad requirements, and transfer priorities since these variables significantly influence strategic behaviour and ultimately determine the success of transfer negotiations.

Future Research

This study presents a simplified two-player game involving release clauses. Future research could expand the model in several ways.

One avenue would be to develop multiplayer game-theoretic models involving several competing clubs, player agents, and the selling club simultaneously. Such models would better reflect the complexity of modern football transfer markets.

Future studies could also incorporate behavioural game theory by considering imperfect information, bounded rationality, media influence, and psychological factors affecting decision-making.

Another promising direction would be the use of repeated-game models that examine transfer relationships between clubs across multiple transfer windows rather than analysing individual transactions in isolation.

Finally, future researchers could validate the proposed model using quantitative methods, such as analysing large datasets of historical transfers to statistically test whether predicted game-theoretic equilibria correspond to actual club behaviour across different leagues and competitive contexts.

Works Cited

Carmichael, F., & Thomas, D. (1993). Bargaining in the transfer market: Theory and evidence. *Applied Economics, 25*(12), 1467–1476.

Dobson, S., & Goddard, J. (2000). *The economics of football*. Cambridge University Press.

European Commission. (2002). *Commission closes investigations into FIFA regulations on international football transfers*. https://ec.europa.eu

Feess, E., Frick, B., & Muehlheusser, G. (2003). Legal restrictions on buyout fees in professional football. *European Economic Review.*

Frick, B. (2007). The football players' labor market: Empirical evidence from the major European leagues. *Scottish Journal of Political Economy, 54*(3), 422–446.

Hoey, L., et al. (2021). Football player contracts and release clauses: Legal and economic perspectives. *International Sports Law Journal.*

Lennart Ante. (2019). Determinants of football players' market values. *Journal of Sports Economics.*

Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics, 54*(2), 286–295.

Osborne, M. J. (2004). *An introduction to game theory*. Oxford University Press.

Transfermarkt. (2025). Player market values and transfer history. https://www.transfermarkt.com

Zhou, Y. (2025). Strategic decision-making in football transfer markets: A game theoretic perspective. *Journal of Sports Economics.*