|CURIEUX|

ACADEMIC REVIEW

November Issue

Issue 13 Part 3

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: The Feasibility of Stratospheric Aerosol Injection (SAI) as an Engineering Solution to Climate Change By Dan Barakat

Page 17: Pepper was Power: How Kanara's Small Kingdoms Negotiated With the Portuguese Empire in the Seventeenth Century By Arjun Shetty

Page 24: From Data to Decisions: Transforming Industrial Emissions Data into Consumer-Friendly Climate Information By Minal Matai

Page 41: The "Buy Now, Pay Later" (BNPL) Debt Trap: A Secondary Research Analysis of Consumer Credit Risk, Behavioural Drivers, and Regulatory Responses By Jasroop Singh Narang

Page 56: Macrophage Plasticity and ECM Remodelling and Solutions in Fibrosis By Sakthivel Heamish Kanha

Page 85: An Analysis of Amyloid, Tau, and Inflammation in Alzheimer's Disease By Shaunak Sen

Page 100: Hydrogel Bioinks: Addressing Cell Density, Viability, and Resolution Trilemma in 3D Bioprinting By Beliz Kasacı

Page 112: The Role of TRP Channels in Alzheimer's Disease Pathogenesis and Potential Therapeutic Implications By Seoyoon Chung

Page 122: Formation Flight Applied to Commercial Aviation: A Review of the Aerodynamic, Operational, and Safety Constraints By Giacomo Guenzi

Page 133: The Impact of Simultaneous Animated Movie Consumption on High School Students' Academic Recollection By Isabella Neill

Page 146: The Baseline of One: A Safety-First Framework for Longitudinal Smartphone Monitoring of the Diabetic Foot By Eshaan Naik[1] and David G. Armstrong, DPM, MD, PhD[2]

Page 168: Sex Specific Implications of Insular Cortex Degradation During Alzheimer's Disease on Interoceptive Accuracy in Decision-Making By Deepika Srilalita Manikonda

Page 180: Differential Inflammatory and Cytotoxic Responses to Nanoparticles in HMDM and THP-1 Macrophages By Nikhil Rajagopalan

Page 197: Examining the Passage and Legacy of the Civil Rights Act of 1964 By Parvati Kaushik

Page 204: Providing a Suitable Explanation For The Abandonment of Fatehpur Sikri By Krishiv Loiya

Page 215: PROK2 and FGFR1 Mutations Responsible for HPG Disruption leading to Kallmann Syndrome in Women By Ashni Bhupathiraju

Page 229: The Effect of Wing Aspect Ratio on Aerodynamic Efficiency of Regional Jets, Narrow-Body, and Wide-Body Commercial Aircraft By Guanchen Wang

Page 238: The Ethics of Artificial Intelligence in Cybersecurity: Toward a Framework for Responsible Integration By Tahir Ibayev, Mentor By Emily Sheetz

Page 246: Crafting Catalonia: Tradition and Reform in the Architecture of Casa Batlló By Kasin Stone Obarski

Page 260: Investigating the Relationship Between Personality Traits, Social Media Use, and Mental Health Outcomes: A Literature Review By Pak Ying Shannon Tsang

Page 275: Consumer Behaviour and Policy Preferences in response to Rising Fuel Prices: Evidence from an Online Survey in India By Danish Pereira Kalia

Page 295: Why Quantum Computing is Not Powerful Enough to Break AES By Seaena Kim

Page 302: Street and Hidden Homelessness in Japan - A Study of Economic Pathways, Mutual Aid, and Policy Recognition By Yudai Tanaka

Page 316: Reduced-Order Modeling of Nano-Aluminum Additives in Ethanol/LOX Rocket Propellants By John Li

Page 328: Perceptions of Learning Difficulties Across Cultures and Generations: A Cross Sectional Survey in Singapore and Indonesia By Michael Vincent Junaidi

Page 345: Personalisation, Judicial Capture, and the Politics of Gender: Democratic Backsliding under Trump and Erdoğan By Sophia Fox

Page 364: Methods in Ground-Effect Aerodynamics: A Chronological Review By Terumasa Iriyama

Page 386: A Directional-Statistics Study of the Normalized Zero Spacings of the Riemann Zeta Function By Carlo Marville Harefa

Page 400: Engineering Challenges Faced in the Implementation of Atomic Clocks for Autonomous Deep Space Navigation By Nomeda Matonyte[1] [*]

Page 411: Democracy Under Pressure: Institutional Decline in the United States and Hungary By Alex Numbers

Page 425: The Rise of Content Creation among Gen Z: the Attractions and the Precarity By Ethan Y. Gu (student)[1*] , Kanishka Elupala (mentor)[2]

Page 435: Tooth Jewels and Enamel Health: Cosmetic Decoration or a Microbiofilm Trap? By Kareena Bakhshi

Page 450: Calibrating Trust in Human-Like AI: Social Cues, Detection Limits, and Open Neurocognitive Questions By Kevin Li Human Chimerism: Advances in Genetic Testing for Diagnosis By Maria Victoria Cardelli De All

Page 464: An Assessment of LSTM and XGBoost's Comparability in Financial Prediction By Vuong Trac Truong

Page 473: Food Insecurity's Impact on Hypotension and its Effects on Dementia By Ronin Back

Page 484: Hues of Attraction: Color Cohesion, Gendered Cues, and Scarcity Framing in Early-Stage Fashion Brand Engagement on Instagram By Juana Colaroff

Page 497: Evaluating the Impact of Prompt Engineering Strategies on Large Language Model Financial Forecasting By Hugo Shek (student)[1], Plínio Zanini (mentor)[2], Siddharth Krishnan (mentor)[2]

Page 517: Traditional and Emerging Technologies for Glucose Management in Type 2 Diabetes By Siddharth Tirumalai

Page 527: How Effective Is the Introduction of Regulatory Sandboxes In Helping UK Businesses Grow in the Fintech Market By Veer S. Kooner (Student)[1] , Heitor G. C. Santos (mentor)[2]

Page 539: A PLS-SEM Based Analysis of Determinants Influencing Portfolio Construction and Risk Minimization Strategies By Atiksh Suneja

Page 557: An Evolutionary Perspective on Human Susceptibility to Alzheimer's Disease By Hannah Kim

Page 567: Urban Design and Human Behavior: To What Extent Does Urban Design Affect Human Behavior? By Ana-Maria Vulturu

Page 588: Non-Familial Transmission of Rice Cultivation Knowledge in Aizu-Wakamatsu, Fukushima By Rikako Umezawa

Page 601: Retail Investors vs Institutional Investors - The Gamestop Saga By Alexander Yan

The Feasibility of Stratospheric Aerosol Injection (SAI) as an Engineering Solution to Climate Change By Dan Barakat

Abstract

With the continuous increase in greenhouse gas (GHG) emissions worldwide leading to the catastrophic consequences of climate change, researchers needed a rapid solution to global warming. Therefore, Stratospheric Aerosol Injection (SAI), a subfield of Solar Radiation Management (SRM), was proposed by researchers to temporarily reduce temperatures by reflecting incoming light rays back out of the atmosphere. This review studies the feasibility of SAI through the examination of current deployment technologies, the effects on atmospheric chemistry, optimal aerosol material selection and size distribution, as well as ethical considerations. Previous research investigates the atmospheric effects of SAI through simulations, or by comparison to natural occurrences, such as the Mount Pinatubo explosion of 1991. This paper determines whether or not SAI can currently be deployed as a realistic solution to climate change. This paper finds that despite the relatively simple challenges faced when dealing with the short-term technological and economical feasibility of SAI, deployment in the near future can not be considered without further research on environmental impacts and ethical considerations.

Introduction

Since the beginning of the Industrial Revolution, global temperatures have seen a gradual increase due to the rise in greenhouse gas (GHG) emissions and what is known as the greenhouse effect. Global temperatures have risen to about 1.2 degrees Celsius, meaning we are drawing nearer and nearer to the U.N. target of a 2.0-degree limit (Ritchie and Roser). If we go over and stay over the 2.0-degree limit, we risk triggering irreversible critical climate tipping points that further contribute to the warming of the earth, which in most cases becomes catastrophic (McKay et al.). It has been shown that simply the reduction of emissions is not guaranteed to be enough to avoid these scenarios, so researchers proposed a different approach: Climate Intervention, otherwise known as geoengineering (Royal Society).

Climate Intervention aims to artificially control and mitigate climate change at a large scale through the utilisation of different methods (Mwende). While these methods usually fall under two categories, Carbon Dioxide Removal (CDR) and Solar Radiation Management (SRM) (Mwende), this paper will be focusing mainly on the feasibility of Stratospheric Aerosol Injection (SAI), a subfield of SRM which focuses on using the deliberate injection of aerosols into the stratosphere to reflect sunlight back out of the atmosphere (Crutzen).

The deployment of SAI could lead to temperature decreases of great magnitude, depending of course on the amount deployed (Bednarz et al.). SAI can clearly be demonstrated by volcanic eruptions, where volcanoes release sulfur aerosols into the stratosphere similar to how we would with SAI (Robock).

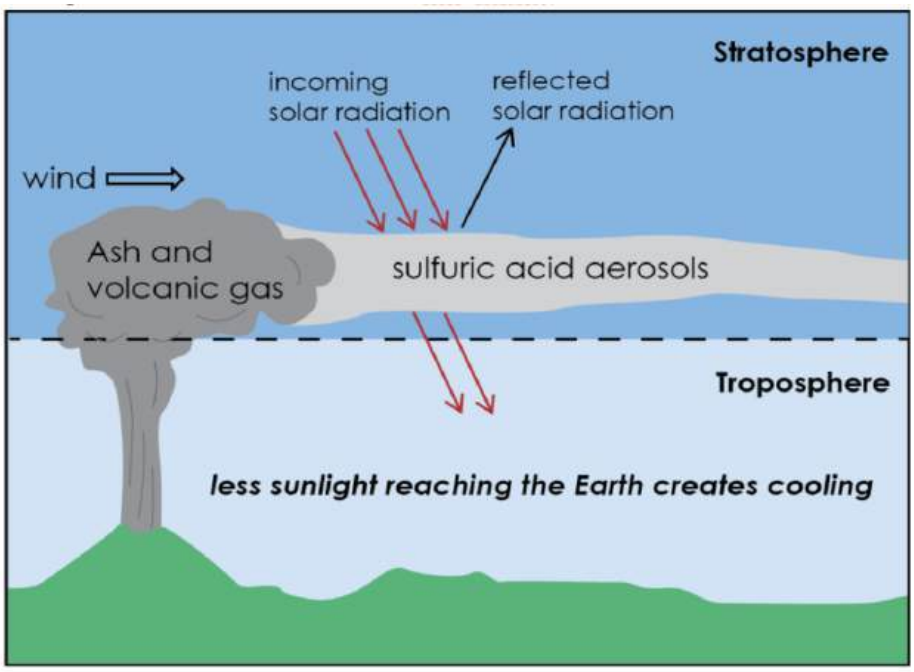

Figure 1. Volcanic sulfur dioxide forms stratospheric aerosols that reflect incoming solar radiation, which demonstrates the natural process that SAI is based off of. Source: Lindsay Iredale (2024), CC BY 4.0; adapted from USGS, public domain.

The above figure shows how the eruption of Mount Pinatubo in 1991 caused a decrease in temperatures (by around 0.5 K), which prompted further research to be done on SAI (Soden et al.).

But of course, there exist multiple challenges, across multiple various fields, that come with the deployment of SAI. If we choose to ignore the ecological, economical, and geopolitical impacts of SAI, we would certainly face catastrophic consequences. Although there exist multiple papers covering the feasibility of SAI in one focused field, the overall feasibility is scarcely studied. This paper aims to judge the overall feasibility of the deployment of SAI, and highlight changes that must be made before global-scale deployment can be considered.

Technological and Economical Feasibility

SAI deployment at a global scale needs to be both technologically and economically feasible, with one major reason being a concept known as Termination shock. Termination shock refers to what happens to temperatures after the sudden halt of SAI deployment. It usually leads to temperatures rising faster and higher than if we had otherwise never implemented SAI, and can result in catastrophe (Jones et al.). To avoid this scenario, we must have deployment technology that is reliable and relatively inexpensive, so that we can continuously inject aerosols into the stratosphere. In other words, it is imperative that we possess SAI deployment technology that is both technologically and economically feasible. This section reviews the technological and economical feasibility of both small-scale and large-scale deployment methods.

SAI seems to be technologically feasible at a small scale because we already have aircraft which can reach high altitudes. However, there is a stark contrast between throwing a kilogram of cheap sulfates off of a weather balloon and continuously releasing tonnes of optimized aerosols around the globe every year. Although existing aircraft such as simple weather balloons could be suitable for small scale regional experimentation, their limited payload capacity would make it almost impossible to scale up globally for continuous deployment. In addition, none of the existing deployment technologies can (or can be modified to) reliably inject aerosols at an altitude of 20 km (Smith and Wagner; McClellan et al.), which is commonly considered to be the baseline for SAI deployment research (Smith et al., "Review of Possible very High Altitude").

The problem grows as the scale of deployment and expected magnitude of cooling increase. Older research by McClellan et al. suggested that modified versions of existing aircraft can be feasible for SAI deployment. Later, however, it was determined by Smith and Wagner that existing deployment technologies would not be enough for global scale development, so they proposed Stratospheric Aerosol Injection Lofter (SAIL). SAIL is a hypothetical deployment technology with a disproportionately large wing to fuselage ratio, designed to be able to deploy 25 tonnes of aerosols at an altitude of 20 km (Smith and Wagner). Under their hypothetical scenario, Smith and Wagner found that approximately 4000 flights would be required in the first year (2033 in their hypothetical scenario), with that number increasing by around an additional 4000 every subsequent year (until 2047 in their hypothetical scenario). For these flights they expect 95 new aircraft to have been built by 2047, which proves that for large scale deployment to be technologically feasible, we would need large amounts of dedicated technology (SAIL) rather than simply modifying existing commercial aircraft.

Additional technological challenges arise if we want to increase the altitude at which SAI is deployed (Smith et al., "Review of Very High Altitude"). It was found that if aerosols were deployed at 25 km instead of 20 km, we could inject a lower amount of aerosols (Smith et al. "Review of Very High Altitude"). Since a lower aerosol mass would directly correlate with lower costs, it would be intuitive to assume that 25 km deployment would be a more economically efficient method. However, Smith et al. concluded that costs would actually be substantially higher due to the requirement of more complex deployment technology and higher operational risk. This showcases an economic trade-off: we could use cheaper technology to deploy at lower altitudes, but we would then need a heavier aerosol payload; or we could carry less aerosol with more advanced technology. Further research must be done in terms of optimizing deployment altitude.

Although it was identified by Smith and Wagner that the aircraft we possess is a technological limitation itself, we could mitigate those limitations by changing the payload carried by the aircraft. In "Production of Sulfates on Board", Smith et al. suggested that we carry sulfur dioxide in its base form and then later convert it to the required sulfuric acid by the time it reaches the stratosphere. Their analysis shows that by virtue of this method, we could approximately halve the mass of aerosols that must be transported, thus making large scale deployment much more feasible. This demonstrates that changes to the aerosol delivery method

could massively reduce innovations needed to the deployment technology. Challenges faced by onboard production, however, include additional machinery and energy needed to power production, meaning that reducing payload mass does not totally eliminate the engineering challenges that come with large scale deployment. Furthermore, although producing sulfur on board may help with the technological feasibility of SAI deployment, its impact on economical feasibility is uncertain, and thus must also be considered (Smith et al., "Production of Sulfates on Board").

After the discussion of technological feasibility, our next focus should be on whether or not SAI is affordable enough for consistent deployment over long durations of time. Surprisingly, the direct economical costs of SAI are much easier to deal with than the technological difficulties. In their hypothetical scenario, Smith and Wagner (2018) estimated early deployment costs to be around $1,500 per tonne of aerosol injected, with total costs over the first 15 years equating to approximately $35.8 billion, or just under $2.4 billion, with an approximate additional $3.6 billion dollars in pre-start costs. However, these costs depend significantly on the deployment method and altitude. As mentioned previously, in "Review of Very High Altitude" Smith et al. reviewed platforms for 25 km deployment, and found that costs increased substantially, despite the higher aerosol reflective efficiency. This proves that estimates such as those given by Smith and Wagner should not be taken as the definitive cost of SAI, as economical feasibility varies greatly depending on a multitude of factors. Moreover, Moriyama et al. argue that previous cost estimates rely heavily on unrealistic and overly optimistic scenarios. They estimate that for a radiative forcing reduction of 2 W/m^2, annual costs could reach approximately $10 billion dollars, which is over four times the amount estimated by Smith and Wagner; this further proves the volatile nature of economical feasibility.

In addition, these estimates only cover the direct costs of SAI. Cost estimates in studies usually disregard expenses incurred by processes such as the monitoring and measuring process, as well as potential emergency expenses required in the case of facing unintended consequences (Huynh and McNeil). As a result, the expected final total economic cost of SAI should be higher than previous estimates which only include initial deployment costs. These costs would be difficult to estimate, however, as they vary wildly depending on the scale of deployment and the magnitude of danger presented by unintended events. Thus, although the direct costs of SAI seem relatively inexpensive, uncertainty with indirect costs makes it difficult to establish overall economical feasibility.

We can conclude that although small scale deployment is technically feasible with existing technology, large scale deployment would require specialized technologies, such as the hypothesized SAIL. It would be recommended for future research not just to look at advancing deployment technology, but also to find creative ways to carry the payload, as found with Smith et al. in "Production of Sulfates on Board". In terms of economical feasibility, we can not come to a final conclusion due to the many variables it relies on, and further research must be conducted on the economic trade-off between these variables.

Climate Consequences and Environmental Feasibility of SAI

The feasibility of SAI should not be determined solely with technological and economical feasibility in mind, but rather should also consider the consequences of deliberately dumping aerosols into the atmosphere on the climate. In addition, the type of aerosol deployed may have different effects on tropospheric chemistry, such as how sulfate aerosols may alter tropospheric ozone levels, and how its impact on incoming solar radiation may change water cycles. Furthermore, these effects can be either mitigated or exacerbated by the features of the aerosol such as its size and composition. Therefore, aerosol selection and its effects on atmospheric chemistry must be considered alongside technological and economical feasibility when discussing the overall feasibility of SAI.

SAI could potentially reduce global temperatures without reverting the hydrological cycle to what it was before the Industrial Revolution. Since greenhouse gases and solar radiation impact the energy balance of the earth in different ways, a decrease in incoming shortwave radiation would not simply undo every climate effect associated with increased atmospheric CO2 levels. SRM techniques in general lead to less energy reaching the surface of the earth, which decreases evaporation and alters precipitation. Thus, according to climate model experiments, the use of SRM, and thus SAI, could reduce global mean precipitation whether or not its cooling effects would counteract the warming induced by greenhouse gases (Bala et al.). This is important because it suggests that simply reducing temperatures would not be enough to correct changes done to the climate (Tilmes et al., "The Hydrological Impact"). Reductions in precipitation are common across many previous studies, however they vary in both magnitude and geographical distribution depending on the magnitude and spatial pattern of aerosol forcing (Huynh and McNeil). The magnitude of this hydrological response is significant, because even small changes in global precipitation can have consequences for water availability. In SAI simulations, the decreased level of solar radiation reaching the surface results in a reduction in precipitation and evapotranspiration, however the magnitude of these changes differ between models and scenarios (Cheng et al.), which highlights uncertainty in the environmental feasibility of SAI.

In addition, SAI can cause massive differences in precipitation between different regions due to shifts in large-scale atmospheric circulation and spatial variations in aerosol forcing (Kravitz et al.; Visioni et al.). For example, uneven aerosol distribution could alter sea surface temperature gradients and shift mid-latitude storm tracks, which could, in turn, lead to drying in regions such as the Amazon basin while leaving other areas unaffected (Simpson et al.; Tilmes et al., "CESM1"). This could be attributed to the uneven regional cooling of the ocean surface compared to land, which thus modifies moisture transport and atmospheric stability (Bala et al.). Disruptions in the hydrological cycle are not the only significant environmental impact of SAI, however.

The sulfate aerosols injected into the stratosphere could also have an effect on chemical reactions that affect atmospheric ozone. As we increase the scale of SAI, the concentration and thus the surface area of the sulfate particles increases. This leads to more opportunities for the

altering of chlorine, bromine, and nitrogen based catalytic cycles that are involved in the destruction of ozone through chemical reactions (Pitari et al.; Huynh and McNeil). In addition, sulfate aerosols can further affect ozone chemistry by heating the lower stratosphere and manipulating atmospheric circulation and water vapor concentrations (Tilmes et al., "Stratospheric Ozone Response"). Initially in the G6sulfur simulations of the Geoengineering Model Intercomparison Project, Antarctic springtime total-column ozone decreased by approximately 10 Dobson Units (DU), and then even further to around 20 DU by the end of the century relative to the high-emissions baseline. It should be noted, however, that these results were based on only three of the six models, as half were not based on interactive chemistry. Additionally, the results across these models were considerably different (Tilmes et al., "Stratospheric Ozone Response"). Importantly, ozone did not decrease everywhere across the globe. Under the same simulations, Northern Hemisphere mid-latitudes experienced increases of up to approximately 20 DU, which demonstrates the magnitude of variety in how SAI affects atmospheric chemical processes. It should therefore be clear that sulfate-based SAI presents many challenges to come in terms of environmental feasibility due to uncertainty in the magnitude and geographical distribution of its consequences (Huynh and McNeil).

Sulfate aerosols are commonly the subject of any SAI study. Due to their constraints on environmental feasibility, however, it would be worthwhile to explore the environmental effects of alternatives, such as calcium carbonate (calcite) and aluminum oxide (alumina) (Huynh and McNeil; Vattioni et al.). In particular, calcite is interesting due to its alkaline surface, which is able to react with acids such as hydrogen chloride and sulfuric acid; this allows for a potential reduction in the frequency of the chemical reactions responsible for ozone depletion. Under their assumptions, Keith et al. found that using calcite could produce cooling while increasing column ozone relative to a sulfate-based approach. Additional modeling of calcite and alumina suggests that they could reduce stratospheric warming compared to sulfate particles, however their effects on ozone remain uncertain due to insufficient data regarding reactions with these materials under stratospheric conditions (Vattioni et al.). Particle size also plays a role: smaller particles have a greater surface area for reactions but usually undergo faster coagulation, while larger particles last a shorter duration in the stratosphere (Huynh and McNeil; Vattioni et al.). Changing the aerosol material could thus potentially reduce the environmental disadvantages of sulfate-based SAI, but instead of removing uncertainty, it would bring about a new set of problems relating to atmospheric chemical processes.

Ethical Feasibility of SAI

Although the technological, economical, and environmental feasibilities of SAI are important, the ethical feasibility must also be discussed before deployment can be considered. Some of the major concerns regarding the ethical feasibility of SAI include distributive justice, procedural and intergenerational justice, and the possibility that it could reduce motivation for greenhouse gas mitigation (National Academies of Science, Engineering, and Medicine).

Climate change is unfair. Countries that have contributed little to nothing in terms of greenhouse gas emissions are facing the same consequences as countries that are majorly responsible for global warming. In addition, countries with less wealth may not be equipped with the technologies necessary for citizens to survive heatwaves, like AC units, which may be commonplace in wealthier countries. On the surface, SAI could be viewed as a potential solution to this problem, as its purpose is to reduce temperature change (Horton and Keith). As discussed before, however, the effects of SAI on the hydrological cycle could create regional differences in precipitation, which would create winners and losers (McLaren; Huynh and McNeil). And thus we face an ethical conundrum: the deployment of SAI could result in a net global benefit, but not without potentially sacrificing the needs of certain regions. It was argued by McLaren (2018) that assessments of not just SAI, but solar geoengineering in general, look through a lens of aggregates while ignoring potential differences between affected populations (McLaren). Consequently, the ethical feasibility of SAI can not be studied through global overall calculations alone; we must consider who experiences the advantages and disadvantages, whether vulnerable populations are disproportionately affected, and whether those populations have the ability to influence deployment decisions (Hourdequin).

Since SAI produces international effects, its consequences would not be restricted to the nation who initially deployed the aerosols. Determining who should have the authority in the deployment of SAI and whose interests are represented concerns an ethical challenge known as procedural injustice (National Academies of Sciences, Engineering, and Medicine). This can become a major issue, since obtaining unanimous agreement from the potentially affected nations is likely to be extremely difficult, especially if these nations view their respective risks and benefits differently. Hourdequin argues that an inclusive approach that gives attention to communities most affected by climate change and those historically excluded from decision making is required for justice in solar geoengineering. Additionally, this problem is intergenerational, meaning that current decisions could affect the lives of those not born yet. The National Academies of Sciences, Engineering, and Medicine thus identifies distributive, procedural, and intergenerational justice as interconnected considerations for solar geoengineering research and governance. SAI thus raises the questions as to whether any group or nation possesses the authority to make decisions with potentially both global and intergenerational consequences.

Perhaps one of the most critical ethical concerns is moral hazard, which is a term for how the development of SAI could reduce commitment to mitigating greenhouse gas emissions. If SAI was seen as simply a rapid method of limiting warming, governments and industries may be less likely to transition out of using fossil fuels (McLaren; National Academies of Sciences, Engineering, and Medicine). This is particularly important because SAI would not solve the root problem behind climate change; atmospheric greenhouse gas concentrations will continue to rise unless emissions are reduced. The National Academies of Sciences, Engineering, and Medicine identifies, however, that the extent to which such a moral hazard would occur would require further investigation. The issue here is that even research and public discussion of SAI could

influence political priorities and investment decisions. McLaren similarly argues that optimistic representation of geoengineering may unintentionally reinforce the use of high-carbon energy systems. The existence of this risk does not necessarily indicate that SAI research should come to a halt, however; intentionally leaving behind a potentially important climate intervention could have its own ethical consequences if generations of the future face severe climate risks (Biermann et al.). The moral hazard argument thus creates a conflict between researching a potentially useful climate intervention and making sure it does not serve as an excuse to delay emissions reductions.

Conclusion

In terms of technological and economical feasibility, SAI would seem simple to deploy compared to other cooling strategies, although further research on the complete economic cost is required, as it has not been studied enough for accurate long-term predictions. Any further research done relating to these areas of feasibility would have to focus on finding the most efficient deployment design, whether it is aeronautical or ground based, and would have to organize a long-term deployment plan. However the ability to simply deliver aerosols to the stratosphere should not be taken as the overall feasibility of SAI, as this would ignore the effects on the environment and ethical policies. The environmental consequences considered in this paper show that SAI would not reverse climate change. In fact, the reduction in incoming shortwave solar radiation can decrease global mean precipitation even when sufficient cooling is achieved, and the hydrological response can vary significantly across regions (Bala et al.; Cheng et al.; Tilmes et al., "The Hydrological Impact"). Sulfate aerosols are known to similarly affect atmospheric chemical processes and alter regional ozone levels, with the magnitude of these effects varying and remaining uncertain across models (Pitari et al.; Tilmes et al., "Stratospheric Ozone Response"). Aerosol materials such as calcite and alumina were suggested as an alternative to sulfate to reduce some of the environmental downsides to sulfate-based SAI, however their effects on atmospheric chemistry are less documented (Keith et al.; Huynh and McNeil; Vattioni et al.). Therefore, SAI could be considered environmentally feasible after sufficient research, but carries too much environmental risk to be considered as a reliable climate intervention before then. Furthermore, its ethical feasibility introduces challenges such as moral hazard, distributive, procedural, and intergenerational justice. This involves making sure that emissions reduction efforts are not replaced with SAI research, and ensuring that all communities who may be affected by SAI have influence on deployment decisions, while making sure those decisions are made with the future generations in mind. Its overall feasibility will ultimately depend on whether or not we can optimize costs, reduce uncertainties in its environmental impacts, and mitigate ethical risks.

Works Cited

Bala, G., et al. “Impact of Geoengineering Schemes on the Global Hydrological Cycle”. Proceedings of the National Academy of Sciences, vol. 105, no. 22, 2008, pp. 7664–69. https://doi.org/10.1073/pnas.0711648105.

Bednarz, Ewa M., et al. “Injection Strategy – a Driver of Atmospheric Circulation and Ozone Response to Stratospheric Aerosol Geoengineering”. Atmospheric Chemistry and Physics, vol. 23, no. 21, 2023, pp. 13665–84. https://doi.org/10.5194/acp-23-13665-2023.

Biermann, Frank, et al. “Solar Geoengineering: The Case for an International Non‐use Agreement”. WIREs Climate Change, vol. 13, no. 3, 2022, e754. https://doi.org/10.1002/wcc.754.

Cheng, Wei, et al. “Soil Moisture and Other Hydrological Changes in a Stratospheric Aerosol Geoengineering Large Ensemble”. Journal of Geophysical Research: Atmospheres, vol. 124, no. 23, 2019, pp. 12773–93. https://doi.org/10.1029/2018JD030237.

Crutzen, Paul J. “Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma?” Climatic Change, vol. 77, nos. 3–4, 2006, pp. 211. https://doi.org/10.1007/s10584-006-9101-y.

Horton, Joshua B., and David W. Keith. “Solar Geoengineering and Obligations to the Global Poor.” Climate Justice and Geoengineering: Ethics and Policy in the Atmospheric Anthropocene, edited by Cristopher J. Preston, Rowman & Littlefield International, 2016, pp. 79-92.

Hourdequin, Marion. “Geoengineering Justice: The Role of Recognition”. Science, Technology, & Human Values, vol. 44, no. 3, 2019, pp. 448–77. https://doi.org/10.1177/0162243918802893.

Huynh, Han N., and V. Faye McNeill. “The Potential Environmental and Climate Impacts of Stratospheric Aerosol Injection: A Review”. Environmental Science: Atmospheres, vol. 4, no. 2, 2024, pp. 114–43. https://doi.org/10.1039/D3EA00134B.

Iredale, Lindsay. “Eruption Products and Hazards.” Environmental Geology, Minnesota State Colleges and Universities, 2024, minnstate.pressbooks.pub/environmentalgeology/chapter/eruption-products-and-hazards/.

Jones, Andy, et al. “The Impact of Abrupt Suspension of Solar Radiation Management (Termination Effect) in Experiment G2 of the Geoengineering Model Intercomparison Project (GeoMIP)”. Journal of Geophysical Research: Atmospheres, vol. 118, no. 17, 2013, pp. 9743–52. https://doi.org/10.1002/jgrd.50762.

Keith, David W., et al. “Stratospheric Solar Geoengineering without Ozone Loss”. Proceedings of the National Academy of Sciences, vol. 113, no. 52, 2016, pp. 14910–14. https://doi.org/10.1073/pnas.1615572113.

Kravitz, Ben, et al. “Comparing Surface and Stratospheric Impacts of Geoengineering With Different SO2 Injection Strategies”. Journal of Geophysical Research: Atmospheres, vol. 124, no. 14, 2019, pp. 7900–7918. https://doi.org/10.1029/2019JD030329.

McClellan, Justin, et al. “Cost Analysis of Stratospheric Albedo Modification Delivery Systems”. Environmental Research Letters, vol. 7, no. 3, 2012, article 034019. https://doi.org/10.1088/1748-9326/7/3/034019.

McKay, David Armstrong, et al. “Updated Assessment Suggests >1.5°C Global Warming Could Trigger Multiple Climate Tipping Points”. Earth System Dynamics, vol. 13, no. 4, 2022, pp. 1485-1501. https://doi.org/10.1002/essoar.10509769.1.

McLaren, Duncan P. “Whose Climate and Whose Ethics? Conceptions of Justice in Solar Geoengineering Modelling”. Energy Research & Social Science, vol. 44, 2018, pp. 209–21. https://doi.org/10.1016/j.erss.2018.05.021.

Moriyama, Ryo, et al. “The Cost of Stratospheric Climate Engineering Revisited”. Mitigation and Adaptation Strategies for Global Change, vol. 22, no. 8, 2017, pp. 1207–28. https://doi.org/10.1007/s11027-016-9723-y.

Mwende, Wairimu G. “The Science of Climate Engineering: Risks and Benefits.” Eurasian Experiment Journal of Engineering, vol. 5, no. 1, 2025, pp. 100-106.

National Academies of Sciences, Engineering, and Medicine. Reflecting Sunlight: Recommendations for Solar Geoengineering Research and Research Governance. National Academies Press, 2021. https://doi.org/10.17226/25762.

Pitari, Giovanni, et al. “Stratospheric Ozone Response to Sulfate Geoengineering: Results from the Geoengineering Model Intercomparison Project (GeoMIP)”. Journal of Geophysical Research: Atmospheres, vol. 119, no. 5, 2014, pp. 2629–53. https://doi.org/10.1002/2013JD020566.

Ritchie, Hannah, and Max Roser. “CO_2 and Other Greenhouse Gas Emissions.” Our World in Data, https://ourworldindata.org/co2-and-greenhouse-gas-emissions. Accessed 6 Sept. 2026.

Robock, Alan. “Volcanic Eruptions and Climate”. Reviews of Geophysics, vol. 38, no. 2, 2000, pp. 191–219. https://doi.org/10.1029/1998RG000054.

Royal Society. Geoengineering the Climate: Science, Governance and Uncertainty. Royal Society, 2009.

Simpson, I. R., et al. “The Regional Hydroclimate Response to Stratospheric Sulfate Geoengineering and the Role of Stratospheric Heating”. Journal of Geophysical Research: Atmospheres, vol. 124, no. 23, 2019, pp. 12587–616. https://doi.org/10.1029/2019JD031093.

Smith, Jordan P., et al. “Production of Sulfates Onboard an Aircraft: Implications for the Cost and Feasibility of Stratospheric Solar Geoengineering”. Earth and Space Science, vol. 5, no. 4, 2018, pp. 150–62. https://doi.org/10.1002/2018EA000370.

Smith, Wake, et al. “Review of Possible Very High-Altitude Platforms for Stratospheric Aerosol Injection”. Environmental Research Communications, vol. 4, no. 3, 2022, article 031002. https://doi.org/10.1088/2515-7620/ac4f5d.

Smith, Wake, and Gernot Wagner. “Stratospheric Aerosol Injection Tactics and Costs in the First 15 Years of Deployment”. Environmental Research Letters vol. 13, no. 12, 2018, article 124001. https://doi.org/10.1088/1748-9326/aae98d.

Soden, Brian J., et al. “Global Cooling After the Eruption of Mount Pinatubo: A Test of Climate Feedback by Water Vapor”. Science, vol. 296, no. 5568, 2002, pp. 727–30. https://doi.org/10.1126/science.296.5568.727.

Tilmes, Simone, et al. “The Hydrological Impact of Geoengineering in the Geoengineering Model Intercomparison Project (GeoMIP)”. Journal of Geophysical Research: Atmospheres, vol. 118, no. 19, 2013, pp. 11036-58. https://doi.org/10.1002/jgrd.50868.

Tilmes, Simone, et al. “CESM1(WACCM) Stratospheric Aerosol Geoengineering Large Ensemble Project”. Bulletin of the American Meteorological Society, vol. 99, no. 11, 2018, pp. 2361–71. https://doi.org/10.1175/BAMS-D-17-0267.1.

Tilmes, Simone, et al. “Stratospheric Ozone Response to Sulfate Aerosol and Solar Dimming Climate Interventions Based on the G6 Geoengineering Model Intercomparison Project (GeoMIP) Simulations.” Atmospheric Chemistry and Physics, vol. 22, 2022, pp. 4557-73. https://doi.org/10.5194/acp-22-4557-2022.

Vattioni, Sandro, et al. “Injecting Solid Particles into the Stratosphere Could Mitigate Global Warming but Currently Entails Great Uncertainties”. Communications Earth & Environment, vol. 6, no. 1, 2025, article 132. https://doi.org/10.1038/s43247-025-02038-1.

Visioni, Daniele, et al. “Seasonally Modulated Stratospheric Aerosol Geoengineering Alters the Climate Outcomes”. Geophysical Research Letters, vol. 47, no. 12, 2020, article e2020GL088337. https://doi.org/10.1029/2020GL088337.

Pepper was Power: How Kanara's Small Kingdoms Negotiated With the Portuguese Empire in the Seventeenth Century By Arjun Shetty

Abstract

This paper looks at how the kingdoms of India's Kanara Coast used their control of pepper, ports, and trade networks to negotiate with the Portuguese Estado da Índia in the early seventeenth century. Portugal had armed fleets, coastal forts, and the cartaz licensing system, but its naval power did not give it control over pepper grown inland or shipped from local ports. This paper examines a seventeenth-century document from a published collection of Portuguese colonial correspondence and places it in the context of secondary historical scholarship. The archival record shows that in 1603, Portuguese officials were told to maintain "peace and good correspondence" with nearby rulers so that pepper shipments would continue. This shows that both sides had something the other wanted. Portugal could offer access to sea routes and trading privileges, while Kanara's rulers controlled access to pepper. Local alliances, independent merchants, and the coastal geography limited Portugal's ability to enforce its cartaz policy. This paper argues that pepper was not only a valuable trade good but was also a source of political power that allowed small kingdoms to influence the decisions of a much larger empire.

Introduction

In February 1603, Portuguese officials sent an instruction that revealed the main priority of their empire in India. Portugal had armed ships, forts, and a cartaz or licensing system (Pearson 38) meant to control trade across the Indian Ocean. Even so, in this archival record, the Crown told its representatives on the Kanara Coast to "conservar a paz e boa correspondencia", or "to conserve peace and good correspondence" with nearby rulers, especially the king of Banga, so that "the loading of pepper made at his ports does not fail" (Documentos Remettidos, vol. 1, pp. 4–6). The same letter explained that the Portuguese fortress at Onor, now Honavar, remained at peace with a local queen of Gersoppa because pepper also had to be loaded at her ports.

Portugal was known for its naval strength, but its officials were not ordering these rulers to be conquered. Instead, they were warning Portuguese representatives not to upset them. Portugal could threaten ships at sea, but it could not grow pepper or move it from inland farms to coastal ports without help from local rulers and merchants. The kingdoms of Kanara controlled pepper production and this gave them the leverage to negotiate with a global empire, because of Portugal's dependence on the pepper trade.

Portuguese power in Kanara was therefore based less on a coercive monopoly than on negotiation. The Estado da Índia wanted to monopolize the pepper trade, but keeping pepper shipments steady required cooperation from rulers whose freedom Portugal was trying to limit. Kanara's kingdoms could not compete directly with Portugal's navy. However, they controlled pepper-growing lands, ports, merchants, and connections to other buyers. They used these advantages to gain political influence. This paper argues that Portugal never gained complete

control of Kanara's trade in pepper and other goods. Instead, this trade became an ongoing negotiation in which both sides had something that the other needed.

Direct Access to Pepper and the Portuguese Sea Route to India

Direct access to pepper was one of the principal economic rationales for Portuguese expansion in Asia, and the reason their colonial adventure in Asia was so financially profitable. Well before Vasco da Gama landed in India in 1498, a taste for pepper was already well established in Europe. Yet, by the time the spice made its way to European markets, it had travelled through a series of Indian, Arab, Persian, Egyptian and Venetian middlemen, with every intermediary adding a price increase due to transport, duties and mark-up (Pearson 26–27). The Portuguese Crown saw an opening in a new sea lane around the Cape of Good Hope, one that would let them do away with those intermediaries. By buying directly from the Malabar and Kanara coasts, their goal was to sell directly in Europe for a substantial return. That profit margin was the goal of Vasco da Gama's trip and protecting this profit led directly to the subsequent infrastructure that came in its wake such as a series of fortresses and fleet of ships armed with cannon (Pearson 30–36; Subrahmanyam 144–46).

Pepper was used in many applications such as cooking, food preservation, and medicine, and its consumption carried real social status (Freedman 50–54). Pepper also has high value relative to its weight making it very well suited for profitable long-distance maritime trade (Freedman 112–15). Hence, a single ship's cargo could yield substantial returns once it reached European markets (Freedman 112–15). The Portuguese Crown treated access to pepper-producing regions as a matter of state, not merely of commerce. In fact, King Manuel I of Portugal formally titled himself "Lord of the Navigation, Conquest, and Commerce of Ethiopia, Arabia, Persia, and India," making the spice trade part of the very identity of the monarchy (Subrahmanyam 82). When it came to policy, the building of forts, the arming of fleets, and the making of treaties all had the primary aim of ensuring that this valuable commodity moved where it served Portugal's interest.

The Importance of Pepper to State Revenues and Kanara's Leverage

The Estado da Índia ('State of India') was the name the Portuguese gave to that part of their empire east of the Cape of Good Hope and at its height in the 16th century, the Estado da Índia stretched from Africa to Japan. More specifically, the Estado da Índia was an administrative apparatus established in 1505 to govern the empire and its trade network from its capital at Goa in India where the viceroy of the Indies was resident (Pearson 20–23, 31–35). The Estado da Índia was largely a string of interconnected ports linked by regular trading fleets. Rather than governing continuous inland territory, Portuguese officials held selected ports and strongholds, such as Goa, Hormuz, Malacca, and on the Kanara coast, Mangalore, from which they could tax shipping and pressure merchants to follow Portuguese controlled routes. Economic historians estimate that over half of the Portuguese crown's state revenue in the sixteenth century derived from two sources, West African gold and Indian pepper and spices,

with the spice trade contributing far more than gold to the royal treasury (Chaudhuri 70–74). At its peak, Portuguese pepper revenues reached over one million cruzados annually, collected through a royal monopoly administered in Lisbon, which taxed and resold every quintal that reached Portugal (Subrahmanyam 63). By the mid-sixteenth century, all spice-bearing ships were required to land at Lisbon and pay a duty of between eighteen and fifty cruzados per quintal, while royal contractors were assigned fixed annual export quotas, often around thirty thousand quintals from the Malabar and Kanara coasts (Subrahmanyam 116–18; Pearson 48). For most of the 16th century, pepper was one of the primary financial pillars of the Portuguese state. Revenue generated through the pepper trade financed much of the infrastructure that sustained the Estado da Índia throughout Asia. Income from pepper helped pay the salaries of royal officials, maintain the fleets that patrolled the Indian Ocean and the fortresses at strategic ports such as Goa, Hormuz, and Malacca, and support the administrative machinery that allowed Portugal to project influence across thousands of miles of coastline.

This scale of Portugal's pepper revenues provides the context for us to understand the 1603 letter's instruction to placate the local Kanara rulers instead of fighting them. The letter's soft language isn't diplomatic courtesy but is actually what fiscal dependence sounds like in an official instruction. Portuguese policy toward the pepper-producing regions of southwestern India consistently emphasized preserving commercial relationships alongside military preparedness. Every interruption in the flow of pepper threatened more than the profits of individual merchants. Reduced shipments affected customs revenues collected by the Estado da Índia, limited the Crown's ability to finance future expeditions, and increased the cost of maintaining the overseas empire. Consequently, securing reliable access to pepper became a strategic objective that influenced diplomacy, naval deployments, and treaty negotiations throughout the sixteenth and early seventeenth centuries.

This led to the Kanara coast having an importance that exceeded its relatively small political size. Pepper was grown inland and then carried through territories controlled by regional rulers. Kingdoms such as Banga, Ullal, and Keladi were small compared with major Indian states, but they held useful harbors, river mouths, and routes to the pepper-growing interior. Historian Nagendra Rao describes them as Kanara's "little kingdoms" (Rao 11). Although they had limited territory, their location gave them much greater importance in trade since any disruption to pepper production or transport threatened one of the key commercial foundations of Portuguese power.

The Cartaz as an Economic and Diplomatic Tool

Portugal tried to regulate Indian Ocean trade through the cartaz, a license issued to merchant ships (Pearson 30). Vessels sailing through waters controlled by Portugal were expected to carry one. The cartaz functioned as a compulsory maritime passport, serving as the primary legal and economic mechanism through which Portugal asserted control over Indian Ocean commerce. Rather than operating as a simple administrative permit, the document legally bound non-Portuguese ships to a strict set of regulations. Each license explicitly recorded the

ship's captain, crew size, and approved itinerary, forcing merchant vessels to call at Portuguese-controlled ports to pay heavy customs duties. Crucially, the system acted as an instrument of commercial monopoly. The terms of the cartaz strictly prohibited local traders from transporting highly valued, crown-monopolized commodities, most notably black pepper, under penalty of vessel seizure and cargo confiscation. Through this regulatory framework, the Estado da Índia generated significant revenue from ships it did not own, projecting diplomatic and financial authority over foreign merchants and independent coastal states alike. A ship found without a valid cartaz could be seized. Pepper was especially strategically important, and Portugal had a policy that tried to prevent non-Portuguese ships from carrying it (Subrahmanyam 64–67).

The system looked powerful on paper. It allowed Portugal to generate revenues from ships it did not own and make rules for merchants from countries it did not govern directly. Portuguese officials could also reward friendly rulers by granting cartazes or pressure them by refusing cartazes. A cartaz was therefore more than a shipping document. It was both a business advantage and a tool of diplomacy.

However, the system was difficult to enforce everywhere. The Kanara Coast had many rivers, backwaters, and small harbors spread across a long shoreline. Portuguese patrols could watch major routes, but they could not inspect every ship or remain in every area at once (Pearson 51–53). Monsoon winds also limited when large fleets could travel safely. A ban on carrying pepper only worked if Portuguese patrols actually found the ship, and a threat against a ruler only worked if Portugal was willing to pay the cost of a conflict.

Merchants created another problem for Portuguese control. Indian Ocean trading communities had existed long before Portugal arrived. They adjusted to new rules while keeping their relationships with local rulers and foreign partners (Das Gupta 24–48; Prakash 87–90). Merchants could use Portuguese trade routes when those routes were useful, but they could also work with Arab, Malabar, Dutch, or English buyers. The cartaz system could affect their choices, but it could not force permanent loyalty.

Portugal still had major naval advantages. Its ships could make resistance costly and could damage ports or capture vessels. Yet the ability to influence trade was not the same as monopolistic control over the trade. The Estado da Índia depended on decisions made by farmers, merchants, and rulers who were outside direct Portuguese control. Kanara's rulers had even more leverage once the Dutch and English merchants began their trading operations in India in the early to mid seventeenth century. Every new buyer gave pepper producers and rulers another choice. This raised the economic and political price Portugal had to pay for their cooperation.

This background enables us to understand the 1603 letter's emphasis on peace with Kanara rulers despite their having access to coercive tools like the cartaz.

Negotiated Commercial Order despite Portuguese Military Superiority

The two sides were far from equal in military power. Portugal had forts, cannon, oceangoing ships, and the ability to capture vessels. It could therefore deploy a serious military force. Kanara's rulers often competed with one another instead of acting as a united group, and Portugal could take advantage of these rivalries.

Even so, unequal military power did not remove the need for negotiation. Portugal could punish a ruler, but punishment did not automatically produce pepper. Destroying or blocking a port might stop the same trade that the Estado da Índia wanted to profit from. Portuguese officials therefore had to ask themselves whether a military victory would cause an economic loss.

This created an ongoing balance. Portugal used its navy and licensing system to shape trade. Kanara's rulers used pepper, geography, merchant networks, and rival buyers to influence Portuguese policy. The Estado da Índia was usually the strongest single power in this system, but it was never the only power that mattered.

Historians M. N. Pearson and Sanjay Subrahmanyam have explained that Portuguese authority often depended on practical relationships with Asian rulers, not only on conquest (Pearson 30–36, 47–48; Subrahmanyam 144–46). The archival record from Kanara presented in this paper shows a specific example of how this relationship worked in practice.

Conclusion

The 1603 instruction from Lisbon on the importance of peace with local rulers, shows that pepper gave Kanara's rulers influence beyond the military size of their kingdoms. Portugal could patrol sea routes, but it still depended on local rulers, merchants, ports, and inland producers to keep pepper moving. This dependence forced Portuguese officials to combine naval pressure with diplomacy. In Kanara, pepper was therefore not only a profitable commodity. It was a source of political leverage.

Portugal had significant influence on Indian Ocean trade, but it did not replace the older system that predated Portugal's arrival by centuries. Kanara's trade networks were too established and flexible to be controlled completely through licenses and coastal forts. Portuguese ships could patrol parts of the ocean, but they could not control the pepper harvest, inland roads, local ports, merchant decisions, and overseas markets all at the same time.

In Kanara, pepper was more than a spice or a source of profit. It was a form of power. The kingdoms that controlled it used that power to negotiate with an empire whose warships were powerful, but whose economy needed the profits from the pepper trade. Naval strength alone could not create a true monopoly when local rulers and merchants still controlled the goods Portugal wanted most.

Acknowledgement

I would like to very deeply and gratefully acknowledge Dr. Tukaram Poojary, retired Head of the Department of History at SVS College in Bantwal, Karnataka, India and founder/curator of the Rani Abbakka museum in Bantwal. His selfless and untiring actions for

over 25 years to keep Rani Abbakka's legacy alive, inspired me to research Rani Abbakka and her interactions with the Portuguese, which in turn led me to write this paper. I would also like to thank him for his guidance on doing archival research in Lisbon and also reviewing and providing feedback on the paper.

Works Cited

Chaudhuri, K. N. Trade and Civilisation in the Indian Ocean: An Economic History from the Rise of Islam to 1750. Cambridge UP, 1985.

Das Gupta, Ashin. The World of the Indian Ocean Merchant, 1500–1800: Collected Essays of Ashin Das Gupta. Edited by Uma Das Gupta, Oxford UP, 2001.

Documentos Remettidos da Índia, ou Livros das Monções. Edited by Raymundo António de Bulhão Pato, Academia Real das Sciências de Lisboa, 1880–1982.

Freedman, Paul. Out of the East: Spices and the Medieval Imagination. Yale UP, 2008.

Pearson, M. N. The Portuguese in India. Cambridge UP, 1987.

Prakash, Om. European Commercial Enterprise in Pre-Colonial India. Cambridge UP, 1998.

Rao, Nagendra. "The Portuguese and the Little Kingdoms of Kanara." Portuguese Studies Review, vol. 27, no. 2, 2019, pp. 7–25.

Subrahmanyam, Sanjay. The Portuguese Empire in Asia, 1500–1700: A Political and Economic History. Longman, 1993.

From Data to Decisions: Transforming Industrial Emissions Data into Consumer-Friendly Climate Information By Minal Matai

Executive Summary

Industrial manufacturing makes up for a large share of global greenhouse gas emissions. However, consumers rarely have access to clear information about the carbon footprint of the products they purchase. While nutrition labels, Energy Star ratings, and fuel economy stickers help Americans make informed decisions in other areas, there is currently no nationally standardized system that provides the carbon intensity of everyday manufactured goods.

The **PROVE IT Act of 2024** is an essential first step toward emissions transparency by directing the U.S. Department of Energy to measure and report greenhouse gas emissions linked with vital industrial products like steel, aluminum, cement, and plastics. However, the Act focuses on gathering emissions data instead of making that information accessible to consumers. Without a clear method of communicating technical emissions reports as understandable product information, much of the Act's potential public benefit may stay unrealized.

This white paper proposes the development of a **voluntary National Consumer Carbon Labeling Framework** that would translate emissions data into standardized, consumer-friendly labels for manufactured products. Like Nutrition Facts labels or Energy Star certifications, these labels would allow consumers to compare the carbon intensity of products using a standardized format based on verified data.

The proposed framework includes four key recommendations:

- Develop a standardized carbon labeling system using emissions data collected through the PROVE IT Act.
- Create a publicly accessible digital database where consumers can compare product emissions across manufacturers and industries.
- Foster voluntary partnerships with manufacturers before considering a wider implementation.
- Increase public education efforts to better the understanding of carbon intensity, climate impacts, and sustainable purchasing decisions.

A national carbon labeling framework would give benefits that go beyond consumer awareness. Greater transparency would reward manufacturers investing in cleaner production methods, strengthen American competitiveness in emerging low-carbon markets, decrease opportunities for greenwashing, and support future climate and trade policies built upon reliable emissions data.

Although challenges continue (including standardizing measurement methodologies, minimizing implementation costs, and ensuring public understanding), these obstacles can be addressed through implementation over time intervals, collaboration with industry, and continued federal leadership.

As governments increasingly implement carbon transparency into trade and manufacturing policy internationally, the United States has the opportunity to transform emissions data into information that empowers consumers, supports clean industry, and strengthens climate accountability. The recommendations presented in this paper offer a practical way for extending the impact of the PROVE IT Act beyond data collection and toward informed public action.

Introduction

Climate change has become one of the biggest challenges of the twenty-first century, and decreasing greenhouse gas emissions demands attention not only toward transportation and electricity but also to industrial manufacturing. The industrial sector (including the production of steel, aluminum, cement, plastics, chemicals, paper, and other materials) makes up for a large share of both local and global greenhouse gas emissions (United States Environmental Protection Agency). Because these products form the foundation of buildings, vehicles, infrastructure, and consumer goods, improving the environmental performance of manufacturing is vital for achieving long-term climate goals. At the same time, many American manufacturers have invested in cleaner methods of production, making parts of the U.S. industry among the least carbon-intensive in the world (Rorke et al.). The challenge lies in showing this advantage through transparent, reliable data that can be understood by all policymakers, businesses, and consumers.

An important concept in measuring industrial sustainability is carbon intensity, which refers to the amount of greenhouse gas emissions made per unit of output rather than the total amount of emissions produced (Richardson). Carbon intensity gives a standardized way to compare how efficiently different manufacturers or countries have made the same material. For example, two producers of steel may manufacture the same amounts of steel while generating distinct levels of greenhouse gas emissions. Measuring emissions intensity rather than total emissions allows governments and industries to find cleaner production methods and encourages technological innovation without failing to notice differences in the scale of production (Richardson). As climate policies rely more on emissions data, carbon intensity has become one of the most important criteria for considering environmental performances.

Despite the increased interest in sustainable purchasing, consumers currently have little access to standardized information about the carbon footprint of manufactured products. Unlike Nutrition Facts labels, Energy Star certification, or fuel economy ratings, there is no nationally recognized system that communicates the greenhouse gas emissions connected with producing common industrial goods. This information gap contributes to greenwashing, where broad or vague environmental claims can make products seem more sustainable than evidence shows (United Nations). Without consistent criteria for reporting emissions, consumers can often struggle to distinguish environmental improvements for marketing claims, decreasing their ability to make informed decisions on purchasing.

Many countries and organizations have begun experimenting with carbon labeling methods that provide consumers with information about the environmental impact of products. These initiatives show that emissions data can be turned into accessible labels and digital tools that encourage informed decision-making while benefitting manufacturers that decrease emissions (Michigan State University). Although these programs differ in methodology, they show a growing global movement toward greater transparency in product-level environmental information.

In the United States, Congress has taken an important first step through the Providing Reliable, Objective, Verifiable Emissions Intensity and Transparency (PROVE IT) Act. The bipartisan bill directs the U.S. Department of Energy to measure and compare greenhouse gas emissions intensity of many industrial products manufactured in the United States and major trading partners (United States, Congress, Senate 1). Products that are included are steel, aluminum, cement, plastics, fertilizers, glass, minerals, batteries, and other important materials (United States, Congress, Senate 2). The Department of Energy would publish these findings through a publicly accessible database and update the information periodically, providing policymakers with standardized emissions data for local and international manufacturing. Moreover, the bill focuses on data collection and transparency; it does not establish a carbon tax or carbon tariff.

The increasing importance of emissions transparency is closely related to developments in international trade. Nations including members of the European Union have begun adding or developing Carbon Border Adjustment Mechanisms (CBAMs) that account for the emissions added to imported goods (Chairman Shelley Moore Capito). As markets globally consider carbon intensity alongside price and quality, reliable emissions data will become vital for maintaining the competitiveness of cleaner manufacturers and fair international trade. Transparent emissions reporting also allows governments to design evidence-based climate policies while decreasing uncertainties for the industry and investors.

Although the PROVE IT Act creates the foundation for measuring industrial emissions, a gap remains between technical emissions reporting and public understanding. Most consumers are not likely to access Department of Energy databases or interpret emissions intensity reports when making purchasing decisions (Izcue et al. 1). As a result, the potential benefits of emissions transparency may remain limited to policymakers and industry experts instead of affecting wider consumer behavior. Research on sustainability labels suggests that consumers are more likely to consider environmental impacts when information is presented in a clear, standardized, and easily comparable format (Izcue et al. 1). Simply collecting emissions data is ineffective if that information is not communicated in a way that supports informed decision-making.

This white paper argues that the next steps beyond emissions measurement is the creation of a National Consumer Carbon Labeling Framework. Building upon the emissions data collected through the PROVE IT Act, such a framework would translate complex industrial emissions information into standardized labels and digital tools that consumers, businesses, and policymakers can easily understand. By connecting scientific measurement with accessibility, the

proposed framework looks to strengthen consumer transparency, benefit cleaner manufacturing, encourage innovation, and support future climate and trade policies rooted in reliable emissions data.

The Problem

Industrial Emissions Continue to Be a Huge Source of Greenhouse Gases

Decreasing greenhouse gas emissions has become one of the most important objectives of national and global climate policy. While a large amount of public attention focuses on transportation and the generation of electricity, industrial manufacturing continues to be one of the largest causes of greenhouse gas emissions. According to the U.S. Environmental Protection Agency (EPA), the industrial sector (including manufacturing, mining, construction, and other industrial activities) makes up for about 23% of total U.S. greenhouse gas emissions, making it the third-largest emitting sector in the country (United States Environmental Protection Agency). Emissions come from both energy consumption and industrial processes, specifically the production of carbon-intensive materials such as steel, cement, aluminum, chemicals, and plastics (United States Environmental Protection Agency).

These industries bring a unique challenge for climate policy because they need high-temperature processes and energy-intensive manufacturing methods that currently have limited low-carbon options (Center for Climate and Energy Solutions). As the national demand for construction materials, vehicles, consumer goods, and infrastructure continues to grow, emissions from heavy industry are expected to stay substantial without targeted policies encouraging cleaner methods of production.

Unlike many consumer products, the environmental impact of industrial materials is mostly invisible to the public. Consumers purchasing vehicles, appliances, buildings, or packaged goods rarely know whether the steel or aluminum used in those products was produced using mostly clean or highly polluting manufacturing practices. This lack of transparency limits the desire for manufacturers to invest in emissions reductions.

Carbon Intensity Provides a Better Measure of Manufacturing Performance

Climate policies are increasingly considering industrial performance through carbon intensity, a measure of greenhouse gas emissions generated per unit of production (Richardson). Instead of simply reporting total emissions, carbon intensity makes up for manufacturing efficiency by comparing emissions linked with producing the same quantity of material.

For example, two steel manufacturers may each make one million tons of steel annually, yet one may emit a lot less carbon dioxide because of cleaner energy sources, more efficient technologies, or better manufacturing processes. Carbon intensity allows these differences to be measured and compared across facilities, companies, and countries.

This evaluation has become especially important because global manufacturing is increasingly interconnected. Many products purchased in the U.S. are manufactured using

international supply chains, and thus, consumers often unknowingly purchase goods produced under vastly different environmental standards (World Economic Forum). Measuring carbon intensity creates a common standard for evaluating these products regardless of where they are produced.

Consumers Lack Reliable Carbon Information

Although consumers have become more interested in sustainability, they currently have limited access to standardized information about the carbon footprint of manufactured goods. Over the past decades, governments have successfully implemented consumer-facing labels that communicate important product information, including Nutrition Facts labels on food, Energy Star certifications for appliances, and fuel economy labels on vehicles (Carbon Footprint International). These programs show that standardized labeling systems can affect purchasing decisions while encouraging manufacturers to improve product performance.

No comparable national system exists for industrial carbon emissions.

While some companies voluntarily communicate sustainability information or take part in carbon certification programs, these efforts continue to be fickle. Existing carbon labels often rely on different calculation methods, cover different stages of a product's life cycle, or use scoring systems that make comparisons difficult (Hood). As a result, consumers cannot easily determine whether one product creates fewer greenhouse gas emissions than another.

This lack of standardized emissions information represents a notable market failure. Without consistent, comparable data, environmentally aware consumers cannot reward cleaner manufacturers through purchasing decisions, decreasing the economic need for companies to invest in emissions reductions.

Greenwashing Weakens Consumer Trust

The absence of standardized emissions reporting also contributes to greenwashing: the practice of presenting products or companies as more environmentally friendly than evidence supports (Gavilá, 2). Organizations such as the United Nations have found that greenwashing is a growing challenge because vague claims such as "eco-friendly", "carbon conscious", or "sustainably made" often lack transparent, verifiable data.

Recent regulatory developments across North America and Europe show increasing concern over misleading environmental claims (NielsenIQ, et al.). Governments are introducing stricter standards that need companies to prove sustainability marketing with measurable evidence rather than general claims.

Consumer trust depends upon transparency. Research consistently shows that individuals are more likely to trust environmental claims when they are supported by standardized criteria, independent verification, and publicly available methodologies (Kühne et al.). Without these standardizations, consumers struggle to differentiate genuine environmental leadership from marketing campaigns made to improve reputation.

Consumers Want Clear Sustainability Information

Many studies show that consumers consider environmental impacts more when making purchasing decisions. McKinsey & Company reports that many consumers express willingness to pay more for products advertised as environmentally sustainable, specifically when supported by trustworthy information (NielsenIQ, et al.). Academic research further shows that carbon footprint labels can have a positive effect on purchasing behavior, especially when labels are easy to understand and shown consistently (NielsenIQ, et al.).

Studies done across Europe have shown that standardized carbon labels improve the ability for consumers to compare products and encourage lower emission purchasing choices. More recent research suggests that carbon labels influence behavior across diverse groups, although their effectiveness relies a lot on label design, clarity, and consumer understanding (Lenk et al.).

Despite the growing demand, consumers currently lack access to consistent carbon information for most manufactured products sold in the U.S. This gap between consumer interest and available information limits the possible impact of market-driven climate action.

Current Policy Stops at Data Collection

Recognizing the importance of emissions transparency, Congress introduced the Providing Reliable, Objective, Verifiable Emissions Intensity and Transparency (PROVE IT) Act. The bill directs the Department of Energy to measure and report the greenhouse gas emissions intensity of many industrial products made both locally and globally (United States, Congress, Senate).

The PROVE IT Act represents a huge advancement because it forms standardized methods for comparing emissions across countries and industries while creating a reliable national database of emissions information. The bill also provides a foundation of science for future climate and trade policies by improving the quality and consistency of data.

However, the Act mainly addresses measurement instead of communication.

Although researchers, industry leaders, and policymakers may benefit from a technical emissions database, the bill does not create the methods needed for translating this information into formats that are accessible to consumers. Most Americans are unlikely to reference government emissions reports when purchasing products, leaving the practical benefits of emissions transparency mostly limited to experts.

Global Trade Increasingly Depends on Carbon Transparency

The importance of emissions transparency goes beyond environmental policy into international trade. The European Union has begun implementing its Carbon Border Adjustment Mechanism (CBAM), requiring importers of carbon-intensive goods to report greenhouse gas emissions and eventually pay carbon-related adjustments on certain imports (European Union).

Policies like these show a wider shift toward including emissions intensity into global trade. Accurate emissions reporting allows governments to develop evidence-based trade policies while also allowing manufacturers to show environmental performance in international markets.

Research also suggests that greater carbon transparency benefits businesses directly. Firms that give verified emissions information often experience lower financing costs, improved confidence for investors, and better reputations (Dionizio & Antoniol). Transparency can thus create economic value in addition to environmental benefits.

The Missing Public Connection

Altogether, these challenges show an important gap in current U.S. climate policy. Industrial emissions continue to be a major contributor to climate change, consumers look for reliable sustainability information more, and governments are investing in more accurate emissions measurement through bills such as the PROVE IT Act (Weidong et al. 165-177). Yet the information found by these efforts rarely reaches the consumers whose purchasing decisions help influence markets.

Without a standardized, consumer-friendly method of communicating emissions data, carbon transparency remains mostly limited to reports rather than making differences to everyday purchases. The U.S. has an opportunity to seal this gap by changing verified emissions data into accessible information that empowers consumers, rewards cleaner manufacturing, and strengthens the connection between scientific data and public action.

Proposed Solution: A National Consumer Carbon Labeling Framework

Building Upon the PROVE IT Act

The PROVE IT Act lays a foundation for emissions transparency by leading the U.S. Department of Energy (DOE) to develop standardized methods for measuring the greenhouse gas emissions intensity of important industrial products. By creating a consistent national database of verified emissions information, the Act addresses one of the greatest limitations of climate-informed trade policy: the lack of comparable, objective emissions data across industries and countries.

However, while the bill improves emissions measurement significantly, it does not address how this information reaches the public. Reports and emissions databases are valuable tools for policymakers, researchers, and manufacturers, but they are not designed for everyday consumer decision-making. As discussed in the previous section, consumers rarely have access to meaningful carbon information at the point of purchase, despite growing interest in sustainable products.

This white paper proposes the creation of a National Consumer Carbon Labeling Framework (NCCLF): a voluntary, government-supported system that translates verified emissions data collected through the PROVE IT Act into standardized, consumer-friendly product labels. Rather than changing existing legislation, the framework expands its impact by translating scientific measurement into information that consumers can understand and use. The proposed framework builds upon three key principles:

- Scientific Accuracy, using emissions data verified by federal methods.
- Consumer accessibility, giving information in a simple and standardized format.
- Market transparency, allowing manufacturers to communicate measurable environmental performance while allowing consumers to compare products more easily.

Framework Design

The proposed framework uses design principles that have proven to be successful in existing consumer labeling programs such as Nutrition Facts labels, Energy Star certifications, and internationally recognized carbon footprint labels.
Each product would have a standardized carbon label containing four primary elements:

1. Carbon Intensity Value: A numerical value that represents kilograms of carbon dioxide equivalent (kg CO2e) associated with producing one unit of the product.
2. Carbon Performance Rating: A simple rating system (for example, A-E or Low-Moderate-High Carbon) that allows consumers to understand emissions information quickly without requiring previous technical knowledge.
3. Government Verification Seal: A certification mark that shows the emissions calculations which follow DOE methods created through the PROVE IT Act.
4. QR Code: A scannable code linking consumers to detailed emissions information, manufacturing methods, comparison tools, and educational resources hosted by a national database.

Research on environmental labeling has consistently found that consumers respond more positively to labels that have both concise visual information and opportunities to access detailed explanations (Guenther et al., 445). A hybrid approach (using both an at-a-glance rating and optional information) balances conciseness with scientific transparency.

Framework

The National Consumer Carbon Labeling Framework would combine directly with the emissions reporting infrastructure established under the PROVE IT Act. Rather than creating an entirely new reporting system, it would go beyond the existing emissions measurement process to communication designed for consumers.
The proposed workflow is shown below:

- Manufacturer Production Data: Manufacturer submits production information under the DOE reporting standards.
- Department of Energy Verification: DOE calculates product carbon intensity using standardized national methods developed under the PROVE IT Act.
- National Carbon Database: Verified emissions information is published on a public database led by the DOE.

- Consumer Carbon Label: Participating manufacturers put standardized carbon labels on qualifying products, including QR codes linking directly to detailed emissions information.
- Consumer Decision-Making: Consumers compare products using consistent carbon information, increasing demand for lower-emission manufacturing.

This framework decreases the administrative burden by using emissions data already collected under existing legislation while extending its purpose beyond government reporting.

Key Features of the Framework

The proposed framework includes many features designed to maximize transparency, trustworthiness, and consumer engagement.

Standardized National Methodology: All participating products would use a single federal method for calculating emissions intensity. Standardization allows for the products to be compared fairly across manufacturers and industries.

Consumer-Friendly Label Design: Rather than presenting complex technical data, labels would communicate emissions information using visual indicators, concise ratings, and numerical values.

Digital Transparency:

QR codes would give consumers access to:

- Detailed emissions calculations
- Summaries of the manufacturing process
- Industry averages
- Product comparisons
- Educational materials explaining carbon intensity
- Updates as emissions data change over time

National Carbon Transparency Database: The DOE would keep a publicly searchable database which allows users to compare emissions across products, industries, manufacturers, and countries.

Voluntary Initial Participation: The framework would begin as a voluntary certification program focused on industries already reporting under the PROVE IT Act. As participation increases, policymakers could find opportunities for wider implementation.

Advantages Over Existing Approaches

Many carbon labeling programs already exist worldwide, including certifications led by the Carbon Trust, CarbonFree, and many other international eco-label organizations (Ecolabel Index). While these programs have increased emissions transparency, they remain underdeveloped.

Existing labels often differ in:

- Certification requirements
- Visual presentation
- Consumer interpretation
- Methodologies that consider emissions

As a result, products that have different carbon labels may be difficult to compare directly.

The proposed National Consumer Carbon Labeling Framework addresses these limitations by creating:

- One standardized methodology
- One recognizable national label
- One public database
- One verification authority

Unlike sustainability reports (which are often lengthy technical documents for investors) the proposed framework translates verified emissions information into an accessible format for everyday purchasing choices.

Importantly, the framework works with rather than replaces existing certifications. Manufacturers could continue participating in programs like Carbon Trust or CarbonFree while showing a nationally standardized label based upon federally verified emissions data.

Benefits to Stakeholders

The proposed framework gives measurable benefits for many stakeholder groups.

- Consumers: Consumers get access to clear, standardized information that supports informed purchasing decisions and lessens confusion caused by inconsistent environmental marketing.
- Manufacturers: Manufacturers investing in cleaner technologies gain greater recognition for emissions reductions. Transparent reporting rewards innovation while strengthening competitiveness in local and international markets.
- Policymakers: Governments benefit from increased public engagement with emissions data while strengthening the larger objectives of the PROVE IT Act and future climate policy efforts.
- Investors: Greater transparency decreases information discrepancies, allowing investors to better consider climate performance and long-term risks.
- Environment: By increasing the demand for lower-carbon products, the framework creates more needs for emissions reductions throughout industrial supply chains without pressing direct purchasing mandates.

Feasibility

Implementing a national carbon labeling framework is both administratively and technically possible because similar systems already exist globally.

The U.S. has years of experience administrating standardized consumer information programs, including Nutrition Facts labels, Energy Star certifications, and fuel economy labels (Ressecary). Similarly, international organizations have demonstrated that carbon footprint labels can be successfully applied across various types of products.

Because the PROVE IT Act already creates emissions measurement methods, much of the technical side needed for carbon labeling would already exist. The proposed framework therefore focuses mainly on translating verified emissions information into accessible consumer communication rather than creating an entirely new emissions reporting system.

Implementation over time intervals, beginning with voluntary participation, would further lower administrative costs while allowing government agencies and manufacturers to improve methods before wider adoption.

Challenges and Mitigation Strategies

Like any national labeling program, the proposed framework above faces many implementation challenges.

Firstly, accurately calculating product-level emissions can be resource-intensive, especially for manufacturers with complex international supply chains. Beginning with large manufacturers and industries already reporting under the PROVE IT Act would help create consistent methodologies before increasing participation.

Second, consumers may struggle to understand carbon information at first if labels become overly technical. Research suggests that simple rating systems along with detailed information improve consumer understanding without losing transparency (Röös & Tjärnemo, 982).

Third, keeping public trust requires rigorous verification. Government supervision, standardized calculation methods, and publicly available methodologies would decrease any inaccurate reporting or greenwashing.

Finally, standardization internationally remains an ongoing challenge because countries currently use different emissions calculation methods. Aligning the framework with globally recognized methodologies and working together with reporting programs would improve the ability to compare across international markets.

Key Points

The National Consumer Carbon Labeling Framework expands upon the PROVE IT Act by transforming verified emissions data into information that consumers can easily understand and apply. By including standardized labeling, transparency, and federal verification, the framework strengthens emissions accountability while supporting cleaner manufacturing, informed purchasing, and evidence-based policy.

Conclusion

The transition to a lower-carbon economy needs more than technological innovation or industrial policy, it also requires transparency. While governments, manufacturers and researchers recognize the importance of measuring greenhouse gas emissions more, consumers continue to be largely disconnected from the information needed to make informed purchasing decisions. As a result, opportunities to reward cleaner manufacturing, minimize greenwashing, and encourage market-driven emissions reductions remain limited.

The PROVE IT Act represents a historic advancement in U.S. climate and trade policy by creating a standardized system for measuring and publicly reporting the carbon intensity of important industrial products. By creating objective and verifiable emissions data, the bill provides an important scientific foundation for future policymaking. However, measurement alone is not enough. Without accessible methods for communicating this information, much of its potential value remains limited to technical reports rather than influencing everyday decisions made by consumers, businesses, and institutions.

This white paper has proposed the creation of a National Consumer Carbon Labeling Framework as an extension of the PROVE IT Act. By translating verified emissions data into standardized product labels, digital resources, and a public database, the framework fills in the gap between scientific measurement and consumer action. Like the success of Nutrition Facts labels and Energy Star certifications, a nationally recognized carbon label would provide clear, trustworthy information that encourages individuals to compare products based on their environmental impact while rewarding manufacturers that invest in cleaner production practices.

Beyond consumer education, the proposed framework offers broader economic and environmental benefits. Greater emissions transparency can strengthen the competitiveness of American manufacturers, support future carbon border adjustment policies, encourage investment in low-carbon technologies, and decrease opportunities for misleading environmental marketing. Standardized carbon labeling would also place the U.S. alongside an increasing number of countries including emissions transparency into trade, manufacturing, and sustainability policy.

Implementing a national carbon labeling system will require collaboration among federal agencies, manufacturers, retailers, researchers, and consumers. Challenges like emissions verification, standardized methodologies, costs of implementation, and public understanding should not be ignored. However, successful existing and international carbon labeling show that these challenges are manageable through implementation over time intervals, standards that are consistent, and engagement from stakeholders.

Overall, climate action depends not only on collecting better data but also on making that data purposeful. When consumers have access to clear and trustworthy information about the environmental impacts of the products they purchase, they become actively involved in driving cleaner markets and more sustainable industries. The National Consumer Carbon Labeling Framework provides a path for transforming emissions transparency into informed action.

As the United States continues to strengthen its climate and manufacturing policies, the next step is not simply to measure emissions, but to communicate them. Building upon the

foundation created by the PROVE IT Act, a national carbon labeling framework can help create a future where transparency, innovation, and informed consumer choice work together to support a more competitive economy and a more sustainable planet.

Final Recommendation

Congress and the DOE should explore the creation of a National Consumer Carbon Labeling Framework that builds upon the emissions data found through the PROVE IT Act. Programs involving manufacturers, retailers, researchers, and consumers should be put into action to evaluate labeling design, public understanding, and the feasibility of implementation. Lessons learned from these programs can inform future standards that improve emissions transparency, strengthen American manufacturing, and encourage consumers to make informed and climate-aware purchasing decisions.

Work Cited

Airaksinen, Saara. "How to Use Carbon Footprint Labels?" Biocode, 30 Nov. 2023, biocode.io/carbon-footprint-labels/.

Carbon Footprint International. "The International Network of Carbon Footprint Communication Programmes." Carbon Footprint International, carbonfootprintinternational.com/. Accessed 6 July 2026.

Chairman Shelley Moore Capito. "Whitehouse, DelBene Reintroduce Carbon Border Adjustment to Boost Domestic Manufacturers, Tackle Climate Change." U.S. Senate Committee on Environment and Public Works, 17 Dec. 2025, www.epw.senate.gov/public/index.cfm/2025/12/whitehouse-delbene-reintroduce-carbon-border-adjustment-to-boost-domestic-manufacturers-tackle-climate-change.

Center for Climate and Energy Solutions. "Controlling Industrial Greenhouse Gas Emissions." Center for Climate and Energy Solutions, www.c2es.org/content/regulating-industrial-sector-carbon-emissions/. Accessed 8 July 2026.

Chen, Jessica, et al. "Business Case on Carbon Labeling." American Sustainable Business Network, asbnetwork.org/wp-content/uploads/2025/02/asbn_carbon_labeling_business_case.pdf. Accessed 9 July 2026.

Craig, Anthony J. Measuring Supply Chain Carbon Efficiency : a Carbon Label Framework. 2012. PhD thesis. Massachusetts Institute of Technology, dspace.mit.edu/entities/publication/b43c4a83-5f88-4ab3-a91c-408c16230631. Abstract.

Dionizio, Emily, and Virginia Antoniol. "The Role of Carbon Transparency and International Cooperation in Supply Chains for an Effective and Just Transition." Forum on Trade, the Environment, and the SDGs, 11 May 2026, tessforum.org/latest/the-role-of-carbon-transparency-and-international-cooperation-in-supply-chains-for-an-effective-and-just-transition.

Directorate-General for Mobility and Transport. "EU Introduces Flight Emissions Label for More Informed and Sustainable Travelling." European Commission, 18 Dec. 2024, transport.ec.europa.eu/news-events/news/eu-introduces-flight-emissions-label-more-informed-and-sustainable-travelling-2024-12-18_en.

Dreist, Denise, et al. "EU Climate Labeling Policy: Analyzing Consumer's Comprehension of CO2 Footprint Labels." Journal of Agriculture and Food Research, vol. 17, Sept. 2024, p. 101248, https://doi.org/10.1016/ j.jafr.2024.101248.

Ecolabel Index. www.ecolabelindex.com/ecolabels/?st=category,carbon. Accessed 9 July 2026.

European Union. "Carbon Border Adjustment Mechanism." European Union, taxation-customs.ec.europa.eu/carbon-border-adjustment-mechanism_en. Accessed 8 July 2026.

Feucht, Yvonne, and Katrin Zander. "Consumers' preferences for carbon labels and the

underlying reasoning. A mixed methods approach in 6 European countries." Journal of Cleaner Production 178 (2018): 740-748.

Finstad, Brad. "Finstad's Prove It Act Passes House Judiciary Committee." Brad Finstad, 2025, finstad.house.gov/2025/05/21/2025-5-finstad-s-prove-it-act-passes-house-judiciary-committee/.

Fu, Shanshan. "Carbon Labeling on Products: Global Practice and Implementation in China." China Briefing, 25 July 2023, www.china-briefing.com/news/carbon-labeling-on-products-global-practice-and-implementation-in-china/.

Garwood, Grace. "Carbon Labeling Gains Momentum Despite Challenges." The Food Institute, 6 Apr. 2021, foodinstitute.com/focus/carbon-labeling-gains-momentum-despite-challenges/.

Gavilá, Pilar Galán. "The Challenge of Greenwashing: An International Regulatory Overview." KPMG Law, 2025, assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2025/12/greenwashing-regulatory-overview.pdf.coredownload.inline.pdf.

Ge Vernova. "Why Accurate Emissions Data Is Crucial for Impactful Decarbonization Efforts." Ge Vernova, 6 Jan. 2026, www.gevernova.com/software/blog/how-accurate-emissions-data-create-impactful-decarbonization-efforts.

Guenther, Meike, Caroline M Saunders, and Peter R Tait. 2012. “Carbon Labeling and Consumer Attitudes.” Carbon Management 3 (5): 445–55. doi:10.4155/cmt.12.50.

Hood, Chris. "The Age of Transparency and Its Impact on Consumer Trust." Chris Hood, 13 Mar. 2024, chrishood.com/the-age-of-transparency-and-its-impact-on-consumer-trust/?srsltid=AfmBOopTv9Z3fwJOl6HvtxueiMvgNANSyj82Siepu75MxWks65BulL0n.

Izcue, Jose, and Yasna Palmeiro-Silva. “A systematic review on how environmental sustainability and social responsibility food labels influence consumers' food choices and purchasing decisions.” Public health nutrition vol. 29,1 e62. 11 Feb. 2026, doi:10.1017/S1368980026101979.

Kennedy, Kevin, et al. "4 US Congress Bills Related to Carbon Border Adjustments in 2023." World Resources Institute, 13 Dec. 2023, www.wri.org/update/4-us-congress-bills-related-carbon-border-adjustments-2023.

Kühne, Swen J et al. “Can carbon labels encourage green food choices?.” Frontiers in psychology vol. 13 902869. 27 Jan. 2023, doi:10.3389/fpsyg.2022.902869.

Lenk, Julia Diana et al. “Which Consumers Change Their Food Choices in Response to Carbon Footprint Labels? The Role of Political Ideology and Other Socio-Demographic Factors.” Nutrients vol. 17,8 1321. 10 Apr. 2025, doi:10.3390/nu17081321.

Li, Michelle. "Carbon Labels." Clever Carbon, clevercarbon.io/carbonlabels. Accessed 9 July 2026.

Liu, Tiantian, Qunwei Wang, and Bin Su. "A review of carbon labeling: Standards, implementation, and impact." Renewable and Sustainable Energy Reviews 53 (2016): 68-79.

Loughlin, Bronagh. "The Role of Carbon Labels in Improving Climate Consciousness." Institute of Sustainability Guides, 22 Feb. 2024, instituteofsustainabilitystudies.com/insights/guides/role-of-carbon-labels/.

Man, Sophia. "Would You Buy Products with Carbon Labels?" Earth.org, 2 June 2021, earth.org/carbon-labels-on-products/.

Meyerding, Stephan GH, Anna-Lena Schaffmann, and Mira Lehberger. "Consumer preferences for different designs of carbon footprint labelling on tomatoes in Germany—does design matter?." Sustainability 11.6 (2019): 1587.

Michigan State University. "What Does Carbon Impact Look like on a Food Label?" Michigan State University Extension Economic Development, 14 Jan. 2021, www.canr.msu.edu/news/what-does-carbon-impact-look-like-on-a-food-label.

NielsenIQ, et al. "Consumers Care about Sustainability—and Back It up with Their Wallets." McKinsey and Company, 6 Feb. 2023, www.mckinsey.com/industries/consumer-packaged-goods/our-insights/consumers-care-about-sustainability-and-back-it-up-with-their-wallets.

One Planet. "Using Carbon Labels to Enable Informed Consumer Choices." One Planet, 12 Aug. 2020, www.oneplanetnetwork.org/news-and-events/news/using-carbon-labels-enable-informed-consumer-choices.

UN Environment Programme. "What Is Climate Transparency and Why Is It Important?" UN Environment Programme, 29 Aug. 2025, www.unep.org/news-and-stories/story/what-climate-transparency-and-why-it-important.

United Nations. "Greenwashing – the Deceptive Tactics behind Environmental Claims." United Nations Climate Action, www.un.org/en/climatechange/science/climate-issues/greenwashing. Accessed 6 July 2026.

United States, Congress, Senate. PROVE IT Act of 2024, S. 1863, 118th Cong., 2nd sess., introduced 7 June 2023. Congress.gov, www.congress.gov/bill/118th-congress/senate-bill/1863.

United States Environmental Protection Agency. "Industry Sector Emissions." United States Environmental Protection Agency, 10 Mar. 2026, www.epa.gov/ghgemissions/industry-sector-emissions.

Reith, Nils von. "Navigating the Future of Green Claims: What Do Automotive Manufacturers and Suppliers Need to Know?" TaylorWessing, 17 Mar. 2025, www.taylorwessing.com/en/insights-and-events/insights/2025/03/navigating-the-future-of-green-claims.

Ressecary. "Carbon Labeling in the United States – Illustrated by CarbonFree." Ressecary,

www.reccessary.com/en/reccpedia/footprint-label/carbon-labeling-in-the-united-states-carbonfree.
Accessed 8 July 2026.
Richardson, Ben. "Carbon Intensity Explained: What It Is and How to Reduce It." Zevero, 21 May 2026, www.zevero.earth/blog/Carbon-intensity-what-it-is-and-how-to-reduce-it.
Röös, Elin, and Heléne Tjärnemo. "Challenges of Carbon Labelling of Food Products: A Consumer Research Perspective." British Food Journal, vol. 113, no. 8, 9 Aug. 2011, pp. 982-96,
https://doi.org/10.1108/00070701111153742.
Rorke, Catrina, et al. "Understanding the U.S. Carbon Advantage: A Foundation for Strategic Policy Decisions." Climate Leadership Council, 14 Apr. 20225, clcouncil.org/publications/blog/understanding-carbon-advantage/.
Soe, Aye Chan, and Chavis Ketkaew. "Understanding Consumer Responses to Carbon Footprint Product Labels: A Multi-group Analysis of Online and Offline Shopping Engagement in Retail
Contexts." Digital Business, vol. 5, no. 2, Dec. 2025, p. 100158,
https://doi.org/10.1016/j.digbus.2025.100158.
Vanderbilt University. "Carbon Labeling: Putting the Power in Consumer's Hands." Vanderbilt University, 29 Mar. 2011, news.vanderbilt.edu/2011/03/29/carbon-labeling/.
Wang, William. "Understanding Carbon Label: Empowering Enterprises for Low-Carbon Transformation." Million Pack, 12 Jan. 2026, millionpack.com/carbon-label/.
World Economic Forum. "If Consumers Care about Sustainability, Why Are Sustainable Choices Still Hard?" World Economic Forum, 15 Mar. 2026, www.weforum.org/stories/nature-and-biodiversity/consumers-sustainability-choices-world-consumer-rights-day/.
Xu, Weidong & Sun, Zixun & Ni, He, 2024. "Transparency pays: How carbon emission disclosure lowers cost of capital," Economic Analysis and Policy, Elsevier, vol. 83(C), pages
165-177.

The "Buy Now, Pay Later" (BNPL) Debt Trap:
A Secondary Research Analysis of Consumer Credit Risk, Behavioural Drivers, and Regulatory Responses By Jasroop Singh Narang

Abstract

The explosion in the use of Buy Now, Pay Later (BNPL) credit products has revolutionised the market for consumer credit due to their capacity to enable consumers to split their payments over short-term periods, typically interest-free, at the time of purchase. While BNPL credit products are often positioned as a low-risk option compared to credit card products, there is increasing evidence drawn from a variety of secondary sources that the combination of structural features of the product, lack of regulation, and consumer behaviour renders BNPL credit products similar to debt traps for financially vulnerable consumers. The current paper presents a synthesis of research evidence from government reports, peer-reviewed journal articles, and institutional data sources, generally dating from 2013 to 2026, on the mechanisms through which BNPL credit use leads to overspending and loan stacking. Narrative systematic review methodology informed by inclusion and exclusion criteria is employed using PubMed, ScienceDirect, Google Scholar, and institutional sources such as the Consumer Financial Protection Bureau (CFPB) and the Federal Reserve System. Findings show that usage of BNPL is prevalent among younger individuals and those who are poor and have bad credit; that installment payments decrease the significance of spending through processes explained by mental accounting and hyperbolic discounting theory; and that there is a correlation between BNPL usage and increased credit card debt, personal loan debt, late payment fees, and likelihood of suffering from depression, anxiety, and posttraumatic stress. The paper discusses how BNPL can offer some advantages, such as cash flow management and interest-free loans for unbanked consumers, and reviews regulatory responses to BNPL in the United States, United Kingdom, Australia, Europe, and India. The paper finds that BNPL is not all negative nor all positive, but that the risks of BNPL usage are disproportionately borne by the most vulnerable consumers.

Keywords buy now pay later, BNPL, consumer debt, fintech, financial literacy, debt trap, behavioral economics

CHAPTER 1. INTRODUCTION

Historically, consumer credit has been available via financial instruments that were regulated through disclosure standards, interest rate limits, and credit bureau reporting in most jurisdictions, including credit cards, personal loans, and retail credit. In the past ten years, an innovative form of point-of-sale credit referred to as "Buy Now, Pay Later" or BNPL has emerged, falling outside much of this regulatory framework. BNPL refers to an arrangement where the consumer receives goods or services on the spot while dividing the total payment into a few scheduled payments, most commonly four, paid biweekly, with the first installment paid upfront at the moment of check-out (Checkbook.org, 2025). Since BNPL does not charge any

interest as long as payments are timely, the financial service has been presented as a financially sensible alternative to consumer revolving credit card debt.

However, the scale of BNPL usage has increased significantly. Industry forecasts mentioned in consumer advocacy research indicate that transaction volume of BNPL in the USA will be approximately $122 billion in 2025 and $184 billion by 2030, with the number of Americans using a BNPL product estimated at 86.5 million in the previous year (Checkbook.org, 2025). There is confirmation of this trend from the government agency's statistics: according to the CFPB (2025b), the six biggest BNPL lenders made 335.8 million loans worth $45.2 billion in 2023 alone, with unique users reaching 53.6 million, which means twelve percent increase from the previous year. In addition, a survey of the Federal Reserve estimates that fifteen percent of American adults, or more than 38 million people, used BNPL at least once in 2024 (Board of Governors of the Federal Reserve System, 2025, as cited in CFPB, 2025b). This is not the case in the United States alone: the same trend of rapid growth has been observed in the UK, Australia, the EU, and South/Southeastern Asia, where digital lending based on fintech increased consumer access to credit beyond the banking system (Reserve Bank of India, 2022).

What drives the focus of this research paper is the possibility of misusing the same properties of BNPL that make it appealing. The fact that not all BNPL loans are registered by the main credit bureaus allows a consumer to take several BNPL loans from various companies simultaneously, while neither of the lenders will be aware of the other outstanding debts. This problem has even received the name "phantom debt" among regulators (Federal Register, 2025). According to the research conducted by the CFPB (2025a), more than sixty percent of BNPL repeat users had more than one loan at a time, and BNPL borrowers had much higher levels of other unsecured debt than consumers who did not use BNPL. Moreover, BNPL was associated with depression, anxiety, and post-traumatic stress according to recent scientific studies (Shah et al., 2025). All of these factors have pushed consumer groups, journalists, and scientists to consider BNPL loans to be potentially a "debt trap," a term which was used only for payday loans before.

The current study represents the synthesis of the literature in the secondary form in order to analyze and assess the validity of the debt trap claim. Due to the relatively recent emergence of BNPL, the lack of primary longitudinal studies is notable; most existing evidence comes from cross-sectional surveys, data highlights related to regulations, and behavioral finance research inspired by credit card debts and payday lending. Thus, the purpose of the paper is not to provide any new data, but to critically synthesize the existing secondary information, government data, and peer-reviewed research to provide an understanding of how BNPL works, who the most vulnerable groups are, and what benefits and regulations can be considered in relation to the issue. The structure of the paper is the following. Chapter 2 represents a literature review in the narrative form. Chapter 3 describes a theoretical framework based on behavioral economics. Chapter 4 describes the research methodology used for searching and synthesizing the sources. Chapter 5 presents the results and discusses the benefits, risks, and trends. Chapter 6 is devoted to conclusions.

CHAPTER 2. LITERATURE REVIEW

In terms of research and literature regarding BNPL credit products, there has been a significant rise in studies since 2022 due to the fast growth of the product. Descriptive early work on the topic was limited to determining basic market facts, such as demographics and uses of BNPL products. Based on a sixteen percent sample of applications by six large BNPL lenders from 2017 to 2022, the CFPB (2025a) stated that twenty-one percent of Americans with a credit record used at least one BNPL product offered by a major lender in 2022, and sixty-one percent of BNPL customers have subprime and deep subprime credit scores; that is, their FICO score was lower than 620 or 580 (cited in Silverstein, 2025). In addition, Stavins (2024) reported, writing for the Federal Reserve Bank of Boston, that the usage of BNPL products tends to be prevalent among young people and those with limited access to credit markets. That was confirmed later by the Federal Reserve Bank of Kansas City (2025), which found a significant correlation between the late payment of BNPL users and overall financial vulnerability.

Another line of literature has studied how BNPL services interact with, but do not substitute for, other forms of debt. In one of the early academic economic studies on BNPL services, Guttman-Kenney et al. (2023) examined United Kingdom credit card transactions and found that a substantial number of consumers financed BNPL installments through a credit card, turning their interest-free arrangement into revolving debt at interest rates that the authors highlighted typically reach up to twenty percent per year. Such practice, according to the authors, casts a shadow of doubt on whether such consumers could afford the respective purchase. Supporting the findings from the United Kingdom with those from the United States, CFPB (2025a) reported that consumers who had taken out at least one BNPL loan per month were holding $453 more in personal loans and $871 more in credit cards on average than other similar consumers who had not used BNPL services. Since BNPLs do not report information about loan payments to credit reporting agencies, an earlier CFPB survey known as Making Ends Meet determined that BNPL customers are more likely to have revolving credit cards, had trouble with delinquencies on conventional credit products, and have used high-cost financial services like payday loans, pawn loans, and overdraft coverage even though many of the differences existed before the consumer's first BNPL product purchase (CFPB, 2023). This discovery challenges a simplistic understanding of causation, as BNPL does not cause financial instability but simply attracts individuals who suffer from such an issue.

The third strand of the literature, which is more recent, examines the psychological/clinical correlates of BNPL. In a study in the JAMA Health Forum, Shah et al. (2025) examined nationally representative data from 2,121 U.S. adults and found that people with depression were 1.91 times more likely than others to engage in BNPL after controlling for demographic covariates, whereas those with anxiety and posttraumatic stress were 1.77 and 2.35 times more likely, respectively. Shah et al., along with co-authors from the Johns Hopkins Bloomberg School of Public Health and de Beaumont Foundation, pointed out that the connection between mental health and BNPL was likely bidirectional: financial insecurity can lead to mental problems, but mental problems can also lead people to seek relief in financial

instruments such as BNPLs that give immediate gratification at the expense of future costs. This is in line with earlier findings, including a systematic review and meta-analysis by Richardson et al. (2013) in the Clinical Psychology Review showing a significant link between unsecured debt and both depression and anxiety in many studies before the advent of BNPL.

Fourthly, there has been research from marketing, consumer psychology, and cultural economy academic literature about what makes BNPL effective from the psychological perspective in inducing consumption, especially among the younger demographic. Relja et al. (2024) conducted research with users of BNPL services in the United Kingdom and noted the key aspects of continued BNPL usage to be related to convenience, control, and diminished pain of paying. Threadgold et al. (2025) expanded on this idea by noting the use of gamification through the use of progress bars and reward messages in BNPL apps, which create a sense of repayment as a task like playing a game, rather than fulfilling a financial obligation, which is especially true for young people forming their financial habits independently. Studies from South and Southeast Asian contexts reinforce this picture: research on young consumers in Malaysia found that limited financial literacy and strong peer influence increase vulnerability to hidden fees, repayment delays, and debt accumulation, even as the same convenience features expand access to goods otherwise unaffordable in a single payment. Similar studies of Millennial and Generation Z consumers in India have linked lower financial literacy directly to higher rates of BNPL misuse and recommend targeted financial education as a mitigating strategy.

Fourth, there is a relatively smaller but growing literature on the social implications of being over-indebted due to BNPL schemes. According to the research report on the topic of financial stress issued by the American Psychological Association (2026), there was qualitative research conducted by the researcher from Lund University, Erik Eriksson, on young adults in Sweden who had gotten into unsecured debts, partly because of short-term borrowing through digital means and were having problems with collection agencies; these individuals found it extremely challenging to secure housing and mortgage loans, and the financial stress was affecting their romantic relationships and desire to start families. It was clear that over-indebtedness in this group of people was due to the need to finance necessities, not discretionary purchases.

In each of these four threads of research, one finds a clear if complicated consensus emerging. The use of BNPL services is highly skewed towards younger and economically vulnerable individuals; the payment model and the way in which BNPL services are designed make loan stacking more likely; the use of such services correlates with signs of psychological distress; and the dangers involved are tied to existing problems of financial vulnerability, making it hard to blame BNPL for "causing" debt.

CHAPTER 3. THEORETICAL FRAMEWORK

Two complementary strands of behavioral economic theory provide the conceptual foundation for understanding why BNPL is capable of inducing overspending even among

consumers who do not intend to overextend themselves: mental accounting theory and hyperbolic discounting.

Mental accounting, first formalized by Thaler (1985, 1999), describes the tendency of individuals to treat money as though it were divided into separate, non-fungible mental categories rather than a single fungible resource, and to evaluate transactions in terms of "transaction utility," meaning the perceived quality of a deal, rather than purely in terms of absolute cost. Under this framework, splitting a $200 purchase into four payments of $50 does not simply defer cost; it reframes the transaction itself, making the immediate outlay appear smaller and more easily absorbed within a consumer's short-term mental budget, even though the total financial obligation is unchanged. Prelec and Loewenstein (1998) extended this reasoning specifically to debt, arguing in their theory of the "pain of paying" that separating the moment of consumption from the moment of payment reduces the psychological discomfort normally associated with spending, and that credit instruments which widen this separation, as BNPL explicitly does by design, tend to increase both the frequency and size of purchases. Zhang and Sussman (2018) further noted that mental accounting is not a fixed trait but is actively shaped by the design of financial products themselves, meaning that BNPL interfaces, which visually present a purchase as a manageable installment rather than a lump-sum debt, are not neutral presentations of the same information but active behavioral interventions that can systematically bias consumer judgment.

The second theoretical pillar, hyperbolic discounting, addresses the time-inconsistency of consumer preferences. Classical economic models assume that individuals discount future costs and benefits at a constant exponential rate; however, foundational work by Strotz (1956) demonstrated that decision-makers exhibit dynamically inconsistent preferences, valuing immediate outcomes disproportionately more than outcomes even a short time in the future, a pattern later formalized as hyperbolic discounting. Applied to BNPL, this theory predicts that a consumer evaluating a purchase at checkout will systematically undervalue the cumulative burden of three or four future installments relative to the immediate gratification of acquiring the product now, even when the consumer would, in a calmer moment of reflection, judge the future payments to be a meaningful financial commitment. Because BNPL is typically offered at the moment of highest purchase motivation, the checkout screen, it is structurally positioned to exploit precisely this bias, a concern echoed in the applied literature on consumer autonomy and algorithmic design, which has begun to model BNPL explicitly as a debt instrument capable of interacting with hyperbolic discounting to distort intertemporal consumption decisions.

These two theories are mutually reinforcing rather than competing explanations. Mental accounting explains why the same total cost feels smaller when partitioned into installments, while hyperbolic discounting explains why consumers accept future installment obligations that a more patient, exponentially discounting agent would reject. Together, they offer a parsimonious behavioral account of the findings reviewed in Section 2: elevated spending following BNPL adoption, a tendency toward loan stacking, and a disproportionate concentration of BNPL use among consumers who may have reduced cognitive bandwidth for long-horizon financial

planning. This framework does not suggest BNPL users are irrational; rather, it suggests BNPL, as currently designed, interacts with well-documented, near-universal cognitive biases in ways that predictably produce systematic overspending for at least some share of the user population.

CHAPTER 4. RESEARCH METHODOLOGY

This study employs a secondary research design, meaning that no primary data were collected from human participants; instead, the paper synthesizes and critically evaluates existing peer-reviewed literature, government reports, and institutional data already published by recognized research bodies. Secondary research of this kind is well suited to an emerging topic such as BNPL because it allows a comprehensive body of otherwise scattered evidence, spanning economics, psychology, public health, and law, to be integrated into a single coherent analysis, while making transparent the provenance and credibility of each source used.

Research Design

The paper adopts a narrative systematic review design. Unlike a purely narrative review, which selects sources at the author's discretion without a documented protocol, this review followed an explicit, reproducible search and screening process, described below, to reduce selection bias. Unlike a strict meta-analysis, however, the review does not attempt to statistically pool effect sizes across studies, because the underlying literature on BNPL uses heterogeneous outcome measures, ranging from odds ratios for psychological symptoms to dollar-denominated debt balances to qualitative interview themes, that are not appropriate for direct quantitative aggregation. The narrative systematic approach therefore combines the topical breadth of a narrative review with the procedural transparency associated with systematic methods.

Data Sources

Sources were drawn from three categories. The first category comprised peer-reviewed academic journal articles indexed in scholarly databases, including studies published in JAMA Health Forum, the Journal of Behavioral and Experimental Finance, the International Journal of Bank Marketing, the Journal of Cultural Economy, Clinical Psychology Review, Marketing Science, and the Journal of Economic Literature. The second category comprised institutional and government reports, including publications and data spotlights issued by the Consumer Financial Protection Bureau, the Federal Reserve Board of Governors, the Federal Reserve Banks of Boston and Kansas City, the Reserve Bank of India, the Australian Securities and Investments Commission, and the Australian federal government's Treasury department. The third category comprised working papers and reputable secondary reporting from established outlets covering financial regulation and consumer affairs, used selectively to corroborate or contextualize primary institutional data rather than as a substitute for it.

Databases Used

Literature searches were conducted primarily through Google Scholar, PubMed, and ScienceDirect, supplemented by direct searches of institutional repositories, including the CFPB's own research and data publications page, the Federal Reserve's publications archive, ResearchGate, and the websites of national financial regulators. PubMed and its associated PubMed Central repository were prioritized for sources addressing the mental health and clinical correlates of BNPL use, consistent with their status as the principal indexing services for biomedical and public health literature, while ScienceDirect and Google Scholar were prioritized for economics, marketing, and behavioral finance sources.

Search Strategy

Search terms combined the core phrase "buy now pay later" or its abbreviation "BNPL" with secondary terms including "debt trap," "overspending," "consumer credit," "financial literacy," "mental health," "regulation," "loan stacking," "hyperbolic discounting," and "mental accounting." Searches were restricted, where filters permitted, to sources published between 2013, the year of the Richardson et al. meta-analysis anchoring the mental health discussion, and 2026, the present. Reference lists of retrieved articles were also examined for additional sources through backward citation searching, though only sources independently verified through direct retrieval were included in the final synthesis to avoid misattribution.

Inclusion and Exclusion Criteria

Sources were included if they met the following criteria: (a) the source directly addressed BNPL, consumer installment credit, or a closely related construct such as unsecured debt and mental health, or a foundational behavioral economics theory applied in the analysis; (b) the source was published by a peer-reviewed academic journal, a recognized financial or health regulator, a central bank or its regional reserve banks, or an established research or policy organization; and (c) the source's key claims could be independently verified against the original publication. Sources were excluded if they consisted of unverified promotional material published by BNPL providers, opinion content without an identifiable factual basis, or content for which full citation details could not be reliably confirmed. Where a claim originated from journalistic coverage of an underlying government report, the original report was cited as the primary source and the journalistic coverage was used only as a secondary corroborating reference.

Data Analysis Method

Because the included sources vary substantially in methodology and outcome type, this paper uses thematic synthesis rather than statistical meta-analysis. Following standard practice for narrative reviews of heterogeneous literatures, findings from each source were coded into four recurring themes that emerged inductively during review: (1) market scale and user demographics, (2) debt accumulation and loan-stacking behavior, (3) psychological and behavioral mechanisms, and (4) regulatory and policy responses. Findings within each theme

were then compared across sources to identify convergence, divergence, and evidentiary strength, with attention to whether independent sources using different datasets reached consistent conclusions. This approach allows defensible conclusions about the weight of evidence for the "debt trap" characterization while remaining transparent about the limits of causal inference that a secondary, largely cross-sectional evidence base can support.

CHAPTER 5. FINDINGS AND DISCUSSION

5.1 Key Findings

Synthesis of the reviewed literature yields several consistent findings regarding the scale, distribution, and structural risk of BNPL credit. First, BNPL usage has grown rapidly and is no longer a marginal payment method. The CFPB (2025b) documented that loan originations among the six largest providers grew from a smaller base in 2019 to 335.8 million loans worth $45.2 billion in 2023 alone, even as the pace of growth decelerated somewhat compared to earlier years. The average dollar value of BNPL borrowing per user also rose, from $745 in 2022 to $848 in 2023, an increase of fourteen percent that outpaced the growth in the number of users, indicating that existing users were borrowing more, not merely that new users were entering the market (CFPB, 2025b). Even so, the CFPB (2025b) cautioned that BNPL spending still represents only about one percent of total credit card spending volume nationally, a useful reminder that BNPL, while significant and fast-growing, remains a smaller component of aggregate consumer credit than headline growth figures might suggest in isolation.

Second, BNPL borrowers are disproportionately drawn from financially constrained segments of the population. The CFPB's analysis of 145 million loan applications found that sixty-one percent of BNPL borrowers held subprime or deep-subprime credit scores (as cited in Silverstein, 2025), while Federal Reserve survey data found that nearly three-quarters of adults earning under $50,000 annually who used BNPL did so because it was the only way they could afford a needed purchase (as cited in Checkbook.org, 2025). This demographic concentration matters because the product's risks are not evenly distributed; they fall most heavily on consumers with the least financial slack to absorb a missed payment or an unexpected stacking of obligations across lenders.

Third, and most directly relevant to the debt-trap framing, loan stacking and cross-product debt accumulation are common rather than exceptional among frequent BNPL users. More than sixty percent of heavy BNPL users held simultaneous loans from multiple providers in 2022 (CFPB, 2025a), and monthly BNPL users carried, on average, several hundred dollars more in personal loan and credit card balances than comparable non-users (CFPB, 2025a). Guttman-Kenney et al. (2023) added a further layer of concern, documenting that a share of UK consumers charge nominally interest-free BNPL installments to a credit card, converting a zero-interest obligation into revolving debt at rates far higher than the underlying BNPL terms and effectively defeating BNPL's interest-free premise. Because most BNPL lenders historically have not reported to major credit bureaus, this accumulation of obligations across providers can

remain invisible to any single lender, a gap regulators describe as "phantom debt" (Federal Register, 2025).

Fourth, the psychological correlates of BNPL use are measurable and clinically meaningful, even if causally ambiguous. Shah et al. (2025) found elevated odds of BNPL use among adults reporting depression, anxiety, and posttraumatic stress, with posttraumatic stress remaining the strongest independent predictor even after adjusting for the other two conditions and for demographic covariates. This finding is consistent with, and extends, the older but well-established association between unsecured debt and poor mental health documented by Richardson et al. (2013). Taken together, these findings suggest that BNPL functions as one node within a broader and long-recognized cycle in which financial strain and psychological distress reinforce one another, rather than as an entirely novel psychological hazard.

5.2 Benefits of BNPL

A balanced secondary synthesis must also account for the documented benefits of BNPL, which are neither trivial nor uniformly outweighed by its risks. For many consumers, BNPL functions as intended: a short-term budgeting tool that allows a predictable expense to be smoothed across a few weeks without incurring the interest charges typical of revolving credit card debt. The CFPB's own Making Ends Meet survey work found that a substantial share of BNPL borrowers used the product without any observable indication of financial distress, and that BNPL borrowers, contrary to a common misconception, generally retained access to traditional forms of credit rather than turning to BNPL as a lender of last resort (CFPB, 2023). This finding is important because it complicates any narrative in which BNPL exclusively serves consumers who have been shut out of the traditional credit system; instead, many users appear to select BNPL deliberately as a preferred short-term financing tool alongside, rather than instead of, conventional credit.

BNPL also offers a genuine financial-inclusion benefit in markets where large segments of the population remain underbanked or lack an established credit history. In several South and Southeast Asian markets, BNPL has expanded access to durable goods, electronics, and educational expenses for consumers who would otherwise struggle to obtain instalment credit through conventional banking channels, a dynamic recognized by the Reserve Bank of India's broader digital-payments strategy (Reserve Bank of India, 2022). Because most BNPL products charge no interest when repaid on schedule, the product can provide a materially cheaper form of short-term credit than either a credit card carrying a double-digit annual interest rate or a payday loan carrying a substantially higher effective rate. From a merchant's perspective, Berg et al. (2024) found that BNPL availability at checkout increases both conversion rates and average order values for retailers, a commercial benefit that has helped finance the interest-free pricing model distinguishing BNPL from most other unsecured credit products.

Finally, credit performance data reported directly by BNPL lenders to the CFPB show measurable improvement over the study period rather than deterioration, a pattern that complicates any assumption that BNPL risk is worsening uniformly over time. The share of

BNPL loans assessed a late fee fell from 5.2 percent in 2022 to 4.1 percent in 2023, and the share of loans ultimately charged off declined from 2.63 percent to 1.83 percent over the same period, the lowest level recorded across the five years of data collected by the Bureau (CFPB, 2025b; Cover et al., 2025). These improving aggregate performance metrics suggest that, at least among the largest and most established providers, credit underwriting and collections practices may be maturing even as the overall volume of lending continues to expand, a nuance often lost in more alarmist popular coverage of the sector.

5.3 Emerging Trends

Several emerging trends are likely to shape the trajectory of BNPL risk and its regulatory treatment in the coming years. The most significant is the global convergence toward formal credit regulation of BNPL, ending the product's original status as a largely unregulated fintech innovation. In Australia, the Treasury Laws Amendment (Responsible Buy Now Pay Later and Other Measures) Act 2024 extended the National Credit Code to BNPL contracts, requiring providers to hold an Australian credit licence, join the Australian Financial Complaints Authority, and conduct a modified form of responsible-lending assessment from June 2025 onward (Australian Securities and Investments Commission, 2025; Hall & Wilcox, 2025). In the United Kingdom, HM Treasury and the Financial Conduct Authority have moved to bring so-called Deferred Payment Credit under formal regulation beginning in mid-2026, following a multi-year consultation process, which will require affordability checks, standardized advertising rules, and access to the Financial Ombudsman Service for BNPL disputes (Addleshaw Goddard LLP, 2025). The European Union's revised Consumer Credit Directive similarly eliminates the previous exemption for short-term, interest-free credit products under three months, bringing BNPL fully within the Directive's responsible-lending and disclosure framework as member states transpose the directive into national law by the end of 2025 (Checkout.com, 2025). In the United States, by contrast, regulatory momentum has been comparatively uneven; while the CFPB issued guidance in 2024 clarifying that BNPL lenders must provide dispute rights and refund protections comparable to those available to credit card holders, subsequent changes in the agency's enforcement posture have introduced uncertainty about the durability of federal oversight (Federal Register, 2025).

A second emerging trend is the gradual extension of credit-bureau reporting to BNPL products, which, if fully implemented, would directly address the "phantom debt" and loan-stacking concerns documented throughout this review. Because major U.S. credit bureaus have historically excluded most BNPL loans from standard credit files, lenders extending new credit, including mortgage lenders, have had limited visibility into a borrower's existing BNPL obligations, a gap regulators have flagged as a specific risk to accurate underwriting (Federal Register, 2025). Policy discussions increasingly focus on standardizing BNPL data reporting, which would improve lender risk assessment and give consumers a clearer, centralized picture of their aggregate short-term debt.

A third trend concerns the psychological design of BNPL products themselves. Building on the gamification concerns raised by Threadgold et al. (2025), consumer advocates have begun scrutinizing interface features such as one-click checkout integration, default BNPL enrollment at the point of sale, and reward-style messaging for on-time payment, arguing these design choices merit the same scrutiny applied to other forms of manipulative digital design. Finally, the mental health findings reported by Shah et al. (2025) point toward an emerging interdisciplinary research agenda at the intersection of consumer finance and public health, likely to generate further evidence on BNPL's welfare effects in the near term.

CHAPTER 6. CONCLUSION

Summary of the Study

This paper undertook a secondary, thematically synthesized review of the academic, governmental, and institutional literature on Buy Now, Pay Later credit in order to evaluate the extent to which BNPL functions as a debt trap for consumers. Drawing on peer-reviewed research from economics, psychology, and public health, together with primary data published by the Consumer Financial Protection Bureau, the Federal Reserve System, and international financial regulators, the review examined four interconnected themes: the scale and demographic distribution of BNPL use, patterns of debt accumulation and loan stacking, the behavioral and psychological mechanisms underlying BNPL spending, and the regulatory responses now emerging across major markets.

Major Conclusions and Practical Implications

The evidence reviewed here supports a qualified rather than absolute version of the debt-trap characterization. BNPL is not, on the aggregate credit-performance data available, an unusually high-default product; charge-off and late-fee rates among the largest providers actually improved between 2022 and 2023 (CFPB, 2025b). At the same time, the product is disproportionately used by financially constrained, subprime, and psychologically distressed consumers, and its installment structure interacts with well-documented cognitive biases, namely mental accounting and hyperbolic discounting, in ways that plausibly encourage overspending and cross-product debt accumulation for at least a meaningful subset of users (Thaler, 1999; Prelec & Loewenstein, 1998). The practical implication for consumers is that BNPL is best understood as a tool whose safety depends heavily on the financial circumstances and self-regulatory capacity of the individual user rather than as a product that is inherently safe or inherently predatory. For lenders and merchants, the findings suggest that responsible design, including clearer aggregate-obligation disclosures and friction at the point of a third or fourth simultaneous loan, could meaningfully reduce stacking-related harm without eliminating the product's legitimate liquidity-smoothing function. For policymakers, the consistent finding that BNPL debt is frequently invisible to other lenders because of gaps in credit-bureau reporting represents perhaps the single most tractable regulatory target identified in this review.

Limitations of the Study

As a secondary research synthesis, this paper is constrained by the limitations of its underlying sources. Most available empirical evidence on BNPL is cross-sectional rather than longitudinal, meaning that associations between BNPL use and outcomes such as debt accumulation or psychological distress cannot be interpreted as establishing causal direction; as CFPB (2023) and Shah et al. (2025) both acknowledge, financial vulnerability and mental health symptoms may precede, rather than result from, BNPL use in many cases. The review is further limited by the concentration of available high-quality data in a small number of jurisdictions, principally the United States, the United Kingdom, and Australia, which restricts the generalizability of findings to markets with different regulatory histories, credit infrastructures, and consumer protection norms, including much of South Asia and Sub-Saharan Africa. Finally, because BNPL is a fast-evolving product and regulatory environment, some of the data synthesized here, particularly market-scale statistics, will inevitably become dated as new reporting periods are published.

Recommendations for Future Research

Future research would benefit substantially from longitudinal panel studies capable of tracking individual consumers' BNPL usage alongside subsequent changes in debt levels, credit scores, and mental health outcomes over multi-year periods, which would allow researchers to better disentangle selection effects from causal harm. Comparative cross-national research examining how differences in regulatory regimes, such as Australia's newly implemented licensing requirements versus the historically lighter-touch U.S. approach, affect consumer outcomes would also provide valuable natural-experiment evidence for policymakers. Additional research is needed on the specific interface and design features of BNPL applications, building on the gamification concerns raised by Threadgold et al. (2025), to determine which design elements most strongly predict compulsive or stacking behavior and which could be modified with minimal loss of the product's legitimate convenience benefits. Finally, given the novel findings of Shah et al. (2025), further interdisciplinary research combining clinical psychology and consumer finance methods is warranted to clarify the bidirectional relationship between mental health and short-term credit use.

Ethical and Regulatory Considerations

The ethical stakes of BNPL regulation center on a tension between consumer autonomy and consumer protection. BNPL provides real value to many users who exercise it responsibly, and restrictions that eliminate the product would remove a genuinely useful, often cheaper, alternative to credit cards. At the same time, the concentration of BNPL's risks among subprime, low-income, and psychologically vulnerable populations raises a distinct ethical concern about extending fast, low-friction credit to consumers least equipped to evaluate its cumulative cost, particularly when offered at the moment of peak purchase motivation rather than after reflection. The emerging international regulatory consensus, reflected in Australia's licensing regime, the

United Kingdom's move toward Deferred Payment Credit regulation, and the European Union's revised Consumer Credit Directive, suggests most major jurisdictions have concluded BNPL should be treated as consumer credit subject to affordability assessment and standard protections rather than as an exempt product category (Australian Securities and Investments Commission, 2025; Addleshaw Goddard LLP, 2025; Checkout.com, 2025). Mandatory, standardized disclosure of a consumer's aggregate outstanding BNPL obligations across providers, together with credit-bureau reporting sufficient to prevent phantom debt, represent minimally intrusive interventions that would preserve consumer choice while addressing the most well-documented mechanism of harm identified in this review.

Works Cited

Addleshaw Goddard LLP. (2025, August). Regulation of buy now pay later and FCA consultation. https://www.addleshawgoddard.com/en/insights/insights-briefings/2025/financial-regulation/financial-regulation-in-the-know-consumer-finance-august-2025/regulation-buy-now-pay-later-fca-consultation/

American Psychological Association. (2026). Buy now, pay later: A growing financial stressor. Monitor on Psychology. https://www.apa.org/monitor/2026/04-05/financially-stressed-digitally-tempted

Australian Securities and Investments Commission. (2025, January 17). ASIC alerts buy now pay later providers to apply for a licence under new laws. https://www.asic.gov.au/about-asic/news-centre/news-items/asic-alerts-buy-now-pay-later-providers-to-apply-for-a-licence-under-new-laws/

Berg, T., Burg, V., Keil, J., & Puri, M. (2024). The economics of "buy now, pay later": A merchant's perspective (NBER Working Paper No. 33152). National Bureau of Economic Research. https://doi.org/10.3386/w33152

Board of Governors of the Federal Reserve System. (2025). Economic well-being of U.S. households in 2024. Federal Reserve System.

Checkbook.org. (2025). Buy now, pay later is often a debt trap. https://www.checkbook.org/national/buy-now-pay-later/

Checkout.com. (2025, December 16). BNPL regulations across major global markets. https://www.checkout.com/blog/bnpl-regulations-across-major-global-markets

Consumer Financial Protection Bureau. (2023). Consumer use of buy now, pay later: Insights from the CFPB Making Ends Meet survey. https://www.consumerfinance.gov/data-research/research-reports/consumer-use-of-buy-now-pay-later-insights-from-the-cfpb-making-ends-meet-survey/

Consumer Financial Protection Bureau. (2025a, January). Consumer use of buy now, pay later and other unsecured debt. https://files.consumerfinance.gov/f/documents/cfpb_BNPL_Report_2025_01.pdf

Consumer Financial Protection Bureau. (2025b, December). The buy now, pay later market. https://www.consumerfinance.gov/data-research/research-reports/the-buy-now-pay-later-market/

Cover, J., Willis, C., & Furletti, M. (2025, December 22). What the CFPB's new buy now, pay later market report shows. Consumer Financial Services Law Monitor. https://www.consumerfinancialserviceslawmonitor.com/2025/12/what-the-cfpbs-new-buy-now-pay-later-market-report-shows/

Federal Register. (2025, June 24). Request for information regarding buy now pay later unsecured debt. https://www.federalregister.gov/documents/2025/06/24/2025-11575/request-for-information-regarding-buy-now-pay-later-unsecured-debt

Federal Reserve Bank of Kansas City. (2025). Financial constraints among Buy Now, Pay Later users. Economic Review. https://www.kansascityfed.org/research/economic-review/financial-constraints-among-buy-now-pay-later-users/

Guttman-Kenney, B., Firth, C., & Gathergood, J. (2023). Buy now, pay later (BNPL)…on your credit card. Journal of Behavioral and Experimental Finance, 37, Article 100788. https://doi.org/10.1016/j.jbef.2023.100788

Hall & Wilcox. (2025). Buy-now pay-later reform: What you need to know. https://hallandwilcox.com.au/news/buy-now-pay-later-reform-what-you-need-to-know/

Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. Journal of Economic Literature, 52(1), 5–44. https://doi.org/10.1257/jel.52.1.5

NLS Forum. (2025, April 15). Digital lending in India: Analysis and implications. https://forum.nls.ac.in/ijlt-blog-post/digital-lending-in-india-analysis-and-implications/

Prelec, D., & Loewenstein, G. (1998). The red and the black: Mental accounting of savings and debt. Marketing Science, 17(1), 4–28. https://doi.org/10.1287/mksc.17.1.4

Relja, R., Ward, P., & Zhao, A. L. (2024). Understanding the psychological determinants of buy-now-pay-later (BNPL) in the UK: A user perspective. International Journal of Bank Marketing, 42(1), 7–37.

Reserve Bank of India. (2022, September 2). Guidelines on digital lending (Circular No. DOR.CRE.REC.66/21.07.001/2022-23).

Richardson, T., Elliott, P., & Roberts, R. (2013). The relationship between personal unsecured debt and mental and physical health: A systematic review and meta-analysis. Clinical Psychology Review, 33(8), 1148–1162.

Shah, H., Prus, E., Castrucci, B., Galea, S., & Ettman, C. K. (2025). Poor mental health and the use of buy now, pay later loans. JAMA Health Forum, 6(12). https://doi.org/10.1001/jamahealthforum.2025.5620

Silverstein, S. (2025, January 14). Buy now, pay later users pile on debt, CFPB finds. Payments Dive. https://www.paymentsdive.com/news/cfpb-buy-now-pay-later-bnpl-use-debt-research-report/737285/

Stavins, J. (2024). Buy now, pay later: Who uses it and why (Current Policy Perspectives No. 24-3). Federal Reserve Bank of Boston.

Strotz, R. H. (1956). Myopia and inconsistency in dynamic utility maximization. Review of Economic Studies, 23(3), 165–180. https://doi.org/10.2307/2295722

Thaler, R. H. (1985). Mental accounting and consumer choice. Marketing Science, 4(3), 199–214. https://doi.org/10.1287/mksc.4.3.199

Thaler, R. H. (1999). Mental accounting matters. Journal of Behavioral Decision Making, 12(3), 183–206.

Threadgold, S., et al. (2025). Buy now, pay later technologies and the gamification of debt in the financial lives of young people. Journal of Cultural Economy, 18(1), 52–67.

Treasury Laws Amendment (Responsible Buy Now Pay Later and Other Measures) Act 2024 (Cth) (Austl.).

Zhang, C. Y., & Sussman, A. B. (2018). Perspectives on mental accounting: An exploration of budgeting and investing. Financial Planning Review, 1(1–2), Article e1011. https://doi.org/10.1002/cfp2.1011

Macrophage Plasticity and ECM Remodelling and Solutions in Fibrosis
By Sakthivel Heamish Kanha

Keywords Fibrosis, Macrophages, ECM, Treatment, novel therapeutics

ABSTRACT

Fibrosis is a chronic and progressive disease that occurs in many different organs. Excessive extracellular matrix deposition results in organ stiffening and eventually, organ failure. For a healthy person, the wound healing process is rigidly controlled with the inflammatory, proliferative, resolution and remodelling phase happening sequentially and with coordination. However, repeated injury can cause these mechanisms to become dysregulated with continued signalling from cytokines such as TGF-β, macrophage polarisation to M2 like state, myofibroblasts resisting apoptosis and the resolution and remodelling phase of ECM components being impaired. The pathways such as the PI3K -AKT, JAK-STAT, NF-κB all contribute to myofibroblast intrinsic and extrinsic apoptosis resistance. Without the proper mechanisms in place, a continued deposition of ECM components such as collagen result in tissue and organ stiffening.

Fibrosis persists when macrophage-derived profibrotic programs and apoptosis resistance lock myofibroblasts in a survival state. Thus, while there is no current cure for fibrosis, therapies that interrupt TGF-β signaling and/or reprogram macrophages represent the most promising route to disease modification.

Anti-fibrosis treatments such as nintedanib and pirfenidone help decrease the progression of fibrosis through targeting intracellular kinase receptors, TGF-β signalling and acting as an antioxidant. Emerging biologics such as fresolimumab, a human monoclonal antibody that targets and neutralises all TGF-β isoforms. Novel therapeutic strategies involving nanoparticles and hydroxychloroquine aim to change the polarisation of M2 macrophages to M1 macrophages.

1.0 Introduction

Fibrosis causes an over deposition of scar or connective tissue in an organ/tissue leading to the stiffening and impairment of organ function over time[1]. The malfunctioning of the human healing process leads to excess deposition of extracellular matrix (ECM) and gradual loss of organ function due to stiffening[2]. It can be triggered by persistent inflammation as a consequence of repeated injury[3].

Process in healthy person overview

A number of processes take place when a tissue is injured. Firstly, inflammation happens causing swelling and the area to become hot[4] . The neutrophils and macrophages help engulf pathogens and so are recruited to site on injury where they secrete specific cytokines such as IL-1, IL-6, IL-8, and IL-12[5]. The macrophages then transition from M1(inflammatory) to M2 stage releasing TGF-β as a result of interaction with cytokines[6]. In addition, the fibroblasts

increase in number, differentiating into myofibroblasts, before creating ECM [7] The fibroblasts create type 3 collagen. ECM remodelling then takes place with the MMPs breaking down the ECM and the TIMPS inhibit the activity of the MMPS[8]. At the same time, the fibroblasts (myofibroblasts) and immune cells undergo apoptosis[9].

Fibrosis

Fibrosis is when there is excessive ECM deposition such as collagen within the connective tissue space[10] resulting in tissue stiffening. In patients who have fibrosis, the normal wound healing sequence is affected[11]. The myofibroblasts resist apoptosis, causing them to continually produce ECM components and the macrophages are polarised into a persistent M2 state that don't revert back to their other phenotypes.[12] MMPS regulate the amount of the ECM deposition by degrading it[13]. In fibrosis TIMPS are upregulated (produced more) causing action of MMPs to fall and for surplus of ECM deposition to remain[14].

Significance of fibrosis

Fibrosis can lead to negative health effects because the thickened stiffening of the tissue causes permanent scarring. The inability for tissue regeneration, changes organ shape and prevents it from carrying out its function. Eventually makes it very serious as there is a risk of organ failure .These are the types of fibrosis:

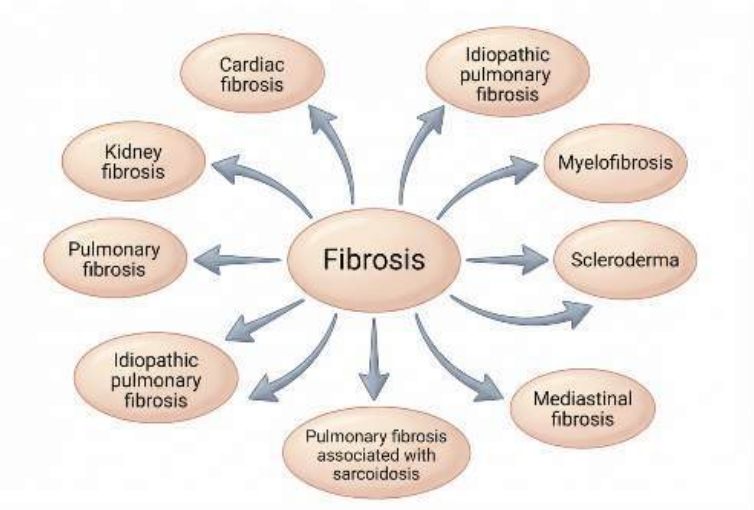

Fig 1: Different types of fibrosis
Citation:Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

An estimate of around 13-20 in every 100,000 people worldwide[15], have pulmonary fibrosis. Around 6-7% of the global population has liver fibrosis[16]. Fibrosis is a serious issue and if left untreated can lead to permanent impairment of organ function and poses a high risk to patients.

Normal Wound Healing

Persistent injury leads to errors in the normal healing process due an overproduction of scar tissue.The process becomes unregulated with mechanisms such as inflammation, fibroblast activation and ECM deposition not being controlled[17]. Understanding the normal healing process provides a comprehensive understanding on how fibrosis develops and differs from normal wound healing mechanisms.

2.1 Inflammatory phase

Inflammation occurs as an immediate response when the body is injured[18]. It helps protect the body from pathogens. Inflammation happens for a duration of around a few days. The first step is short vasoconstriction to lower blood loss into surroundings, followed by a phase of vasodilation which takes place to bring other immune cells to kill pathogens[19]. Platelets are also activated and form an initial blockage. However a proper mesh is formed by fibrin. Fibrin is derived from fibrinogen to form a mesh. The process by which this happens is because of the secretion of thrombin (Factor IIa) which is an enzyme causing fibrinogen to polymerise into fibrin[20][21]. Thrombin is created through the prothrombinase complex that converts prothrombin into thrombin via cleavage of prothrombin in certain conditions[22]. The overall process causing blood clotting to form is known as the coagulation cascade[23]. Mesh formation is also important to prevent bacteria from entering into the body[24].

At the same time, cytokines such as PDGF, TGF B are released by the platelets[25]. The platelets have alpha granules that release these cytokines once the platelets are activated[26]. The immune cells are alerted and go to the inflammatory site.

Immune cells such as neutrophils release reactive oxygen species, enzymes and conduct phagocytosis on bacteria[27]. Neutrophils are able to kill bacteria as they use the oxygen to generate the reactive oxygen species and through a process called respiratory burst, HOCI is formed that can kill bacteria[28].

Afterwards, monocytes differentiate into M1 macrophages that further release cytokines to continue the inflammation before it transitions into M2 macrophages that releases cytokines such as TGF-β, IL-10 to counter inflammation and cause fibroblast activation[30].

2.2 Proliferative phase (fibroblast & ECM synthesis)

From the TGF-β , fibroblasts start to synthesize collagen type III, fibronectin, which is an ECM component[31]. Fibroblasts have their own receptors and when the TGF-β binds to those receptors, a process known as SMAD starts[32]. Once the TGF-β binds to receptor TβRII and activates it, SMAD2 and SMAD3 are phosphorylated which join SMAD4 and then go to the nucleus to modify the gene expression[33]. This leads to increased myofibroblasts and ecm component deposition and production[34]. It is known as the canonical TGF-β/SMAD pathway[35]. The type 3 collagen helps provide temporary support to the cell[36].

Another important part of inflammation is angiogenesis. It is the formation of new blood vessels from ones that already exist[37]. Angiogenesis tends to lengthen the inflammation process and over produce VEGF-A[38]. VEGF-A signaling to VEGFR-2 (resulting from some different pathways) causes endothelial cells to proliferate, and form sprouts allowing lumens to be constructed[40] which become new blood vessels after they become stable. In a healthy individual VEGF-A falls after the vessels are grown but in fibrotic cases it stays for longer that can contribute to increased scarring[41].

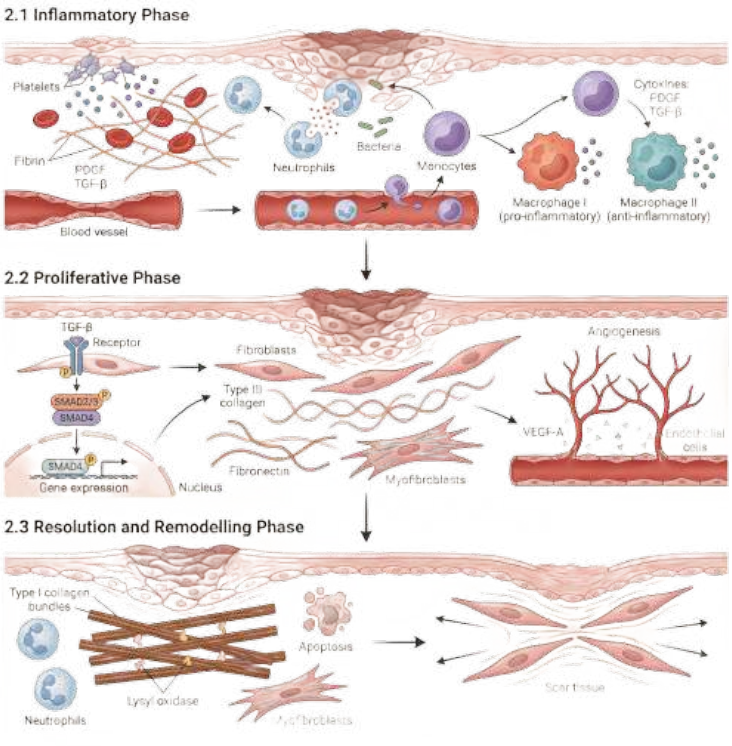

Fig 2: Diagramatic overview of the three processes mentioned, inflammation, proliferation, resolution and remodelling

Citation: Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

2.3 Resolution and remodelling

This stage is around weeks to months. The trigger for this stage is when there is sufficient ECM deposition, resolution and remodelling needs to take place[42]. The ECM contains type 3 collagen which is more of a short term solution. To fix this the tissue resorts to type 1 collagen fibrils which bundle together to form a large fibre[43]. Type 1 collagen is stronger and durable. They are crosslinked through lysyl oxidase and replace more and more of the type 3 collagen till it occupies the entire space[44].

Neutrophils undergo apoptosis once they have cleared the area of pathogens and any infectious particles[45]. The myofibroblasts also undergo apoptosis in order for excessive scarring to not happen[46].

The Myofibroblasts under the influence of TGB-F differentiate from fibroblasts and have a-sma which allows it to contract the edges of the ecm to help close any gaps and pull it to form a densely packed structure to decrease the area of the wound[47][48].

Thus the resolution and remodeling stage comes to an end.

3.0 Fibrosis Pathogenesis

Fibrosis can be pathological when the classical inflammatory fibrosis resolution pathway is altered and prolonged. For instance if myofibroblasts don't undergo apoptosis, ECM components continue to be deposited at a rate greater than ECM degradation, there is continued signalling then it indicates fibrotic development.

3.1 Persistent injury

During fibrosis macrophages become activated as M2 (pro fibrosis). This can happen if IL-13 and IL4 cytokines are present which are mainly responsible for M2 macrophage state[49]. In addition the TIMPS are overexpressed and since they inhibit activity of MMPs who degrade the ECM, there is lower removal of ECM. The Myofibroblasts have BAX and BIM proteins that cause cells to die while also having BLC-2 and Bcl xl that prevent cells from dying[50]. During fibrosis, the myofibroblasts resist apoptosis (programmed cell death) due to increased BLC-2 and Bclxl. From the stiffness of the tissue, the PI3K, YAP/TAZ are activated along with TGF-β and PDGF which increase the expression of BLC-2 and Bcl xl levels thus preventing cells from dying[51]. In addition, Fas or TNF that signal which are meant to activate caspase-8 and caspase-10 that initiate apoptosis are ineffective as the number of fas receptors may reduce so signals never reach them[52]. Fas receptors fall because they are controlled to gene expression and the PI3K /Akt pathway suppresses them.

3.2 Macrophage and T-cell plasticity

Immune cell plasticity: what specifically causes them to undergo change in states

Macrophages become M1 under IFN-γ and TNF-α while Macrophages become M2 under IL-4, IL-10, IL-13, and TGF-β[53]. The cytokines bind to receptors on the macrophage surface resulting in M1 or M2 genes being expressed as a result of signal transduction[54]. Macrophage can change phenotypes from M1 to M2 (yet they may also exist in hybridised states)

Reasons behind why macrophages become persistent M2:

Repeated tissue injury causing M2 macrophages to persistently be present for repair . Normally, those cytokines that cause the switch of state are only released during the repair phase and so once the wound closes, the cytokines signalling would stop so the change of state occurs naturally. Yet since its repeated tissue injury, the cytokines are continually released[55] triggering the macrophages to not switch states.

IL-4 and IL-13 are cytokines produced by TH2 cells, eosinophils and basophil cells when tissue is consistently damaged[56]. DAMPS also play a role as they are released when a cell dies which explains why macrophages resist apoptosis during continued infection of the same site[57]. DAMPS can be identified through the cellular receptors such as pattern recognition receptors (PPR) which then cause stress response mechanisms to be increasingly activated resulting in a positive feedback loop of inflammation and injury of the tissue [58]. This causes macrophages to be present and resist apoptosis.

3.3 The role of ECM remodelling in fibrosis

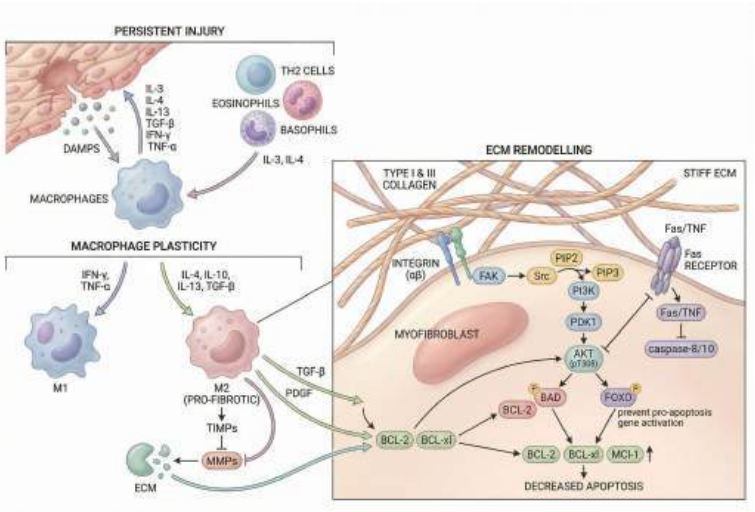

Fig 3: Detailed diagrams behind the processes involved in persistent injury, macrophage plasticity and ECM remodelling

Citation: Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

ECM is an interconnected network of proteins, and other components that surround, and maintain shape for cells and tissues [59]. In fibrosis, specifically type 1 and type 3 collagen is deposited in excess. The tissue stiffening is detected through integrins. Integrins are the main receptor proteins for which the cells bind and respond to the extracellular matrix[60] . They join the cytoskeleton inside the cell to the ECM proteins to enable it to form a grip. The integrin binds to the ligand(a molecule that is complementary to receptor protein) forming a linkage that leads to local adhesions to form[61]. The cytoplasmic tail of the β integrin subunit binds to the talin, α-actinin, and filamin.[62] The cross linking proteins help connect the integrin to actin filaments in the cell cortex[63][64].This triggers biochemical processes within the cell. ,

Firstly, FAK is activated (Focal Adhesion Kinase). FAK binds to the talin and goes to that location through the Focal adhesion targeting. It creates sites for other signaling proteins and this recruits SRC kinases leading to FAK-Src complex[65]. PI3K is activated causing the AKT pathway. The activated PI3K k converts PIP2 to PIP3 which leads to the phosphoinositide-dependent kinase-1 (PDK1) influencing the Akt[66]. The AKT binds to the PIP3 and is activated through phosphorylation of T308 by PDK1[67][68]. AKT reduces

apoptosis by causing multiple events to happen. BAD which causes apoptosis is phosphorylated by the AKT[69]. The BAD detaches from Bcl-2 causing Bcl-2, Bcl-Xl and MCl which results in the survival[70]. The AKT phosphorylates FOXO proteins and prevents them from activating the pro apoptosis genes. This directly prevents apoptosis from happening and causes increased fibrosis.

For healthy individual, eventually in the wound healing phase, the integrin activation is reduced as ecm tension falls which causes the AKT signaling to get weaker resulting in myofibroblasts undergoing apoptosis[71] and the MMPS help degrade the excess ecm.However during fibrosis, that can be caused through repeated injury the aforementioned cycle causing the myofibroblasts to survive is repeated and ecm remains stiff causing the integrins to constantly be activated.

4.0 Molecular Regulation of Apoptosis in Myofibroblasts

The prolonged survival of myofibroblasts causes chronic scarring and increased deposition of the extracellular matrix which leads to fibrosis[72]. It is important to understand these processes so that specific approaches can be developed at a molecular level to regulate myofibroblast survival and cause their apoptosis at a correct time to reduce the progression of fibrosis.

4.1 Myofibroblast activation & survival and apoptosis

Myofibroblasts are very important in fibrosis because they lead to deposition of collagen that can cause tissue stiffening. They undergo apoptosis after their job in the wound healing stage is completed however in fibrosis they don't undergo apoptosis.There are two main pathways the extrinsic and intrinsic apoptosis

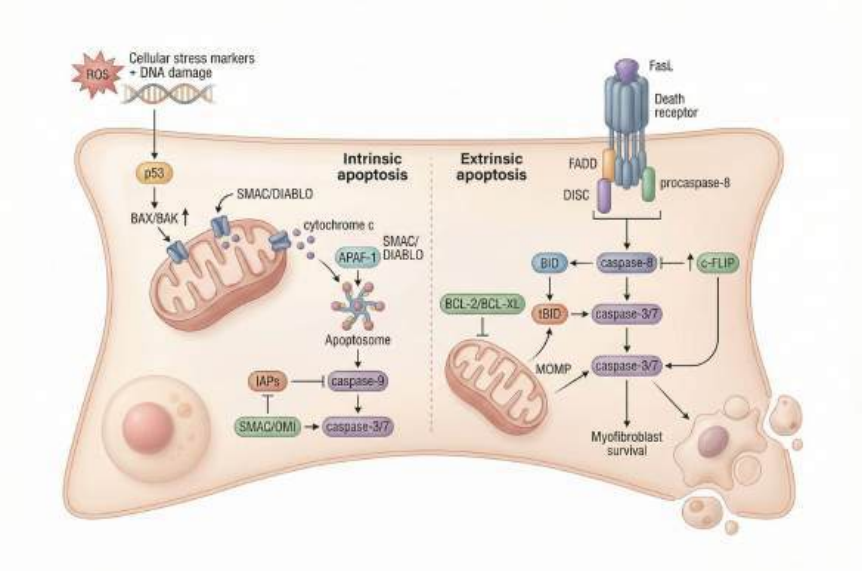

Fig 4:Intrinsic and extrinsic apoptosis pathways
Citation: Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

4.2 Intrinsic pathway

The intrinsic pathway is triggered through a cellular injury. This can be triggered through DNA damage and oxidative stress along with other reasons[73].

- **P53 protein**

A crucial protein called p53 induces the intrinsic apoptosis pathway[74]. When internal damage has happened, p53 ubiquitylation is suppressed which is critical as that is what keeps p53 levels low[75]. MDM2 is unable to degrade the p53 as ATM and CHK2 has phosphorylated the p-53 making it difficult to bind to MDM2[76]. This allows for p53 levels to rise steadily. Protein p53 is important because at high levels it can transcriptionally activate prop apoptotic Bcl2 members[77]. At the same time, it represses anti-apoptotic family members. The Bcl-2 are a group of proteins that are some pro apoptotic while some are antiapoptotic[78].

- **MOMP**

Before continuing, it is important to look at MOMP (Mitochondrial Outer Membrane Permeabilisation). If the membrane becomes permeable, it releases cytochrome c which then binds to Apaf1 and creates a complex with Casapase-9 under specific circumstances such as ATP being present[79]. Caspase 9 causes other caspases to be activated leading to DNA fragmentation and the cleavage of other cellular proteins[80]. The permeabilisation also causes the release of OMI/HtrA2 which cleaves multiple IAPs (Inhibitor of apoptosis proteins) which is irreversible and is very efficient in disabling them[81][82]. It does so through destroying the specific peptide bonds of the IAP and thus causes its shape to change and for it to lose its function[83]. IAP's can prevent apoptosis by inhibiting caspases that can cause cell death. They can do this by using their BIR domains to bind to the catalytic sites (these sites are what allows caspases to cleave the substrates and cause death of cells) and attach ubiquitin molecules to caspases[84][85]. It also controls the activation of NF-κB whose role will be explained later[86]. That is why the OMI/HtrA2 role in blocking IAP's is very important. Also, as a result of the membrane becoming permeable SMAC/DIABLO is released that also inhibits IAPS by the SMAC binding to the BIR domains of the IAP and thus allowing the action of caspases[87].

- **Bcl-2 family**

Of the Bcl-2 Members, BAX, BAK and PUMA are proapoptotic[88] as they promote MOMP while Bcl-2 and Bcl-XL are anti-apoptotic. The MOMP is regulated through the Bcl-2 proteins as the BAX and BAK proteins of the family help form the pores that lead to the ensuing processes causing the cell death.

Normally, for myofibroblasts the intrinsic pathway is activated through cellular stress which can be caused after the withdrawal of survival signals and ROS accumulation[89]. However in fibrosis, the main reason behind why the myofibroblasts resist the aforementioned processes that cause apoptosis is because of the persisting survival signals.

- **TGF-β1 signaling**

The persistent TGF-β1 signaling in fibrosis binds to the receptor on the surface of myofibroblast[90]. It activates the SMAD dependent pathways which then goes to the nucleus and cause the transcription of anti apoptotic genes[91]. It does through binding to the DNA and promotes the target genes that leads to production of Bcl-2 and Bcl-XL through increasing their transcription[92].

- **NF-κB pathway**

The NF-κB pathway which can be triggered IL-1 (pro fibrotic cytokine) is activated through the phosphorylation of the IκB kinase into IKB proteins which is then degraded[93]. The NF-κB goes into the nucleus and usually the IKB inhibits it but since the IKB is degraded the NF-κB can't be stopped from going into the nucleus[94]. The NF-κB binds to the DNA and causes the transcription of anti apoptotic genes such as Bcl-2 and Bcl-XL and upregulates IAPs which as mentioned before inhibits caspase[95].

- **JAK STAT pathway**

Also, in some cases TGF-β can activate the JAK STAT pathway[96]. Once activated, JAK phosphorylates STAT[97]. The phosphorylated STAT then moves into the nucleus and causes transcription of the anti apoptotic genes[98]. RNA polymerase II starts the transcription of the anti apoptotic genes by making mRNA which is translated in ribosomes in the cytoplasm into making anti apoptotic proteins such as Bcl-2[09]. The myofibroblasts resist apoptosis and live longer leading to increased deposition of collagen and ECM causing fibrosis.

4.3 Extrinsic pathway

The extrinsic pathway happens through FAS ligand which binds to the cell surface receptors[100]. Caspace-8, an initiator caspace, is activated[101]. The FAS domain protein then binds with procaspase 8 which is then activated by forming a death-inducing signaling complex (DISC)[102]. Caspace 8 causes the caspase 3 and caspace 7 to form resulting in dna fragmentation and the destruction of the cell[103]. Or it can cleave the BID into tBID causing the MOMP aforementioned process to happen[104]. It must be noted though, that c-FLIP is an antiapoptotic protein that can prevent the activation of caspase 8 through joining the DISC[105]. In fibrosis, c-FLIP is upregulated through the PI3K pathway and the NF-κB pathway that is activated during the TGF-β cytokines[106].

Mechanical causes of ECM stiffness such as activated integrins and focal adhesion signaling leads to rise in NF-κB and PI3K K pathways[107]. This causes lower expression and response to those particular ligands such as Fas-L so the fibroblasts undergo lesser apoptosis and increased cell survival resistance[108].

The DISC process of formation is controlled as if it isn't then it could lead to premature cell death. After Fas ligand binds to Fas receptor, the FADD causes the procaspace 8 to come to the receptor[109]. This allows for the dimerisation and cleavage before activating caspase

8[110]. However it is important to understand the role played by c-FLIP. As it doesn't have catalytic activity, it can bind to DISC and if its levels are high, then the DISC process is disrupted and caspase 8 activation is prevented111][112]. Thus c-FLIP plays a very prominent role in the cells survival as if it is upregulated then it can cause the cell to live longer by competing for the binding with the DISC.

In fibrosis, c-FLIP is upregulated. The resistance to apoptosis of the fibroblasts that become myofibroblasts causes them to be ever present in the tissue and continue to produce collagen III which leads to excess ecm deposition. Normally , TRAIL or FAS ligands send signals that should lead to caspase 8 activation through forming a complex with FADDosome. cFLIP is catalytically inactive but it is complementary to the shape of the adaptor proteins[113][114]. So, cFLIP can bind to it thus occupying the binding sites that should have been used for the creation of caspase 8.

Finally, the extrinsic pathway is crucial for leading to the apoptosis of the cell however during fibrosis, the upregulation of c-FLIP leads to the cell to resist apoptosis by preventing caspase 8 activation by preventing the dimerisation of procaspase 8 molecules at the DISC as it binds to the sites present because of its shape. This has led scientists to try finding therapeutic approaches to prevent the upregulation of c-FLIP as if they can control it, then apoptosis would happen which would lead to myofibroblast death and reduced ecm deposition. The solutions will be explored in the next section.

5.1 Nintedanib

It is a clinical drug used to treat idiopathic pulmonary fibrosis[115]. It is a tyrosine kinase inhibitor and inhibits receptor tyrosine kinases (RTKS)[116]. The overall effects of nintedanib include: lower fibroblast activation and increased fibroblast apoptosis while also potentially reducing angiogenesis[117]. As mentioned earlier in the paper, all these help reduce fibrosis.

Nintedanib is able to do this by binding to the intracellular kinase domains of RTKS and so occupies the ATP binding site thus preventing the autophosphorylation of tyrosine that could have happened from the kinase[118][119]. This overall process is called competitive inhibition because in this case the molecule competes with the ATP[120]. It reduces the free sites for binding and lowers the chance of autophosphorylisation[121]. Autophosphorylation causes phosphotyrosine sites to form, allowing for intracellular adaptor proteins to attach to RTKS[122]. Other downstream effectors that enhance transcription factors are initiated, leading to increased cell survival and proliferation[123]. The downstream pathways that are significant include the PI3K -AKT pathway contributing to fibrosis through cell survival [124] .Thus the chance of this process is lowered through the competitive binding of the drug molecule.

Nintedanib is a multikinase inhibitor which explains why the molecule is able to simultaneously inhibit multiple processes from happening such as the PDGFR and VEGFR[104]. Also without autophosphorylation RAS–RAF–MEK–ERK is blocked because the adaptor proteins cannot latch onto the surface SOS is not coming to the site so the RAS is inactive[125].

The entire process is prevented from happening, so the transcription factors and signals that would have been expressed are stopped from being exhibited or transcripted.

However nintedanib is not a direct cure[126]. Yet it decreased the annual decline rate by 51% for FVC vs placebo[127]. In relation to fibrosis, the mortality rate of people who had fibrotic treatment vs those who did not were studied over a 5 year period. The results were that those with the Nintendanib treatment had a 55% lower mortality rate[128]. The average survival was also 11.6 compared to without treatment of 3.7 years[129].

Based on the Weibull distribution, mean (95% CI) survival was estimated as 11.6 (9.6, 14.1) years in nintedanib-treated patients and 3.7 (2.5, 5.4) years in placebo-treated patients[130].

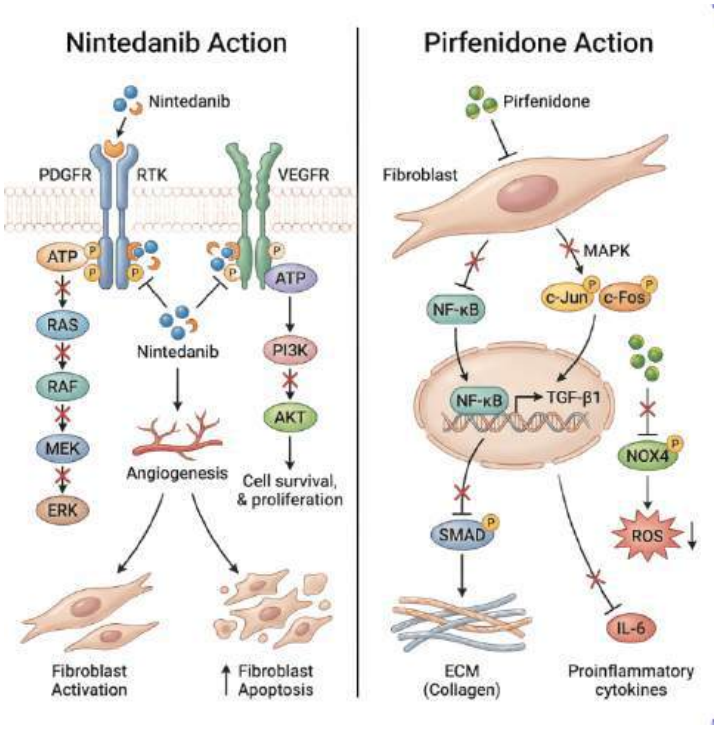

Fig 5: Nintenanib and pirfenidone mechanism of action

Citation: Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

5.2 Pirfenidone

Pirfenidone is an oral drug that is used to treat Idiopathic pulmonary fibrosis and reduces the fibrotic progression and disease progression through controlling the fibrotic signaling and other pathways that lead to inflammation[131].

Pirfenidone downregulates the TGF-β1 [132]which as mentioned previously is pro fibrotic and drives fibroblast activation. Pirfenidone also downregulates the SMAD phosphorylation thus reducing the deposition of extracellular matrix components such as collagen[133]. It also directly suppresses the other pro-inflammatory cytokines like IL-6 thus preventing the progression of fibrosis and lowers the reactive oxygen species[134]

Pirfenidone is able to achieve the above mentioned things through a variety of processes. Firstly, TGF-β1 transcription is due to NF-κB and AP-1 both of which activate TGF-β1[135]. pirfenidone interferes with this process by inhibiting the NF-κB nuclear translocation through preventing its phosphorylation of IKBA[136]. It also downregulates AP-1 activity . The chances of AP-1 binding to DNA and transcript TGF-β1 is enhanced by MPK-pathway which phosphorylates the c-Jun and c-Fos(AP-1 is made of these proteins)[136][137]. pirfenidone acts by inhibiting the MPK phosphorylation [138]thus reducing AP-1 DNA complex formation which leads to lower transcription of TGF-β1 and so lower fibrotic progression. In addition, pirfenidone also reduces NOX4 expression in fibroblasts which leads to lower ROS levels which is crucial in the fight against fibrosis[139] because ROS increases the development of fibroblast to myofibroblast that leads to fibrosis. While ROS is important for immune cells and is important for cellular functions, in fibrosis the ROS is in excess due to lower antioxidants and so lead to oxidative stress[140] and increased proliferation of fibroblasts into myofibroblasts. This is why pirfenidone acts as an antioxidant and keeps ROS levels in check.

The ASCEND trial data can be used to assess the clinical relevance of Pirfenidone. In this trial, patients were randomly chosen to be in placebo vs pirfenidone and it was found out that there was a 47.9%relative reduction in the ratio of patients with ≥10% absolute decline in predicted FVC[141].

Another research study pooled the participants of CAPACITY and ASCEND study to draw conclusions. It was found out that In the first year, pirfenidone lowered the ratio of the people in study with a ≥10% decline in predicted FVC or death by 43.8% and elevated the ratio of patients with no decline by 59.3%[142] . This showed that pirfenidone is effective in increasing survival chances for patients with IPF and reduces disease progression.

However it must be noted that side effects that affect the patients consuming this drug include nausea, diarrhea reactions however they are frequently within tolerable limits and only a minority had to undergo dosage reduction. There was around 18.5% dosage reduction for full dosage patients compared to 5.8% for placebo effect[143].

It is given usually in dosages of 2403 mg per day orally and was FDA approved in 2014 for ESBRIET[124].

Comparison between Nintedanib and pirfenidone

Nintedanib is able to block a wider range of fibrotic pathways compared to pirfenidone which acts on a lesser number of pathways. For instance, Nintedanib acts to work against fibroblast survival, angiogenesis, and fibroblast multiplication whereas pirfenidone acts on TGF-β signaling mainly. Furthermore, nintedanib has a slightly lower FVC decline rate compared to Pirfenidone. For nintedanib, the annual rate of change was −114.7 ml with nintedanib compared to 239.9 ml[145] with placebo (difference of 125.3 ml) which is roughly around 50%. For Pirfenidone, the FVC decline from baseline to 12 months was −109.0 (13.6) mL for pirfenidone vs −207.5 (13.7) mL for placebo, which shows a difference of 98.5 (17.5) mL[146]. This shows nintedanib is slightly better due to lower mean annual decline rate. In

addition, based on side effects, they are both similar as they don't have any severe side effects despite nintedanib leading to gastronomical issues such as diarrhea. Both nintedanib and pirfenidone slow FVC decline although direct comparison between these compounds is limited. Thus, the choice often depends on tolerability, comorbidities, and clinician preference although, through the data presented here, Nintedanib has seemed to be a better option.

5.3 Emerging biologics (anti-TGF-β?)

Fresolimumab is a human monoclonal antibody that is in early phase of clinical trials. It targets all TGF-β isoforms and neutralises all the forms[147]. To understand how fresolimumab works, it is important to understand how TGF-β forms. TGF-β is secreted and is dormant (inactive form) in the extracellular matrix with the LAP (latent associated peptide) helping it remain inactive[148]. The LTBP (latent TGF-β binding protein) helps keep the TGF-β attached to the extracellular matrix[149]. TGF-β can be activated through MMPs that cleave the LAP thus allowing TGF-β to become active[150], integrins can also help free TGF-β by exerting mechanical tension on the LAP to pull it and free the TGF-β[151]. The TGF-β is then able to bind to TGF-β receptor II leading to phosphorylation of SMAD complexes and gene transcription[152].

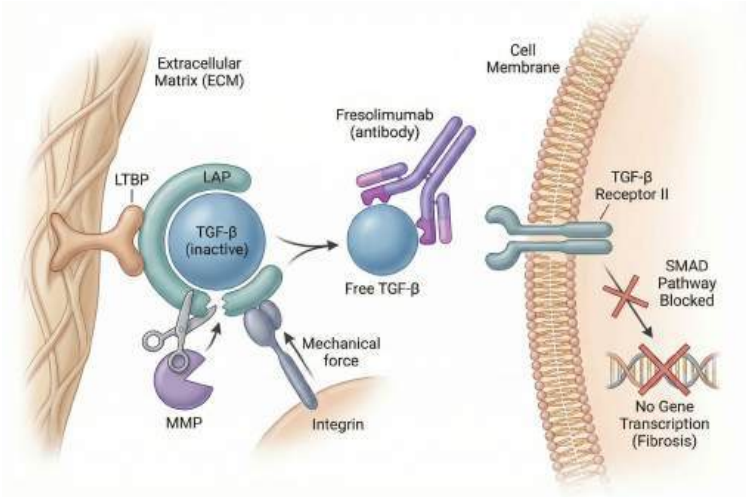

Figure 6: Action of Fresolimumab in inhibiting fibrosis.
Citation: Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

Now fresolimumab has a shape complementary to the free TGF-β[153]. The TGF-β has binding sites for it to bind to for instance the TGF-β receptor II. These binding sites are called epitope and fresolimumab has a complementary shape to these sites. This allows it to bind to the free TGF-β and prevents it from binding to TGF-β receptor II thus blocking the initiation of the SMAD pathways that lead to fibrosis.

In terms of clinical trials, fresolimumab has been used in a single dose study involving 16 patients. The patients had primary FSGS and got a dose of fresolimumab for 120 days[134]. The results revealed that the fresolimumab was safe as the only negative thing was a patient getting a pustular rash but apart from that it was well received by the body[154]. In terms of results related to anti TGF-β, there the median was a slight reduction of 1.2mg/mg in proteinuria with 3 patients having an average decline of 3.6mg/mg[154]. This indicates that TGF-β was inhibited. TGF-β encourages podocyte injury, which downregulates nephrin that causes apoptosis of podocytes, meaning higher albumin permeability leading to higher proteinuria[155]. However, since proteinuria decreased, there is a correlation between fresolimumab inhibiting TGF-β signaling as this monoclonal antibody as the usual cascade leading to higher proteinuria was stopped.

However it must be noted that the sample size for this experiment was only 16 patients, and the result of slight median decrease of 1.2mg/mg could be due to random fluctuations rather than fresolimumab's success as anti-tgf beta, there was also no control group so results cant be compared against and the patients received only a single dose which is insufficient as some of this could have been lost in the body or did not reach the target organ. So while a direct correlation between fresolimumab inhibiting TGF-β signaling is unclear, it can be concluded that it is safe for administration and more trials need to be done to see if there is a potential link between fresolimumab and anti TGF-β.

5.4 Approaches targeting macrophage re-programming

Reprogramming macrophages can help inhibit the growth of fibrosis. Macrophages have a double role for fibrosis. Starting with M1 macrophages, these macrophages generally have an anti fibrosis role[156]. Initially though, they have proinflammatory properties by secreting TNF-α, IL-1β and IL-6 which are pro fibrosis[157] as these chemokines activate fibroblasts by triggering NF-κB pathways. However, M1 macrophages play a role in degrading extracellular matrices through MMPs[158][159] which indicate that they help resolve fibrosis. Between the anti fibrosis and fibrosis role of the M1 macrophages, the M1 macrophages role in resolving fibrosis outweighs the indirect contributions that may cause fibrosis. This means that M1 macrophages can be viewed as antifibrosis[160].

M2 macrophages play a more dominant pro-fibrotic role as they secrete TGF-β1 which causes fibroblast to myofibroblasts transition[161], and secrete other growth factors that lead to excess deposition of ecm. Furthermore they also secrete IL-13 and IL-10 and produce TIMPS which decreases the removal of ecm by inhibiting the activity of MMPs[162][163]. Thus, M2 macrophages are predominantly profibrosis.

M1 macrophages can switch to M2 macrophages in a process known as macrophage polarization, this drives fibrosis rather than resolving it[164]. Instead it would be useful to consider pathways that cause M2 macrophages to switch to M1 so that the production of pro-fibrosis signals are reduced and the fibrosis growth can be stopped.

A preclinical solution has been developed on macrophage reprogramming, by using nanoparticles that contain IRF5 MRNA and IKKB mrna with a targeting system[165]. The nanoparticles then enter into the body and specifically into the M2 macrophages which results in the ribosomes within the macrophages translating the IRF5 and IKKB[166]. Those are proteins that cause the M1 genes to be expressed, while M2 gene expression stops because they are able to stop the M2 transcription factors from binding to those particular gene promoters causing the macrophages to switch polarisation from M2 to becoming M1[146].

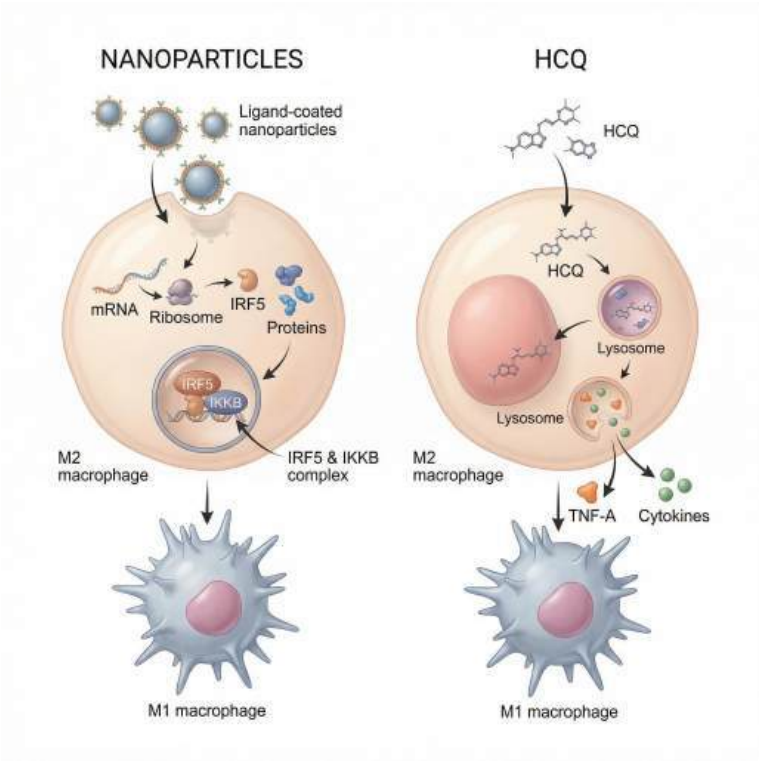

Figure 7: Mechanism of action for nanoparticles and HCQ
Citation: Illustrae.co. (2025). *Illustrae - Professional Scientific Illustration Platform*. [online] Available at: https://illustrae.co/.

There has been a study done involving in-vivo mice experiences where the control mice group had initially 0.5% M1 like population while afterwards the results showed an increase to 10.2% of M1 population as a result of this method[167]. Yet while there is proof of concept for the mice population, this type of method has not been done for humans and since the IRF5 lasts only lasts for a few days[147], it may mean that it would need a number of repeated dosages and there may be problems involving the targeting of the cells, as in the human body there may be other cells that take these up such as neutrophils that may reduce the efficiency[168][169]. So while this method has potential it is important to do further testing with humans for its efficiency to be determined.

Another method for macrophage reprogramming is Hydroxychloroquine (HCQ). HCQ works by causing the transition from M2 to M1 macrophages. A proposed mechanism of action

is that since HCQ is weakly basic and soluble in membranes allowing it to pass through the macrophage plasma membrane[170]. HCQ then accumulates in lysosomes and disrupts signalling pathways causing lower release of inflammatory causing cytokines such as TNF-α[171]. By reducing pro fibrotic signalling, it discourages the switch of M1 macrophages to M2 and may encourage a higher M1 like macrophage population.

An experiment was conducted to investigate the renal protective effect on creatine clearance and Blood Urea Nitrogen levels after IRI on mice with HCQ involving in-vivo mice.The experiment was done for 7 days and was given through dissolving hCQ in CMC with different dosages of HCQ[172]. In terms of results, HCQ inhibited M2 expression and reduced macrophage activation and there was an increase in macrophage apoptosis meaning lesser fibrosis and more resolution[172][173].In addition, HCQ has already been approved for use by the US FDA as a medication for malaria, and a few other diseases (not fibrosis though) [174] while nanoparticles are still in the experimental stage. Also, nanoparticles may be inefficient as a solution because they could be taken up by non-target species in the body (e.g. macrophage in non fibrosis area) while HCQ doesn't depend on reaching targeted cells so its effects are more predictable. Lastly, since fibrosis is a chronic disease it requires a long term treatment plan and HCQ is suitable for this kind of task and is orally administered while nanoparticles only work for a short time. Thus HCQ, while still needing more clinical trials to be done involving humans shows more promise than the other solution involving nanoparticles,

Discussion

Fibrosis is an important medical condition that is a result of a wound healing mechanism failing to function properly as a result of repeated injury. The consequences include excessive ECM deposition, increased survival of myofibroblasts and inflammation. Prolonged TGF-β signaling, macrophage polarisation to M2 like state and other pathways cause resistance to apoptosis and increase in deposition of components such collagen. In the long run, fibrosis leads to loss of organ function and serious medical conditions.

Normal wound healing mechanisms have an inflammatory, proliferative, resolution and remodelling phase that is regulated by cytokines, fibrin, immune cells and more. The proliferative phase includes TGF-β and the SMAD pathway while the resolution phase revolves apoptosis and remodelling of the ECM.

However during fibrosis pathogenesis, the mechanisms above fail and the first two stages are prolonged. The myrofibroblasts resist apoptosis and the macrophages become polarised due to a M2 state due to specific cytokines being expressed. Instead of the ECM being remodelled, during fibrosis the DAMPS and the integrins are activated leading to slower resolution.

Nidenaib and Pirfenidone are both approved antifibrosis solutions that slow down its progression. Nintedanib is an inhibitor that blocks PDGFR and VEGFR signalling, reduces fibroblasts survival, and angiogenesis by inhibiting pathways that lead to an improvement in survival and lower FVC decline. pirfenidone done action on fibrosis is by downregulating TGF-β1 signalling through inhibiting NF-κB pathway, reducing collagen deposition and

inflammation Clinical Trials show both drugs have potential for success with non severe side effects, but in my opinion, nintedanib should be preferred as it is more successful than pirfenidone in lowering FVC capacity and covers more anti fibrosis pathways.

Another solution is to target TGF-β. Fresolimumab, a human monoclonal antibody, is a potential solution to reducing fibrosis. Fresolimumab neutralises all TGF-β isoforms, blocks receptor II signaling and affects other pathways. A FSGS trial has also supported the anti-fibrosis potential for fresolimumab. However, going forward it is important to continue research in this area to see if this a long term viable solution as the sample size for that experiment was too small and more trialling is required for a conclusive result.

A third solution is Macrophage reprogramming. Macrophage reprogramming works by shifting M2 macrophages to M1 like phenotype macrophages who play a more anti-fibrosis role compared to M2 macrophage states. Nanoparticles, a relatively new solution that lacks trialling in humans, can be used as a means of delivering IRF5/IIKB cytokines that promote transition to M1 states which, as seen in an experiment involving mice, showed an increased M1population. Likewise Hydroxychloroquine also promotes the macrophage transition by disrupting signalling from proinflammatory cytokines . HCQ unlike nanoparticles is a FDA approved drug and is more efficient than nanoparticles as there is higher efficiency due to lower chance of wastage and is more suitable for a long term treatment, and since fibrosis is a chronic disease HCQ shows more promise.

Overall, with the current solutions and therapeutic strategies available with research underway, it is very foreseeable that in the near future a viable solution for fibrosis will be found.

Works Cited

1.FAU. Connective tissue on the wrong road – when organs start to scar. Faculty of Medicine. Published February 6, 2019. Accessed March 7, 2026. https://www.med.fau.eu/2019/02/06/connective-tissue-on-the-wrong-road-when-organs-start-to-scar/

2.Eming SA, Martin P, Tomic-Canic M. Wound repair and regeneration: Mechanisms, signaling, and translation. Science Translational Medicine. 2014;6(265):265sr6-265sr6. doi:https://doi.org/10.1126/scitranslmed.3009337

3.Adler M, Mayo A, Zhou X, et al. Principles of Cell Circuits for Tissue Repair and Fibrosis. iScience. 2020;23(2):100841. doi:https://doi.org/10.1016/j.isci.2020.100841

4.Clinic C. Inflammation: What is it, causes, symptoms & treatment. Cleveland Clinic. Published March 22, 2024. https://my.clevelandclinic.org/health/symptoms/21660-inflammation

5.Arango Duque G, Descoteaux A. Macrophage Cytokines: Involvement in Immunity and Infectious Diseases. Frontiers in Immunology. 2014;5(491). doi:https://doi.org/10.3389/fimmu.2014.00491

6.Eom YW, Hong JE, Jung PY, et al. TGF-β expressed by M2 macrophages promotes wound healing by inhibiting TSG-6 expression by mesenchymal stem cells. Yuan Z, ed. PLOS ONE. 2025;20(4):e0316692. doi:https://doi.org/10.1371/journal.pone.0316692

7.Klingberg F, Hinz B, White ES. The myofibroblast matrix: implications for tissue repair and fibrosis. The Journal of Pathology. 2012;229(2):298-309. doi:https://doi.org/10.1002/path.4104

8.Cabral-Pacheco GA, Garza-Veloz I, Castruita-De la Rosa C, et al. The Roles of Matrix Metalloproteinases and Their Inhibitors in Human Diseases. International Journal of Molecular Sciences. 2020;21(24):9739. doi:https://doi.org/10.3390/ijms21249739

9.Chávez‐Galán L, Becerril C, Ruíz A, et al. Fibroblasts From Idiopathic Pulmonary Fibrosis Induce Apoptosis and Reduce the Migration Capacity of T Lymphocytes. Frontiers in Immunology. 2022;13. doi:https://doi.org/10.3389/fimmu.2022.820347

10.Wight TN, Potter-Perigo S. The extracellular matrix: an active or passive player in fibrosis? American Journal of Physiology-Gastrointestinal and Liver Physiology. 2011;301(6):G950-G955. doi:https://doi.org/10.1152/ajpgi.00132.2011

11.El Ayadi A, Jay JW, Prasai A. Current Approaches Targeting the Wound Healing Phases to Attenuate Fibrosis and Scarring. International Journal of Molecular Sciences. 2020;21(3):1105. doi:https://doi.org/10.3390/ijms21031105

12.Li X, Liu Y, Tang Y, Xia Z. Transformation of macrophages into myofibroblasts in fibrosis-related diseases: emerging biological concepts and potential mechanisms. Frontiers in Immunology. 2024;15. doi:https://doi.org/10.3389/fimmu.2024.1474688

13.Kanniyappan H, Chathurika Rathnayake RA, Osamor J, Islam M, Wang RR. The role of collagen and collagen I/III ratio in pathological conditions: insights into molecular mechanisms and therapeutic approaches. Frontiers in Bioengineering and Biotechnology. 2025;13. doi:https://doi.org/10.3389/fbioe.2025.1679625

14.Liu J, Khalil RA. Matrix Metalloproteinase Inhibitors as Investigational and Therapeutic Tools in Unrestrained Tissue Remodeling and Pathological Disorders. Progress in molecular biology and translational science. 2017;148:355-420. doi:https://doi.org/10.1016/bs.pmbts.2017.04.003
15.Plus M. Idiopathic pulmonary fibrosis: MedlinePlus Genetics. medlineplus.gov. https://medlineplus.gov/genetics/condition/idiopathic-pulmonary-fibrosis/
16.Ginès P, Graupera I, Lammert F, et al. Screening for liver fibrosis in the general population: a call for action. The Lancet Gastroenterology & Hepatology. 2016;1(3):256-260. doi:https://doi.org/10.1016/s2468-1253(16)30081-4
17.Wynn T. Cellular and molecular mechanisms of fibrosis. The Journal of Pathology. 2008;214(2):199-210. doi:https://doi.org/10.1002/path.2277
18.Hannoodee S, Nasuruddin DN. Acute Inflammatory Response. PubMed. Published 2024. https://www.ncbi.nlm.nih.gov/books/NBK556083/
19.Ansar W, Ghosh S. Inflammation and Inflammatory Diseases, Markers, and Mediators: Role of CRP in Some Inflammatory Diseases. Biology of C Reactive Protein in Health and Disease. Published online 2016:67-107. doi:https://doi.org/10.1007/978-81-322-2680-2_4
20.Weisel JW, Litvinov RI. Fibrin Formation, Structure and Properties. SpringerLink. 2017;82:405-456. doi:https://doi.org/10.1007/978-3-319-49674-0_13
21.Zhang B, Pang Z, Hu Y. Targeting hemostasis-related moieties for tumor treatment. Thrombosis Research. 2020;187:186-196. doi:https://doi.org/10.1016/j.thromres.2020.01.019
22.Krishnaswamy S. The Transition of Prothrombin to Thrombin. Journal of thrombosis and haemostasis : JTH. 2013;11(0 1):265-276. doi:https://doi.org/10.1111/jth.12217
23.Tarantino C. Coagulation Cascade: What Is It, Steps, and More | Osmosis. www.osmosis.org. Published 2021. https://www.osmosis.org/answers/coagulation-cascade
24.Keric N, Campbell A. Meshing around: high-risk hernias and infected mesh. Trauma surgery & acute care open. 2024;9(Suppl 2):e001379-e001379. doi:https://doi.org/10.1136/tsaco-2024-001379
25.Europe PMC. Europe PMC. Europepmc.org. Published 2016. Accessed March 7, 2026. https://europepmc.org/article/pmc/3857353
26.Blair P, Flaumenhaft R. Platelet α-granules: Basic biology and clinical correlates. Blood Reviews. 2009;23(4):177-189. doi:https://doi.org/10.1016/j.blre.2009.04.001
27.Segal AW. HOW NEUTROPHILS KILL MICROBES. Annual Review of Immunology. 2005;23(1):197-223. doi:https://doi.org/10.1146/annurev.immunol.23.021704.115653
28.T.K. Chow V. Neutrophils and NETs in Bacterial Infection and Pneumonia. https://medicine.nus.edu.sg/mbio/assets/documents/news/news/medicine_vincentchow.pdf
29.Yang J, Zhang L, Yu C, Yang XF, Wang H. Monocyte and macrophage differentiation: circulation inflammatory monocyte as biomarker for inflammatory diseases. Biomarker Research. 2014;2(1):1. doi:https://doi.org/10.1186/2050-7771-2-1

30.Sziksz E, Pap D, Lippai R, et al. Fibrosis Related Inflammatory Mediators: Role of the IL-10 Cytokine Family. Mediators of Inflammation. 2015;2015:1-15. doi:https://doi.org/10.1155/2015/764641
31. Tracy LE, Minasian RA, Caterson EJ. Extracellular Matrix and Dermal Fibroblast Function in the Healing Wound. Advances in Wound Care. 2016;5(3):119-136. doi:https://doi.org/10.1089/wound.2014.0561
32.Chia ZJ, Cao Y, Little PJ, Kamato D. Transforming growth factor-β receptors: versatile mechanisms of ligand activation. Acta Pharmacologica Sinica. 2024;45. doi:https://doi.org/10.1038/s41401-024-01235-6
33.Hata A, Chen YG. TGF-β Signaling from Receptors to Smads. Cold Spring Harbor Perspectives in Biology. 2016;8(9). doi:https://doi.org/10.1101/cshperspect.a022061
34.Klingberg F, Hinz B, White ES. The myofibroblast matrix: implications for tissue repair and fibrosis. The Journal of Pathology. 2012;229(2):298-309. doi:https://doi.org/10.1002/path.4104
35.Chen F, Lyu L, Xing C, et al. The pivotal role of TGF-β/Smad pathway in fibrosis pathogenesis and treatment. Frontiers in Oncology. 2025;15. doi:https://doi.org/10.3389/fonc.2025.1649179
36.Singh D, Rai V, K Agrawal D. Regulation of Collagen I and Collagen III in Tissue Injury and Regeneration. Cardiology and Cardiovascular Medicine. 2023;07(01). doi:https://doi.org/10.26502/fccm.92920302
37.Adair TH, Montani JP. Overview of Angiogenesis. Nih.gov. Published 2010. https://www.ncbi.nlm.nih.gov/books/NBK53238/
38.Johnson KE, Wilgus TA. Vascular Endothelial Growth Factor and Angiogenesis in the Regulation of Cutaneous Wound Repair. Advances in Wound Care. 2014;3(10):647-661. doi:https://doi.org/10.1089/wound.2013.0517
39.Abhinand CS, Raju R, Soumya SJ, Arya PS, Sudhakaran PR. VEGF-A/VEGFR2 signaling network in endothelial cells relevant to angiogenesis. Journal of Cell Communication and Signaling. 2016;10(4):347-354. doi:https://doi.org/10.1007/s12079-016-0352-8
40.Iruela-Arispe ML, Davis GE. Cellular and Molecular Mechanisms of Vascular Lumen Formation. Developmental Cell. 2009;16(2):222-231. doi:https://doi.org/10.1016/j.devcel.2009.01.013
41.Shibuya M. Vascular Endothelial Growth Factor (VEGF) and Its Receptor (VEGFR) Signaling in Angiogenesis: A Crucial Target for Anti- and Pro-Angiogenic Therapies. Genes & Cancer. 2011;2(12):1097-1105. doi:https://doi.org/10.1177/1947601911423031
42.Diller RB, Tabor AJ. The Role of the Extracellular Matrix (ECM) in Wound Healing: A Review. National Library of Medicine. Published June 1, 2022. https://pmc.ncbi.nlm.nih.gov/articles/PMC9326521/
43.Greenlee-Wacker MC. Clearance of apoptotic neutrophils and resolution of inflammation. Immunological Reviews. 2016;273(1):357-370. doi:https://doi.org/10.1111/imr.12453

44.Herchenhan A, Uhlenbrock F, Eliasson P, et al. Lysyl Oxidase Activity Is Required for Ordered Collagen Fibrillogenesis by Tendon Cells. Journal of Biological Chemistry. 2015;290(26):16440-16450. doi:https://doi.org/10.1074/jbc.m115.641670
45.Tu H, Ren H, Jiang J, Shao C, Shi Y, Li P. Dying to Defend: Neutrophil Death Pathways and their Implications in Immunity. Advanced science. 2023;11(8). doi:https://doi.org/10.1002/advs.202306457
46.Schuster R, Fereshteh Younesi, Ezzo M, Hinz B. The Role of Myofibroblasts in Physiological and Pathological Tissue Repair. Cold Spring Harbor Perspectives in Biology. 2022;15(1):a041231-a041231. doi:https://doi.org/10.1101/cshperspect.a041231
47.Putra A, Alif I, Hamra N, et al. MSC-released TGF-ß regulate a-SMA expression of myofibroblast during wound healing. Journal of Stem Cells and Regenerative Medicine. 2020;16(2):73-79. doi:https://doi.org/10.46582/jsrm.1602011
48.Li B, Wang JHC . Fibroblasts and myofibroblasts in wound healing: Force generation and measurement. Journal of Tissue Viability. 2011;20(4):108-120. doi:https://doi.org/10.1016/j.jtv.2009.11.004
49.Dyken S. INTERLEUKIN-4- AND INTERLEUKIN-13-MEDIATED ALTERNATIVELY ACTIVATED MACROPHAGES: ROLES IN HOMEOSTASIS AND DISEASE [Internet]. Nih.gov. 2024. Available from: https://pmc.ncbi.nlm.nih.gov/articles/PMC3606684/Zha J, Harada H, Yang E, Jockel J, Korsmeyer SJ. Serine Phosphorylation of Death Agonist BAD in Response to Survival Factor Results in Binding to 14-3-3 Not BCL-XL. Cell. 1996;87(4):619-628. doi:https://doi.org/10.1016/s0092-8674(00)81382-3
50.Datta SR, Dudek H, Tao X, et al. Akt phosphorylation of BAD couples survival signals to the cell-intrinsic death machinery. Cell. 1997;91(2):231-241. doi:https://doi.org/10.1016/s0092-8674(00)80405-5
51.Richard Seonghun Nho, Hong X, Kahn J, Kleidon J, Diebold D, Henke CA. Role of Integrin-linked Kinase in Regulating Phosphorylation of Akt and Fibroblast Survival in Type I Collagen Matrices through a β1 Integrin Viability Signaling Pathway. Journal of Biological Chemistry. 2005;280(28):26630-26639. doi:https://doi.org/10.1074/jbc.m411798200
52.Schuster R, Fereshteh Younesi, Ezzo M, Hinz B. The Role of Myofibroblasts in Physiological and Pathological Tissue Repair. Cold Spring Harbor Perspectives in Biology. 2022;15(1):a041231-a041231. doi:https://doi.org/10.1101/cshperspect.a041231
53.Wu CC, Bratton SB. Regulation of the Intrinsic Apoptosis Pathway by Reactive Oxygen Species. Antioxidants & Redox Signaling. 2013;19(6):546-558. doi:https://doi.org/10.1089/ars.2012.4905
54.Aubrey BJ, Kelly GL, Janic A, Herold MJ, Strasser A. How does p53 induce apoptosis and how does this relate to p53-mediated tumour suppression? Cell Death & Differentiation. 2017;25(1):104-113. doi:https://doi.org/10.1038/cdd.2017.169

55.Yang Y, Li CCH, Weissman AM. Regulating the p53 system through ubiquitination. Oncogene. 2004;23(11):2096-2106. doi:https://doi.org/10.1038/sj.onc.1207411
56.Chen L, Gilkes DM, Pan Y, Lane WS, Chen J. ATM and Chk2-dependent phosphorylation of MDMX contribute to p53 activation after DNA damage. The EMBO Journal. 2005;24(19):3411-3422. doi:https://doi.org/10.1038/sj.emboj.7600812
57.Godefroy N. Transcriptional repression by p53 promotes a Bcl-2-insensitive and mitochondria-independent pathway of apoptosis. Nucleic Acids Research. 2004;32(15):4480-4490. doi:https://doi.org/10.1093/nar/gkh773
58.Qian S, Wei Z, Yang W, Huang J, Yang Y, Wang J. The role of BCL-2 family proteins in regulating apoptosis and cancer therapy. Frontiers in Oncology. 2022;12. doi:https://doi.org/10.3389/fonc.2022.985363
59.Jiang X, Wang X. Cytochrome c Promotes Caspase-9 Activation by Inducing Nucleotide Binding to Apaf-1. Journal of Biological Chemistry. 2000;275(40):31199-31203. doi:https://doi.org/10.1074/jbc.c000405200
60.Avrutsky MI, Troy CM. Caspase-9: A Multimodal Therapeutic Target With Diverse Cellular Expression in Human Disease. Frontiers in Pharmacology. 2021;12. doi:https://doi.org/10.3389/fphar.2021.701301
61.Khan FS, Fujioka M, Datta P, Fernandes-Alnemri T, Jaynes JB, Alnemri ES. The interaction of DIAP1 with dOmi/HtrA2 regulates cell death in Drosophila. Cell Death & Differentiation. 2008;15(6):1073-1083. doi:https://doi.org/10.1038/cdd.2008.19
62.Yang QH . Omi/HtrA2 catalytic cleavage of inhibitor of apoptosis (IAP) irreversibly inactivates IAPs and facilitates caspase activity in apoptosis. Genes & Development. 2003;17(12):1487-1496. doi:https://doi.org/10.1101/gad.1097903
63.van Loo G, van Gurp M, Depuydt B, et al. The serine protease Omi/HtrA2 is released from mitochondria during apoptosis. Omi interacts with caspase-inhibitor XIAP and induces enhanced caspase activity. Cell Death & Differentiation. 2002;9(1):20-26. doi:https://doi.org/10.1038/sj.cdd.4400970
64.Yeretssian G, Labbé K, Saleh M. Molecular regulation of inflammation and cell death. Cytokine. 2008;43(3):380-390. doi:https://doi.org/10.1016/j.cyto.2008.07.015
65.Verhagen AM, Coulson EJ, Vaux DL. Inhibitor of apoptosis proteins and their relatives: IAPs and other BIRPs. Genome Biology. 2001;2(7):reviews3009.1. doi:https://doi.org/10.1186/gb-2001-2-7-reviews3009
66.Varfolomeev E, Goncharov T, Vucic D. Roles of c-IAP proteins in TNF receptor family activation of NF-κB signaling. Methods in molecular biology (Clifton, NJ). 2015;1280:269-282. doi:https://doi.org/10.1007/978-1-4939-2422-6_15
67.Bender T, Martinou JC . Where Killers Meet--Permeabilization of the Outer Mitochondrial Membrane during Apoptosis. Cold Spring Harbor Perspectives in Biology. 2013;5(1):a011106-a011106. doi:https://doi.org/10.1101/cshperspect.a011106

68. Vogler M, Braun Y, Smith VM, et al. The BCL2 family: from apoptosis mechanisms to new advances in targeted therapy. Signal transduction and targeted therapy. 2025;10(1):91. doi:https://doi.org/10.1038/s41392-025-02176-0
69. Wang YC, Wang LT, Hung TI, et al. Severe cellular stress drives apoptosis through a dual control mechanism independently of p53. Cell Death Discovery. 2022;8(1):1-5. doi:https://doi.org/10.1038/s41420-022-01078-2
70. Kim KK, Sheppard D, Chapman HA. TGF-β1 Signaling and Tissue Fibrosis. Cold Spring Harbor Perspectives in Biology. 2017;10(4):a022293. doi:https://doi.org/10.1101/cshperspect.a022293
71. Luo K. Signaling Cross Talk between TGF-β/Smad and Other Signaling Pathways. Cold Spring Harbor Perspectives in Biology. 2016;9(1):a022137. doi:https://doi.org/10.1101/cshperspect.a022137
72. Silva JPN, Pinto B, Silva PMA, Bousbaa H. BCL-2 and BCL-xL in Cancer: Regulation, Function, and Therapeutic Targeting. International Journal of Molecular Sciences. 2026;27(2):1123. doi:https://doi.org/10.3390/ijms27021123
73. Lawrence T. The Nuclear Factor NF- B Pathway in Inflammation. Cold Spring Harbor Perspectives in Biology. 2009;1(6):a001651-a001651. doi:https://doi.org/10.1101/cshperspect.a001651
74. Oeckinghaus A, Ghosh S. The NF- B Family of Transcription Factors and Its Regulation. Cold Spring Harbor Perspectives in Biology. 2009;1(4):a000034-a000034. doi:https://doi.org/10.1101/cshperspect.a000034
75. Luo JL . IKK/NF- B signaling: balancing life and death - a new approach to cancer therapy. Journal of Clinical Investigation. 2005;115(10):2625-2632. doi:https://doi.org/10.1172/jci26322
76. Tang LY, Heller M, Meng Z, et al. Transforming Growth Factor-β (TGF-β) Directly Activates the JAK1-STAT3 Axis to Induce Hepatic Fibrosis in Coordination with the SMAD Pathway. Journal of Biological Chemistry. 2017;292(10):4302-4312. doi:https://doi.org/10.1074/jbc.m116.773085
77. Seif F, Khoshmirsafa M, Aazami H, Mohsenzadegan M, Sedighi G, Bahar M. The role of JAK-STAT signaling pathway and its regulators in the fate of T helper cells. Cell Communication and Signaling. 2017;15(1). doi:https://doi.org/10.1186/s12964-017-0177-y
78. Lam DK, Schmidt BL. Molecular Biology of Head and Neck Cancer. Current Therapy In Oral and Maxillofacial Surgery. Published online 2012:92-101. doi:https://doi.org/10.1016/b978-1-4160-2527-6.00010-4
79. Egan B, Sharples AP. RNA polymerase II - Knowledge and References | Taylor & Francis. Taylor & Francis. Published 2021. Accessed March 7, 2026. https://taylorandfrancis.com/knowledge/Medicine_and_healthcare/Physiology/RNA_polymerase_II/

80.Haymour L, Jean M, Smulski C, Legembre P. CD95 (Fas) and CD95L (FasL)-mediated non-canonical signaling pathways. Biochimica Et Biophysica Acta Reviews on Cancer. 2023;1878(6):189004. doi:https://doi.org/10.1016/j.bbcan.2023.189004
81.Kisenge RR, Toyoda H, Kang J, et al. Expression of short‐form caspase 8 correlates with decreased sensitivity to Fas‐mediated apoptosis in neuroblastoma cells. Cancer Science. 2003;94(7):598-605. doi:https://doi.org/10.1111/j.1349-7006.2003.tb01489.x
82.Roy A, Hong JH, Lee JH, Lee YT, Lee BJ, Kim KS. In vitro activation of procaspase-8 by forming the cytoplasmic component of the death-inducing signaling complex (cDISC). Molecules and cells. 2008;26(2):165-170. https://pubmed.ncbi.nlm.nih.gov/18596409/
83.Orning P, Lien E. Multiple roles of caspase‐8 in cell death, inflammation, and innate immunity. Journal of Leukocyte Biology. Published online June 12, 2020. doi:https://doi.org/10.1002/jlb.3mr0420-305r
84.Huang K, Zhang J, O'Neill KL, et al. Cleavage by Caspase 8 and Mitochondrial Membrane Association Activate the BH3-only Protein Bid during TRAIL-induced Apoptosis. Journal of Biological Chemistry. 2016;291(22):11843-11851. doi:https://doi.org/10.1074/jbc.m115.711051
85.Yang CY, Tseng YC, Tu YF, et al. Reverse hierarchical DED assembly in the cFLIP-procaspase-8 and cFLIP-procaspase-8-FADD complexes. Nature Communications. 2024;15(1). doi:https://doi.org/10.1038/s41467-024-53306-1
86.Mao H, Zhao X, Sun S. NF-κB in inflammation and cancer. Cellular & Molecular Immunology. Published online June 25, 2025. doi:https://doi.org/10.1038/s41423-025-01310-w
87.Di X, Gao X, Luo D, et al. Cellular mechanotransduction in health and diseases: from molecular mechanism to therapeutic targets. Signal Transduction and Targeted Therapy. 2023;8(1). doi:https://doi.org/10.1038/s41392-023-01501-9
88.Dimberg LY, Dimberg AI, Ivarsson K, et al. Ectopic and IFN-induced expression of Fas overcomes resistance to Fas-mediated apoptosis in multiple myeloma cells. Blood. 2005;106(4):1346-1354. doi:https://doi.org/10.1182/blood-2004-04-1322
89.Sobrido-Cameán D, Antón Barreiro-Iglesias. Role of Caspase-8 and Fas in Cell Death After Spinal Cord Injury. Frontiers in Molecular Neuroscience. 2018;11. doi:https://doi.org/10.3389/fnmol.2018.00101
90.Chang DW, Xing Z, Capacio VL, Peter ME, Yang X. Interdimer processing mechanism of procaspase-8 activation. The EMBO Journal. 2003;22(16):4132-4142. doi:https://doi.org/10.1093/emboj/cdg414
91.Oberst A, Dillon CP, Weinlich R, et al. Catalytic activity of the caspase-8–FLIPL complex inhibits RIPK3-dependent necrosis. Nature. 2011;471(7338):363-367. doi:https://doi.org/10.1038/nature09852
92.Hughes Michelle A, Powley Ian R, Jukes-Jones R, et al. Co-operative and Hierarchical Binding of c-FLIP and Caspase-8: A Unified Model Defines How c-FLIP Isoforms

Differentially Control Cell Fate. Molecular Cell. 2016;61(6):834-849. doi:https://doi.org/10.1016/j.molcel.2016.02.023
93.Tsuchiya Y, Nakabayashi O, Nakano H. FLIP the Switch: Regulation of Apoptosis and Necroptosis by cFLIP. International Journal of Molecular Sciences. 2015;16(12):30321-30341. doi:https://doi.org/10.3390/ijms161226232
94.PETER Marcus E. The flip side of FLIP: Figure 1. Biochemical Journal. 2004;382(2):e1. doi:https://doi.org/10.1042/bj20041143
95.Richeldi L, du Bois RM, Raghu G, et al. Efficacy and Safety of Nintedanib in Idiopathic Pulmonary Fibrosis. New England Journal of Medicine. 2014;370(22):2071-2082. doi:https://doi.org/10.1056/nejmoa1402584
96.Gole S, Bankole A. Nintedanib. PubMed. Published 2022. https://www.ncbi.nlm.nih.gov/books/NBK585049/
97.Wollin L, Wex E, Pautsch A, et al. Mode of action of nintedanib in the treatment of idiopathic pulmonary fibrosis. European Respiratory Journal. 2015;45(5):1434-1445. doi:https://doi.org/10.1183/09031936.00174914
98.Nessrine Bellamri, Morzadec C, Joannes A, et al. Alteration of human macrophage phenotypes by the anti-fibrotic drug nintedanib. International Immunopharmacology. 2019;72:112-123. doi:https://doi.org/10.1016/j.intimp.2019.03.061
99.Jamadar A, Suma SM, Mathew S, et al. The tyrosine-kinase inhibitor Nintedanib ameliorates autosomal-dominant polycystic kidney disease. Cell Death & Disease. 2021;12(10). doi:https://doi.org/10.1038/s41419-021-04248-9
100.Olivieri C, Geoffrey Ye Li, Wang Y, et al. ATP-competitive inhibitors modulate the substrate binding cooperativity of a kinase by altering its conformational entropy. Science Advances. 2022;8(30). doi:https://doi.org/10.1126/sciadv.abo0696
101.Hubbard SR, Miller WT. Receptor tyrosine kinases: mechanisms of activation and signaling. Current Opinion in Cell Biology. 2007;19(2):117-123. doi:https://doi.org/10.1016/j.ceb.2007.02.010
102.Ciprian Tomuleasa, Adrian-Bogdan Tigu, Munteanu R, et al. Therapeutic advances of targeting receptor tyrosine kinases in cancer. Signal Transduction and Targeted Therapy. 2024;9(1). doi:https://doi.org/10.1038/s41392-024-01899-w
103.Emad Shamsan, Maged Almezgagi, Gamah M, et al. The role of PI3k/AKT signaling pathway in attenuating liver fibrosis: a comprehensive review. Frontiers in Medicine. 2024;11. doi:https://doi.org/10.3389/fmed.2024.1389329
104.Wind S, Schmid U, Freiwald M, et al. Clinical Pharmacokinetics and Pharmacodynamics of Nintedanib. Clinical Pharmacokinetics. 2019;58(9):1131-1147. doi:https://doi.org/10.1007/s40262-019-00766-0
105.Kolch W, Berta D, Rosta E. Dynamic regulation of RAS and RAS signaling. Biochemical Journal. 2023;480(1):1-23. doi:https://doi.org/10.1042/bcj20220234

106.Man RK, Amaresh Gogikar, Nanda A, et al. A Comparison of the Effectiveness of Nintedanib and Pirfenidone in Treating Idiopathic Pulmonary Fibrosis: A Systematic Review. Curēus. 2024;16(2). doi:https://doi.org/10.7759/cureus.54268
107.Bonella F, Cottin V, Valenzuela C, et al. Meta-Analysis of Effect of Nintedanib on Reducing FVC Decline Across Interstitial Lung Diseases. Advances in Therapy. 2022;39(7):3392-3402. doi:https://doi.org/10.1007/s12325-022-02145-x
108.Štefániková M, Doubková M, Ovesná P, et al. The effect of nintedanib on lung functions and survival in idiopathic pulmonary fibrosis: real-life analysis of the Czech EMPIRE registry. BMC Pulmonary Medicine. 2023;23(1). doi:https://doi.org/10.1186/s12890-023-02450-3
109.Lancaster L, Crestani B, Hernandez P, et al. Safety and survival data in patients with idiopathic pulmonary fibrosis treated with nintedanib: pooled data from six clinical trials. BMJ Open Respiratory Research. 2019;6(1):e000397. doi:https://doi.org/10.1136/bmjresp-2018-000397
110.Clinic M. Pirfenidone (oral route). Mayo Clinic. Published 2025. https://www.mayoclinic.org/drugs-supplements/pirfenidone-oral-route/description/drg-20122610
111.Ballester B, Milara J, Cortijo J. Pirfenidone anti-fibrotic effects are partially mediated by the inhibition of MUC1 bioactivation. Oncotarget. 2020;11(15):1306-1320. doi:https://doi.org/10.18632/oncotarget.27526
112.Mora DAL la, Sanchez-Roque C, Montoya-Buelna M, et al. Role and New Insights of Pirfenidone in Fibrotic Diseases. International Journal of Medical Sciences. 2015;12(11):840-847. doi:https://doi.org/10.7150/ijms.11579
113.Babariya H, Gaidhane SA, Acharya S, Kumar S. Pirfenidone as a Cornerstone in the Management of Fibrotic Interstitial Lung Diseases and Its Emerging Applications: A Comprehensive Review. Cureus. Published online September 30, 2024. doi:https://doi.org/10.7759/cureus.70497
114.Zhou P, Wan X, Zou Y, Chen Z, Zhong A. Transforming growth factor beta (TGF-β) is activated by the CtBP2-p300-AP1 transcriptional complex in chronic renal failure. International Journal of Biological Sciences. 2020;16(2):204-215. doi:https://doi.org/10.7150/ijbs.38841
115.Zhang Z, Song J, Kwon SH, et al. Pirfenidone Inhibits Alveolar Bone Loss in Ligature-Induced Periodontitis by Suppressing the NF-κB Signaling Pathway in Mice. International Journal of Molecular Sciences. 2023;24(10):8682-8682. doi:https://doi.org/10.3390/ijms24108682
116.of Medicine NL. JUN Jun proto-oncogene, AP-1 transcription factor subunit [Homo sapiens (human)] - Gene - NCBI. www.ncbi.nlm.nih.gov. https://www.ncbi.nlm.nih.gov/gene/3725

117.Guo B, Inoki K, Isono M, et al. MAPK/AP-1-dependent regulation of PAI-1 gene expression by TGF-β in rat mesangial cells. Kidney International. 2005;68(3):972-984. doi:https://doi.org/10.1111/j.1523-1755.2005.00491.x
118.Guo X, Yang Y, Liu L, et al. Pirfenidone Induces G1 Arrest in Human Tenon's Fibroblasts In Vitro Involving AKT and MAPK Signaling Pathways. Journal of Ocular Pharmacology and Therapeutics. 2017;33(5):366-374. doi:https://doi.org/10.1089/jop.2016.0107
119.Estornut C, Milara J, Bayarri MA, Belhadj N, Cortijo J. Targeting Oxidative Stress as a Therapeutic Approach for Idiopathic Pulmonary Fibrosis. Frontiers in Pharmacology. 2022;12. doi:https://doi.org/10.3389/fphar.2021.794997
120.de Almeida AJPO, de Oliveira JCPL, da Silva Pontes LV, et al. ROS: Basic Concepts, Sources, Cellular Signaling, and its Implications in Aging Pathways. Giustarini D, ed. Oxidative Medicine and Cellular Longevity. 2022;2022(1):1-23. doi:https://doi.org/10.1155/2022/1225578
121.Hadda V, Guleria R. Antifibrotic drugs for idiopathic pulmonary fibrosis: What we should know? Indian Journal of Medical Research. 2020;152(3):177. doi:https://doi.org/10.4103/ijmr.ijmr_90_20
122.Noble PW, Albera C, Bradford WZ, et al. Pirfenidone for idiopathic pulmonary fibrosis: analysis of pooled data from three multinational phase 3 trials. European Respiratory Journal. 2015;47(1):243-253. doi:https://doi.org/10.1183/13993003.00026-2015
123.Plus M. Pirfenidone: MedlinePlus Drug Information. medlineplus.gov. https://medlineplus.gov/druginfo/meds/a615008.html
124.FDA. HIGHLIGHTS OF PRESCRIBING INFORMATION. https://www.accessdata.fda.gov/drugsatfda_docs/label/2019/022535s012,208780s002lbl.pdf
125.Richeldi L, du Bois RM, Raghu G, et al. Efficacy and Safety of Nintedanib in Idiopathic Pulmonary Fibrosis. New England Journal of Medicine. 2014;370(22):2071-2082. doi:https://doi.org/10.1056/nejmoa1402584
126.Noble P, Albera C, Chou W, et al. Annual rate of FVC decline in patients with IPF treated with pirfenidone: Pooled analysis from 3 pivotal studies. 15 Diffuse Parenchymal Lung Disease. Published online September 2016:OA1810. doi:https://doi.org/10.1183/13993003.congress-2016.oa1810
127.Cravedi P, Angeletti A, Remuzzi G. New biologics in the treatment of rare glomerular diseases of childhood. Current Opinion in Pharmacology. 2017;33:27-33. doi:https://doi.org/10.1016/j.coph.2017.03.010
128.Robertson IB, Rifkin DB. Regulation of the Bioavailability of TGF-β and TGF-β-Related Proteins. Cold Spring Harbor Perspectives in Biology. 2016;8(6):a021907. doi:https://doi.org/10.1101/cshperspect.a021907
129.Hinz B. The extracellular matrix and transforming growth factor-β1: Tale of a strained relationship. Matrix Biology. 2015;47:54-65. doi:https://doi.org/10.1016/j.matbio.2015.05.006

130.Chia ZJ, Cao Y, Little PJ, Kamato D. Transforming growth factor-β receptors: versatile mechanisms of ligand activation. Acta Pharmacologica Sinica. Published online February 13, 2024. doi:https://doi.org/10.1038/s41401-024-01235-6
131.Massagué J, Sheppard D. TGF-β signaling in health and disease. Cell. 2023;186(19):4007-4037. doi:https://doi.org/10.1016/j.cell.2023.07.036
132.Buwaneka P, Ralko A, Gorai S, Pham H, Cho W. Phosphoinositide-binding activity of Smad2 is essential for its function in TGF-β signaling. The Journal of biological chemistry. 2021;297(5):101303. doi:https://doi.org/10.1016/j.jbc.2021.101303
133.Morris JC, Tan AR, Olencki TE, et al. Phase I Study of GC1008 (Fresolimumab): A Human Anti-Transforming Growth Factor-Beta (TGFβ) Monoclonal Antibody in Patients with Advanced Malignant Melanoma or Renal Cell Carcinoma. PLoS ONE. 2014;9(3). doi:https://doi.org/10.1371/journal.pone.0090353
134.Trachtman H, Fervenza FC, Gipson DS, et al. A phase 1, single-dose study of fresolimumab, an anti-TGF-β antibody, in treatment-resistant primary focal segmental glomerulosclerosis. Kidney International. 2011;79(11):1236-1243. doi:https://doi.org/10.1038/ki.2011.33
135.Barutta F, Bellini S, Gruden G. Mechanisms of podocyte injury and implications for diabetic nephropathy. Clinical Science. 2022;136(7):493-520. doi:https://doi.org/10.1042/cs20210625
136.Ma PF, Gao CC, Yi J, et al. Cytotherapy with M1-polarized macrophages ameliorates liver fibrosis by modulating the immune microenvironment in mice. Journal of Hepatology. 2017;67(4):770-779. doi:https://doi.org/10.1016/j.jhep.2017.05.022
137.Periklis Dousdampanis, Ioanna Aggeletopoulou, Athanasia Mouzaki. The role of M1/M2 macrophage polarization in the pathogenesis of obesity-related kidney disease and related pathologies. Frontiers in Immunology. 2025;15. doi:https://doi.org/10.3389/fimmu.2024.1534823
138.Liang S, Wang F, Zhai D, Meng X, Li J, Xiongwen Lv. Matrix metalloproteinases induce extracellular matrix degradation through various pathways to alleviate hepatic fibrosis. PubMed. 2023;161:114472-114472. doi:https://doi.org/10.1016/j.biopha.2023.114472
139.Yan L, Wang J, Cai X, et al. Macrophage plasticity: signaling pathways, tissue repair, and regeneration. MedComm. 2024;5(8). doi:https://doi.org/10.1002/mco2.658
140.Deepak Pokhreal, Crestani B, Doumet Georges Helou. Macrophage Implication in IPF: Updates on Immune, Epigenetic, and Metabolic Pathways. Cells. 2023;12(17):2193-2193. doi:https://doi.org/10.3390/cells12172193
141.Tung NTC, Nogami M, Iwasaki M, et al. M2-like macrophages derived from THP-1 cells promote myofibroblast differentiation of synovial fibroblasts in association with the TGF-β1/SMAD2/3 signaling pathway. Scientific reports. 2025;15(1):25505. doi:https://doi.org/10.1038/s41598-025-10858-6
142.Wynn T, Barron L. Macrophages: Master Regulators of Inflammation and Fibrosis. Seminars in Liver Disease. 2010;30(03):245-257. doi:https://doi.org/10.1055/s-0030-1255354

143.Qi L, Yu H, Zhang Y, et al. IL-10 secreted by M2 macrophage promoted tumorigenesis through interaction with JAK2 in glioma. Oncotarget. 2016;7(44):71673-71685. doi:https://doi.org/10.18632/oncotarget.12317
144.Jumana Alhamdi, Peng T, Al-Naggar IM, Hawley KL, Spiller KL, Kuhn LT. Controlled M1-to-M2 transition of aged macrophages by calcium phosphate coatings. Controlled M1-to-M2 transition of aged macrophages by calcium phosphate coatings. 2019;196:90-99. doi:https://doi.org/10.1016/j.biomaterials.2018.07.012
145.Zhang F, Parayath NN, Ene CI, et al. Genetic programming of macrophages to perform anti-tumor functions using targeted mRNA nanocarriers. Nature Communications. 2019;10(1). doi:https://doi.org/10.1038/s41467-019-11911-5
146.Chong L, Jiang YW, Wang D, Chang P, Xu K, Li J. Targeting and repolarizing M2-like tumor-associated macrophage-mediated MR imaging and tumor immunotherapy by biomimetic nanoparticles. Journal of Nanobiotechnology. 2023;21(1). doi:https://doi.org/10.1186/s12951-023-02122-8
147.Zhang F, Parayath NN, Ene CI, et al. Genetic programming of macrophages to perform anti-tumor functions using targeted mRNA nanocarriers. Nature Communications. 2019;10(1). doi:https://doi.org/10.1038/s41467-019-11911-5
148.Almuttaqi H, Udalova IA. Advances and challenges in targeting IRF5, a key regulator of inflammation. The Febs Journal. 2019;286(9):1624-1637. doi:https://doi.org/10.1111/febs.14654
149.Lihui X, Zhao J, Skonieczna M, Zhou P, Huang R. Nanoparticles‐induced potential toxicity on human health: Applications, toxicity mechanisms, and evaluation models. MedComm. 2023;4(4). doi:https://doi.org/10.1002/mco2.327
150.Liang Ee Low, Chee Kei Kong, Yap WH, et al. Hydroxychloroquine: Key therapeutic advances and emerging nanotechnological landscape for cancer mitigation. Chemico-biological interactions. 2023;386:110750-110750. doi:https://doi.org/10.1016/j.cbi.2023.110750
151.Moosa Al-Hamadani, Mahmoud Darweesh, Mohammadi S, Al-Harrasi A. Chloroquine and hydroxychloroquine: Immunomodulatory effects in autoimmune diseases. World Journal of Biological Chemistry. 2025;16(2). doi:https://doi.org/10.4331/wjbc.v16.i2.107042
152.Zheng H, Zhang Y, He J, et al. Hydroxychloroquine Inhibits Macrophage Activation and Attenuates Renal Fibrosis After Ischemia-Reperfusion Injury. Frontiers in Immunology. 2021;12:645100-645100. doi:https://doi.org/10.3389/fimmu.2021.645100
153.Zheng H, Zhang Y, He J, et al. Hydroxychloroquine Inhibits Macrophage Activation and Attenuates Renal Fibrosis After Ischemia-Reperfusion Injury. Frontiers in Immunology. 2021;12:645100-645100. doi:https://doi.org/10.3389/fimmu.2021.645100
154.Commissioner O of the. Coronavirus (COVID-19) Update: FDA Revokes Emergency Use Authorization for Chloroquine and Hydroxychloroquine. FDA. Published June 16, 2020. https://www.fda.gov/news-events/press-announcements/coronavirus-covid-19-update-fda-revokes-emergency-use-authorization-chloroquine-and

An Analysis of Amyloid, Tau, and Inflammation in Alzheimer's Disease By Shaunak Sen

Abstract

Alzheimer's disease (AD) is a progressive brain disorder that leads to memory loss, cognitive decline, and difficulty performing daily activities. Although current treatments mainly focus on managing symptoms, understanding the underlying causes of the disease can lead to more effective therapies and earlier detection. This report examines the roles of amyloid-beta plaques, tau protein tangles, and neuroinflammation in the development and progression of Alzheimer's disease. Amyloid-beta plaques build up between neurons and disrupt communication, while tau tangles form inside neurons and block essential transport systems, leading to cell death. In addition, chronic neuroinflammation caused by overactive immune cells further damages brain tissue. By studying how these three factors interact, researchers can develop targeted treatments, such as antibody therapies, small-molecule drugs, and stem cell approaches, that aim to slow or stop disease progression. This research also supports the development of biomarkers, including protein levels in blood or cerebrospinal fluid and brain imaging techniques, which can detect Alzheimer's disease at an early stage before major symptoms appear. This report also covers how lifestyle modifications, such as diet, regular exercise, and proper sleep, can help lower the risk of AD or slow its progression (Kivipelto et al. 398). Overall, a better understanding of the biological processes of AD and factors which can ameliorate AD progression offers hope for earlier diagnosis, improved treatments, and potentially preventing Alzheimer's disease in the future.

Introduction

Alzheimer's disease (AD) is a serious and progressive disorder of the brain that slowly damages memory, thinking skills, and the ability to perform everyday tasks. It is the most common cause of dementia, which is a general term for the decline of mental abilities that is severe enough to interfere with daily life (Lane, Hardy, & Schott. 61). In the United States, an estimated 7.2 million people aged 65 or older are currently living with Alzheimer's disease, which equates to approximately one in nine adults in this age group. Women are affected more often than men, making up almost two-thirds of diagnosed cases (Benilova. 349). Researchers posit that sex differences exist because women live longer on average and experience major hormonal changes (Gordon). The number of people living with Alzheimer's continues to grow as life expectancy increases. Experts predict that by 2060, the number of people with Alzheimer's in the U.S. could rise to nearly 14 million if no major treatment breakthroughs occur. Currently, there is no cure for AD, and most treatments only address the symptoms rather than the underlying cause of the disease. In other words, medications can help patients remember things better or manage confusion, but they cannot stop the disease from progressing in the brain (National Institute on Aging). Age is the biggest risk factor for developing Alzheimer's disease, which is why the disease is becoming more common as the population ages (National Institute on Aging).

Neuropathology

The brain changes that occur in Alzheimer's disease develop gradually over many years. At first, neurons struggle to communicate effectively, and over time, they begin to die. This progressive loss of cells leads to noticeable brain shrinkage, particularly in areas involved in memory, such as the hippocampus. Three major biological changes contribute to Alzheimer's disease: amyloid plaques, tau tangles, and chronic inflammation. Amyloid plaques are sticky clumps of beta-amyloid protein that collect between neurons and disrupt communication. Tau tangles are twisted strands of tau protein found inside neurons that block the transport of nutrients and cause cell death. Both abnormal protein formations interfere with normal brain function. Chronic inflammation occurs when overactive immune cells in the brain accidentally damage healthy brain tissue and is associated with neurodegeneration, which is the gradual death of neurons or nerve cells (Chen et al. 165). These changes interfere with brain function long before any noticeable symptoms appear (Braak & Braak. 239).

As shown in Figure 1, the earliest structural and functional changes in Alzheimer's disease occur in the entorhinal cortex and hippocampus, two regions that are essential for learning and the formation of new memories. The entorhinal cortex serves as the primary communication pathway between the hippocampus and the rest of the cerebral cortex, allowing information from different parts of the brain to be integrated into long-term memories. These regions are affected first because their neurons are highly interconnected, have high metabolic demands, and are particularly vulnerable to the accumulation of abnormal proteins such as amyloid-beta and hyperphosphorylated tau. As neurons in these areas begin to degenerate, patients typically experience short-term memory loss, one of the earliest clinical signs of Alzheimer's disease. As the disease progresses, degeneration spreads to the temporal, parietal, and frontal lobes, leading to impairments in language, reasoning, decision-making, and other cognitive functions (Braak & Braak. 239; Lane, Hardy, & Schott. 61; Hampel et al).

Figure 1. Early Anatomical Progression of Alzheimer's Disease

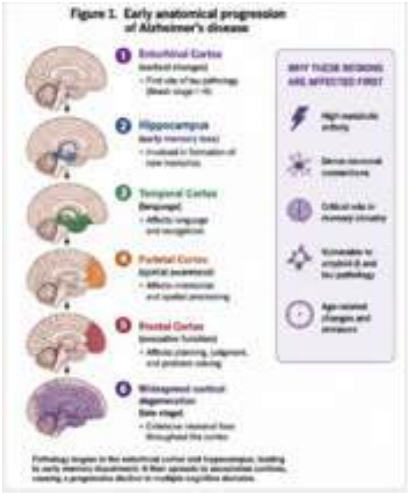

How each part of the brain is affected by Alzheimer's (Adapted from Hampel et al)

Genetic Etiology of AD

Genetics influence a person's chance of developing Alzheimer's disease, but most cases are not completely inherited. Scientists group genetic involvement into two main categories: deterministic genes and risk genes (Corder et al; Goate et al.).

The first category includes rare inherited mutations that directly cause Alzheimer's disease, usually leading to early-onset Alzheimer's before age 65 (Hardy & Selkoe). These mutations affect how the brain processes a protein called amyloid precursor protein (APP), which is normally cut into smaller pieces. When this cutting process goes wrong, it produces too much beta-amyloid, a sticky protein fragment that clumps together and forms plaques in the brain—one of the main features of Alzheimer's disease (Hardy & Selkoe).

Several specific mutations have been studied in mouse models. For example, the Swedish mutation of the APP gene (APPswe) causes cells to produce higher levels of beta-amyloid. Mice engineered with this mutation develop amyloid plaques much earlier than normal mice and show memory and learning problems (Oakley et al.). Another example is the London APP mutation (V717I), which changes how APP is cut and increases the more harmful type of beta-amyloid. Transgenic mice with this mutation show early plaque buildup and changes in brain cell communication (Goate et al.).

Other deterministic mutations affect the gamma-secretase complex, a group of proteins that also helps cut APP. Important genes here include presenilin-1 (PSEN1) and presenilin-2 (PSEN2). Mutations such as PSEN1 M146V and PSEN1 ΔE9 have been tested in mouse models and were shown to increase production of the more toxic beta-amyloid form (Aβ42). These mice develop plaques earlier, show neuron damage, and perform worse on memory tests (Yokoyama et al., 2022). PSEN2 mutations, such as PSEN2 N141I, have also been modeled in mice and lead to similar plaque formation and brain changes, though often more slowly (Yokoyama et al.).
The second category includes risk genes, which increase the likelihood of developing Alzheimer's disease but do not guarantee it. The most well-known risk gene is the apolipoprotein E epsilon-4 (ApoE4) allele. ApoE helps transport fats and cholesterol in the brain. In mouse studies, animals with the ApoE4 version show reduced clearance of beta-amyloid and more inflammation in brain tissue compared to other ApoE types (Corder et al). In humans, having one copy of ApoE4 increases risk, and having two copies increases it even more — but many people with ApoE4 never develop Alzheimer's, and some people without it still do (Corder et al).

Overall, mouse model research has helped scientists understand how specific gene mutations change protein processing, speed up plaque formation, and damage brain cells. These models are important for testing possible treatments and learning how Alzheimer's disease develops at the biological level (Yokoyama et al; Oakley et al).

Diagnosis

Biomarkers

Biomarkers are measurable indicators that can help detect Alzheimer's disease before symptoms become severe (Hampel et al; Jack et al). One major group of biomarkers is proteins, including amyloid-beta and tau (Murphy & LeVine. 311; Hampel et al). Changes in the levels of these proteins in blood or cerebrospinal fluid can signal early stages of the disease (Palmqvist et al; Brier et al). Tracking these protein changes is important for both diagnosis and research, as they can reveal how the disease is progressing or how a treatment is working (Hampel et al). Biomarkers are becoming an increasingly important tool because they provide a way to detect Alzheimer's at an early stage before significant brain damage occurs (Jack et al; Brier et al).

Current diagnostic approach

Diagnosing Alzheimer's disease can be challenging because the symptoms usually develop slowly and can resemble normal aging or other health conditions (Lane, Hardy, & Schott. 61). Memory lapses, confusion, or difficulty with problem-solving may be the first signs, but they are not always specific to Alzheimer's. To make an accurate diagnosis, doctors use several tools and tests together. One common approach is cognitive testing, which includes memory and thinking assessments to evaluate learning ability, attention, language, and problem-solving skills (Jack et al).

In addition to cognitive tests, brain imaging techniques are often used. MRI scans can show shrinkage in certain areas of the brain, while PET scans can detect the buildup of amyloid plaques, which are a hallmark of Alzheimer's (Brier et al; Villemagne et al. 24). These imaging methods allow doctors to see physical changes in the brain that might explain a person's symptoms.

Doctors may also use blood tests or analyze cerebrospinal fluid from a spinal tap to check for biomarkers linked to Alzheimer's disease. These biomarkers can include abnormal levels of beta-amyloid or tau proteins, which are associated with the progression of the disease (Hansson et al; Hampel et al). Combining these tests helps doctors rule out other medical conditions, such as vitamin deficiencies, thyroid problems, or other forms of dementia, and provides stronger evidence for an accurate Alzheimer's diagnosis.

Overall, diagnosing Alzheimer's requires a careful, step-by-step approach that includes looking at symptoms, performing cognitive tests, examining the brain with imaging, and checking for biomarkers. Using all of these tools together increases the chances of identifying the disease early, which can help patients receive treatment and support sooner (Jack et al).

Proteins

Proteins in the brain and cerebrospinal fluid are some of the most important biomarkers for detecting Alzheimer's disease. Two of the key proteins are amyloid-beta and tau, which are directly involved in the formation of plaques and tangles, respectively. In Alzheimer's disease, amyloid-beta can accumulate in the spaces between neurons, while tau tangles form inside neurons. Measuring changes in the levels of these proteins in cerebrospinal fluid or blood can help identify the disease even before significant symptoms appear. For example, elevated levels

of phosphorylated tau and reduced levels of amyloid-beta 42 in cerebrospinal fluid are strong indicators of Alzheimer's pathology. Tracking these protein levels over time can also help researchers monitor disease progression and test the effectiveness of new treatments. Protein biomarkers are therefore crucial for early detection, diagnosis, and research into therapies for Alzheimer's disease (Murphy & LeVine. 311; Hampel et al).

Mechanism of neurodegeneration

Figure 2. Roles of Astrocytes and Microglia in Neurodegeneration

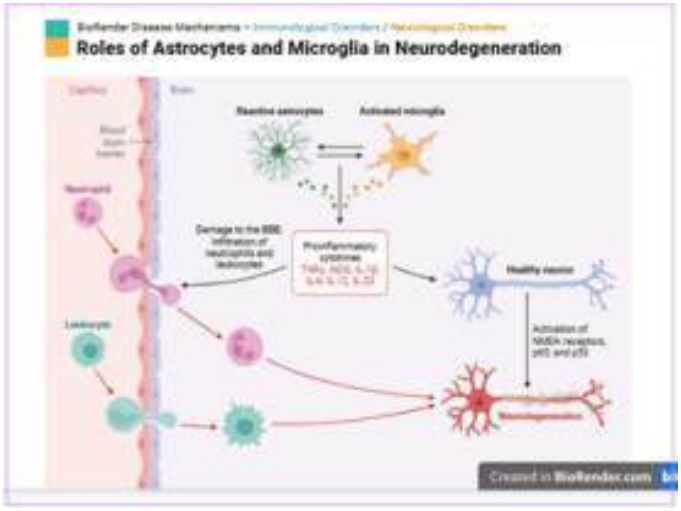

As shown in Figure 2, microglia are important immune cells in the brain.

Amyloid hypothesis

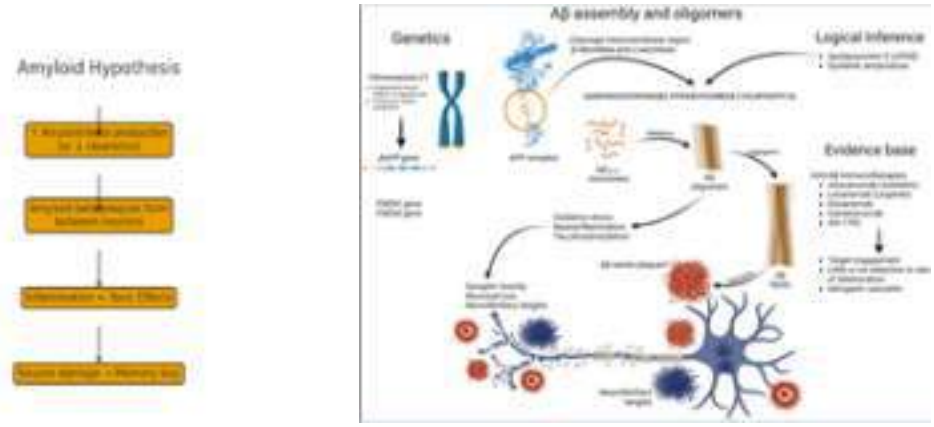

The steps of Amyloid Hypothesis and what it is about.

The amyloid hypothesis suggests that amyloid-beta buildup is the earliest biological event that triggers Alzheimer's disease, often beginning 15–20 years before symptoms appear. Small amyloid-beta oligomers are especially toxic because they interfere with communication

between neurons before joining together to form larger plaques. These plaques activate microglia and astrocytes, the brain's immune cells, leading to chronic neuroinflammation. The inflammatory molecules released by these cells can then promote abnormal changes in tau proteins, allowing tau tangles to spread more rapidly throughout the brain. In many inherited forms of Alzheimer's disease, this sequence of amyloid-beta accumulation - neuroinflammation - tau pathology is very common. However, scientists now recognize that Alzheimer's disease is more complex than originally thought. Rather than being caused by amyloid-beta alone, the progress of the disease is driven by continuous interactions among amyloid-beta, tau, and neuroinflammation (Selkoe & Hardy; Hampel et al; Chen et al 165).

Figure 3. Amyloid-beta aggregation and plaque formation in Alzheimer's disease.

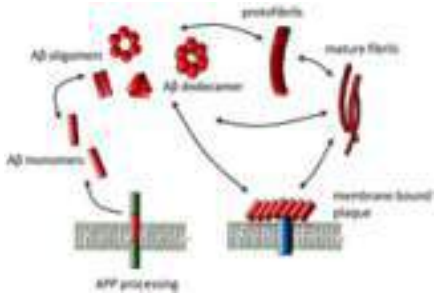

How the oligomers affect the neurons and form large plaques

Figure 4. Amyloid-beta aggregation and fibril formation in Alzheimer's disease.

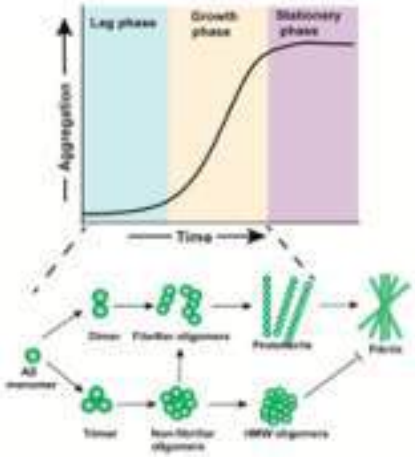

The growth and phases that take place during the formation of the fibrils from the amyloid beta monomers.

Abeta plaques

Abeta plaques are sticky clusters of a protein called amyloid-beta that build up between neurons, the cells in the brain that transmit information (Murphy & LeVine. 311). These plaques interfere with how neurons communicate, which is a key factor in the memory loss and cognitive decline seen in Alzheimer's disease. Amyloid-beta proteins are first produced as single units called monomers. Normally, these monomers float harmlessly in the brain, but under certain conditions, they can stick together to form small clusters called oligomers (Benilova,

Karran, & De Strooper. 349). These oligomers are especially toxic because they can damage neurons and disrupt the signals that allow brain cells to talk to each other.

Over time, the oligomers can join together to form long, thread-like structures called fibrils. Fibrils then accumulate into larger deposits known as plaques. As these plaques grow, they physically block connections between neurons and interfere with the flow of important chemical signals in the brain. In addition to disrupting communication, the plaques can trigger inflammation in brain tissue. This inflammation causes the brain's immune cells to become overactive. While immune cells are supposed to protect the brain, in this case, they can mistakenly attack healthy neurons, causing further damage (x).

Research using brain imaging techniques, such as PET scans, shows that people with high levels of amyloid-beta plaque buildup are more likely to experience memory problems and cognitive decline (Brier et al). Studies also suggest that plaques may appear years before symptoms become noticeable, meaning that amyloid-beta accumulation is an early step in the development of Alzheimer's disease (Selkoe & Hardy). Scientists continue to study Abeta plaques to better understand how they form, why they are toxic, and how targeting them could help prevent or slow the progression of Alzheimer's (Gulisano, Melone, Bernardo, & Ciotti. 195).

Tau tangles

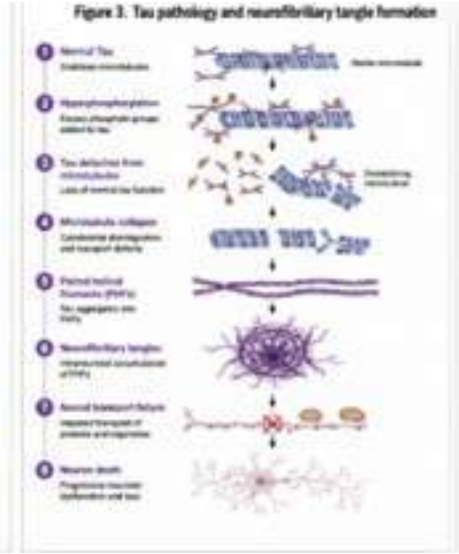

Tau pathology and how Tau tangles form in the brain.

Tau tangles form inside nerve cells when normal tau proteins become twisted and start clumping together (Chen et al. 165). Under normal conditions, tau helps maintain the internal structure of neurons, acting like scaffolding that keeps the cell stable and supports transport pathways for nutrients, proteins, and chemical signals. However, when tau proteins become abnormal, they lose their ability to function properly and instead stick together to form tangles (Villemagne et al. 24). These tangles block the pathways that neurons use to transport essential

materials, which disrupts communication between cells and prevents the neurons from staying healthy.

Unlike amyloid-beta, mutations in the tau gene are not a common cause of typical Alzheimer's disease. Instead, abnormal tau develops mainly because of excessive phosphorylation, a chemical process in which too many phosphate groups attach to the tau protein. This abnormal phosphorylation is driven by several protein kinases, including glycogen synthase kinase-3β (GSK-3β) and cyclin-dependent kinase 5 (CDK5), while phosphatases that normally remove phosphate groups become less active. As a result, tau detaches from microtubules, loses its normal stabilizing function, and begins to misfold into paired helical filaments that eventually accumulate as neurofibrillary tangles. In rare neurodegenerative disorders known as frontotemporal dementia, mutations in the MAPT gene directly cause tau aggregation; however, in Alzheimer's disease, abnormal phosphorylation rather than MAPT mutations is considered the primary mechanism leading to tau pathology (Gulisano et al. 195; Chen et al. 165).

Tau pathology also follows a characteristic pattern throughout the brain. According to the Braak staging mode (Chen et al. 165 & Yu et al. 165) abnormal tau first appears in the transentorhinal cortex and entorhinal cortex, regions that connect the hippocampus with the cerebral cortex. It then spreads into the hippocampus, where memories are formed and consolidated, before progressing to the temporal, parietal, and frontal association cortices. These regions are affected first because they contain highly interconnected excitatory neurons with high metabolic activity, making them especially vulnerable to abnormal tau accumulation and the spread of misfolded tau proteins between neurons. As tau pathology advances through these neural networks, patients gradually develop worsening memory impairment, followed by deficits in language, reasoning, executive function, and eventually widespread cognitive decline (Braak & Braak. 239; Villemagne et al. 24; Hampel et al)

Over time, this disruption causes damage to the neurons and can eventually lead to their death. The spread of tau tangles through the brain follows a predictable pattern, often starting in areas involved in memory, such as the hippocampus, and then spreading to regions responsible for thinking and decision-making (Braak & Braak. 239). Studies have shown that the amount and distribution of tau tangles are closely linked to how severe a person’s cognitive decline becomes. In fact, while amyloid-beta plaques are thought to trigger the disease process, tau tangles are more strongly associated with the progression of symptoms, including memory loss, confusion, and difficulty with problem-solving (Gulisano et al, Melone et al, Bernardo et al, & Ciotti et al). Understanding how tau tangles form and spread is an important area of research for developing treatments that could slow or stop neuron damage in Alzheimer’s disease (Chen et al. 165).

Neuroinflammation

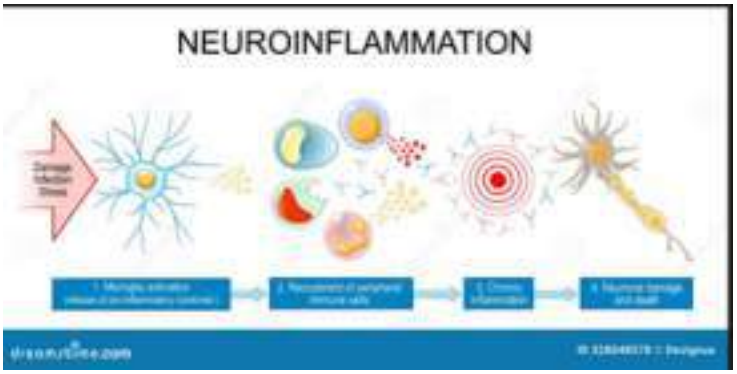

The steps in Neuroinflammation and how it leads to Neuronal damage.

Although amyloid-beta plaques and tau tangles are the most well-known features of Alzheimer's disease, research suggests that chronic neuroinflammation is also a major contributor to disease development and progression (Chen, Li, Liu, & Wang. 165; Heneka, Golenbock, & Latz). Specialized immune cells called microglia normally protect the brain by removing damaged cells, clearing harmful proteins such as amyloid-beta, and responding to injury or infection. Another type of brain cell called astrocytes helps maintain the brain's chemical environment and supports normal neuron function. In Alzheimer's disease, however, both microglia and astrocytes can become chronically overactivated. Instead of protecting the brain, they continuously release inflammatory signaling molecules called cytokines, including tumor necrosis factor-alpha (TNF-α), interleukin-1 beta (IL-1β), and interleukin-6 (IL-6), which damage nearby neurons and weaken the connections between them (Heneka et al).

Although amyloid-beta accumulation has traditionally been considered the first step in Alzheimer's disease, researchers have found that the order of these biological changes is not always the same. In inherited forms of Alzheimer's disease caused by mutations in genes such as APP, PSEN1, or PSEN2, amyloid-beta plaques usually appear first. The buildup of amyloid-beta activates microglia and astrocytes, leading to chronic neuroinflammation. This inflammatory environment promotes abnormal changes in tau proteins, allowing tau tangles to spread more rapidly through the brain and causing progressive neuron loss (Selkoe & Hardy; Hampel et al). However, evidence suggests that some cases of late-onset Alzheimer's disease may begin differently. Aging, reduced blood flow, infections, traumatic brain injury, or genetic risk factors such as ApoE4 can activate microglia before large amounts of amyloid-beta have accumulated. Once these immune cells remain activated for long periods, they become less effective at clearing amyloid-beta from the brain while continuing to release inflammatory molecules. As a result, chronic neuroinflammation can actually cause more amyloid-beta accumulation as well as abnormal tau phosphorylation, eventually leading to plaque formation, tau tangles, and neuron death (Heneka et al; Chen et al. 165).

Because both pathways have been observed, many scientists no longer believe that every patient follows one single sequence of disease progression. Instead, Alzheimer's disease is increasingly viewed as a cycle. Amyloid-beta plaques stimulate neuroinflammation, neuroinflammation promotes additional amyloid-beta buildup and tau pathology, and tau tangles

damage neurons, releasing signals that activate even more inflammation. Each process strengthens the others, creating a positive feedback loop that accelerates neurodegeneration and cognitive decline over time (Chen et al. 165; Hampel et al; d'Errico, Meyer-Luehmann, & Schilling. 265).

Reduced blood flow to the brain can also limit the delivery of oxygen and nutrients, making neurons more vulnerable to injury (Iadecola). In addition, oxidative stress, which involves harmful molecules called free radicals which are highly reactive, can damage brain cells and accelerate degeneration (Chen et al. 165). In addition to the biological mechanisms underlying Alzheimer's disease, several genetic, lifestyle, and environmental factors contribute to an individual's risk of developing the condition. Genetic predisposition, particularly inherited variants associated with Alzheimer's disease, can significantly influence susceptibility. Lifestyle factors, including poor dietary habits, physical inactivity, smoking, and inadequate cardiovascular health, have also been linked to an increased risk of cognitive decline.

Furthermore, environmental exposures, such as air pollution and certain neurotoxic substances, may contribute to neurodegeneration by promoting inflammation and oxidative stress. Although some risk factors, such as genetics and age, are non-modifiable, many lifestyle-related factors can be addressed through preventive interventions, highlighting the importance of healthy behaviors in reducing Alzheimer's disease risk (Kivipelto, Mangialasche, Ngandu, & Solomon. 398). Because the processes are highly interconnected, scientists now view Alzheimer’s disease as a complex condition with multiple causes rather than a disorder caused by a single protein or genetic mutation (Chen et al. 165; Heneka et al).

Figure 5. The role of neuroinflammation in the development and progression of Alzheimer’s disease.

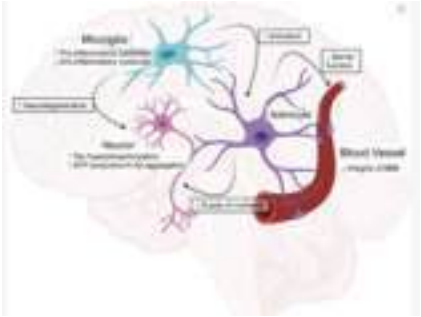

A chain showing how inflammation in the brain leads to Alzheimer's

MRI/PET scans

Brain imaging is another key tool in diagnosing and monitoring Alzheimer’s disease. MRI (Magnetic Resonance Imaging) scans can detect shrinkage in brain regions that are important for memory and thinking, such as the hippocampus and cortex. PET (Positron Emission Tomography) scans can detect the buildup of amyloid-beta plaques and tau tangles, providing a visual representation of the disease in the brain. These imaging techniques allow doctors to assess how far the disease has progressed and can help guide treatment decisions.

Brain scans also play a critical role in research by helping scientists study the effectiveness of new therapies over time.

Treatment of AD

Treatment of Alzheimer's Disease

Antibody-Based Approach

One of the main approaches to treating Alzheimer's disease (AD) involves antibody-based therapies, which are designed to target specific proteins involved in the development of the disease. These laboratory-produced antibodies target amyloid-beta, a protein that can accumulate in the brain and form plaques, one of the key pathological features of Alzheimer's disease. The antibodies bind to amyloid-beta and help the immune system recognize and remove these plaques from the brain. This approach is considered a disease-modifying strategy because it aims to target an underlying feature of Alzheimer's disease rather than only treating its symptoms. Although antibody therapies do not cure the disease, they represent an important development in efforts to slow its progression and preserve cognitive function.

Current Treatments

Currently, there is no cure for Alzheimer's disease, so treatment focuses on managing symptoms, maintaining cognitive function, and improving quality of life. Medications can help address memory and thinking problems, while other treatments may be used to manage behavioral and psychological symptoms such as anxiety or depression. Non-drug approaches, including regular exercise, a healthy diet, social engagement, and structured memory routines, may also support overall cognitive and daily functioning (Kivipelto et al. 398). In addition to these approaches, disease-modifying antibody therapies are now being used for some individuals with early Alzheimer's disease. These treatments represent a shift in Alzheimer's care from primarily managing symptoms toward also attempting to slow disease progression.

Latest Drugs – Antibodies

Recent advances in Alzheimer's research have led to the development of antibody-based drugs such as lecanemab, which specifically targets amyloid-beta in the brain. Clinical research has shown that lecanemab can reduce amyloid-beta plaque levels and modestly slow cognitive and functional decline in individuals with early Alzheimer's disease (van Dyck et al. 9) However, these treatments do not reverse existing brain damage or cure Alzheimer's disease, and they may not be appropriate for every patient. Their development nevertheless represents an important step toward disease-modifying treatment. Alongside antibody therapies, researchers are also investigating emerging approaches such as stem-cell therapy, which aims to replace or repair damaged neurons and restore neural connections involved in memory and learning (Ou et

al). Small-molecule drugs are also being studied for their potential to enter brain cells and target abnormal proteins, inflammation, and other mechanisms involved in Alzheimer's disease.

Stem-cell based

Stem-cell therapy is an emerging strategy that aims to replace damaged neurons with healthy new cells. Stem cells have the unique ability to develop into different types of cells, including neurons, which makes them a potential tool for repairing brain damage caused by Alzheimer's disease. In experimental studies, stem cells can be guided to integrate into the brain and restore some lost neural functions, including synaptic connections that are critical for learning and memory. While this approach is still largely in the research and clinical trial stage, it offers hope for not just slowing the progression of Alzheimer's disease, but potentially restoring some lost cognitive function in the future (Ou et al).

Small-molecule based approach

Small-molecule drugs represent a novel and promising approach in the treatment of Alzheimer's disease. These drugs are small enough to easily enter brain cells, allowing them to directly interact with abnormal proteins such as amyloid-beta and tau, which are central to the development and progression of Alzheimer's. By targeting these harmful proteins, small-molecule therapies aim to prevent their aggregation into plaques and tangles, which disrupt communication between neurons and contribute to cognitive decline. In addition to targeting proteins, some small-molecule drugs are designed to reduce neuroinflammation, which occurs when immune cells in the brain become overactive and mistakenly damage healthy neurons. Because these drugs can penetrate brain tissue efficiently, they are able to reach affected areas and potentially slow or stop the progression of the disease at a cellular level. Researchers are currently testing a variety of small-molecule compounds in both laboratory and clinical settings to determine their safety and effectiveness. While still largely experimental, this approach provides hope that Alzheimer's disease could one day be treated more effectively, rather than just managing its symptoms (Chen et al. 165; Hampel et al; Selkoe & Hardy)

Conclusion

Alzheimer's disease is a progressive brain disorder that causes memory loss, confusion, and cognitive decline, most often in older adults. The disease involves several biological processes, including amyloid-beta plaques, tau tangles, and neuroinflammation, as well as genetic, environmental, and lifestyle factors. Although scientists continue to study which process begins first, research shows that amyloid-beta, tau, and neuroinflammation are closely connected. In some cases, amyloid-beta buildup appears first and triggers neuroinflammation and tau pathology. In other cases, chronic neuroinflammation may occur earlier and contribute to amyloid-beta accumulation and abnormal tau changes. Regardless of which process begins first, these mechanisms eventually interact in a cycle that speeds up neuron loss and cognitive decline. Because these processes are so closely linked, it remains difficult to determine exactly how

Alzheimer's disease begins and which biological changes should be targeted first. This complexity has made it challenging to develop treatments that consistently slow or stop disease progression.

Understanding how these biological processes work together has helped scientists develop better ways to detect and treat Alzheimer's disease. New biomarkers, including blood tests, cerebrospinal fluid analysis, and brain imaging, can help identify the disease before major symptoms appear. Researchers are also developing new treatments, such as antibody therapies, stem cell approaches, and small-molecule drugs, that target amyloid-beta, tau, and neuroinflammation. However, another challenge is diagnosing the disease early enough for these treatments to be most effective. Many people are not diagnosed until memory loss and other symptoms become noticeable, by which point significant brain damage has already occurred. Although there is currently no cure for Alzheimer's disease, continued research offers hope for earlier diagnosis, more effective treatments, and new ways to slow or prevent the progression of the disease in the future.

Works Cited

Benilova, I., et al. "The Toxic Aβ Oligomer and Alzheimer's Disease: An Emperor in Need of Clothes." Nature Neuroscience, vol. 15, no. 3, 2012, pp. 349–357.

Braak, H., and E. Braak. "Neuropathological Staging of Alzheimer-Related Changes." Acta Neuropathologica, vol. 82, no. 4, 1991, pp. 239–259.

Brier, M. R., et al. "Tau and Aβ Imaging, CSF Measures, and Cognition in Alzheimer's Disease." Science Translational Medicine, vol. 8, no. 338, 2016, article 338ra66.

Chen, Y., et al. "Tau and Neuroinflammation in Alzheimer's Disease: Interplay Mechanisms and Clinical Translation." Journal of Neuroinflammation, vol. 20, 2023, article 165.

Corder, E. H., et al. "Gene Dose of Apolipoprotein E Type 4 Allele and the Risk of Alzheimer's Disease." Science, vol. 261, no. 5123, 1993.

d'Errico, P., et al. "Mechanisms of Pathogenic Tau and Aβ Protein Spreading." Frontiers in Aging Neuroscience, vol. 12, 2020, article 265.

Goate, A., et al. "Segregation of a Missense Mutation in the Amyloid Precursor Protein Gene with Familial Alzheimer's Disease." Nature, vol. 349, no. 6311, 1991

Gulisano, W., et al. "Role of Amyloid-β and Tau Proteins in Alzheimer's Disease." Frontiers in Aging Neuroscience, vol. 10, 2018, article 195.

Hampel, H., et al. "The Amyloid-β Pathway in Alzheimer's Disease." Molecular Psychiatry, vol. 26, 2021,

Hansson, O., et al. "CSF Biomarkers of Alzheimer's Disease Concord with Amyloid-β PET and Predict Clinical Progression." Alzheimer's & Dementia, vol. 14, no. 11, 2018, pp

Hardy, J., and D. J. Selkoe. "The Amyloid Hypothesis of Alzheimer's Disease: Progress and Problems on the Road to Therapeutics." Science, vol. 297, no. 5580, 2002,

Iadecola, C. "The Pathobiology of Vascular Dementia." Neuron, vol. 80, no. 4, 2013

Jack, C. R., Jr., et al. "NIA-AA Research Framework: Toward a Biological Definition of Alzheimer's Disease." Alzheimer's & Dementia, vol. 14, no. 4, 2018

Kivipelto, M., et al. "Lifestyle Interventions to Combat Cognitive Decline and Alzheimer's Disease." Lancet Neurology, vol. 17, no. 5, 2018, pp. 398–412.

Lane, C. A., et al. "Alzheimer's Disease." European Journal of Neurology, vol. 25, no. 1, 2018, pp. 59–70.

Murphy, M. P., and H. LeVine. "Alzheimer's Disease and the Amyloid-Beta Peptide." Journal of Alzheimer's Disease, vol. 19, no. 1, 2010, pp. 311–323.

National Institute on Aging. "What Happens to the Brain in Alzheimer's Disease?" U.S. Department of Health & Human Services, 2024.

Oakley, H., et al. "Intraneuronal β-Amyloid Aggregates, Neurodegeneration, and Neuron Loss in Transgenic Mice with Five Familial Alzheimer's Disease Mutations." Journal of Neuroscience, vol. 26, no. 40, 2006

Ou, C.-M., et al. "Stem Cell Therapy in Alzheimer's Disease: Current Status and Perspectives." Frontiers in Neuroscience, vol. 18, 2024, article 1440334.

Palmqvist, S., et al. "Discriminative Accuracy of Plasma Phospho-Tau217 for Alzheimer

Disease vs Other Neurodegenerative Disorders." JAMA, vol. 322, no. 6, 2019
Selkoe, D. J., and J. Hardy. "The Amyloid Hypothesis of Alzheimer's Disease at 25 Years." EMBO Molecular Medicine, vol. 8, no. 6, 2016
van Dyck, C. H., et al. "Lecanemab in Early Alzheimer's Disease." The New England Journal of Medicine, vol. 388, no. 1, 2023, pp. 9–21.
Villemagne, V. L., et al. "Tau Imaging: Early Progress and Future Directions." Lancet Neurology, vol. 12, no. 1, 2013, pp. 24–34.
Yokoyama, M., et al. "Mouse Models of Alzheimer's Disease." Frontiers in Molecular Neuroscience, vol. 15, 2022, article 912995.

Hydrogel Bioinks: Addressing Cell Density, Viability, and Resolution Trilemma in 3D Bioprinting By Beliz Kasacı

Abstract

The goal of 3D bioprinting is to produce living tissues that have the cellular density and structural organization required for future clinical uses such as tissue repair and organ replacement. As the cell density increases, the printing process finds it harder to keep cells alive and to reproduce fine features. This problem leads to what is known as a density-viability-resolution trilemma. This review discusses this issue in the context of 3D bioprinting using hydrogels and sets out three criteria for it, which are a cell concentration of at least 20 million cells per milliliter, a cell viability of at least 80 percent, and features of 50 micrometers or smaller in size. In light-based bioprinting, cells that are closely packed scatter the light and this reduces the accuracy of printing. Refractive-index matching can reduce this effect and thus makes it possible to produce features of about 50 micrometers at a cell density as high as 100 million cells per milliliter. Another method involves the use of computational correction. However, these have so far mostly been tested at lower cell densities. When five recently printed constructs are compared, it turns out that none of them can be shown to satisfy all three criteria at the same time, and four of them do not report at least one of the required measurements. As the size of the constructs increases, further difficulties arise because of problems with nutrient delivery, mechanical support and long-term cell survival. The results indicate that improvements in resolution can be achieved, but it remains unsolved how to keep high density, high viability and fine resolution together in larger functional tissues.

Introduction

One of the major issues yet remaining is the lack of available donor organs in the field of transplantation. By September 2026, 110,449 people were on the organ transplant waiting list in the United States, and recent figures from Türkiye indicate that about 32,000 people were waiting for an organ transplant, and over 20,000 of them were waiting for a kidney (United Network for Organ Sharing; T.C. Sağlık Bakanlığı Konya Şehir Hastanesi). Three-dimensional bioprinting is a promising way of reducing the reliance on donated organs. It is a technique that involves using cells and biomaterials to build up tissue-like structures one layer at a time. Scientists have already applied bioprinting to the creation of liver models, cardiac structures, spinal-cord scaffolds and other tissues (Ma et al.; Lee et al.; Koffler et al.). However, it is much more difficult to produce a whole organ that is able to survive, carry out its function and eventually serve as a replacement for a donated organ.

A reason for this is that a printed organ has to satisfy multiple requirements at the same time. For instance, it has to have a high enough cell density to serve as a biological tissue, the cell has to stay alive both during and after the printing process, and the printing process itself must accurately reproduce small structures. These requirements create what is called the “density-viability-resolution” challenge or “trilemma” that summarizes the requirements that

ensure a usable printed organ (Fig. 1). In the context of this review, high cell density is taken to mean at least 20 million cells per milliliter, high viability at least 80%, and high resolution 50 micrometers or finer. This field is striving to meet all three of these conditions at once through stepwise explorations into each aspect of the trilemma.

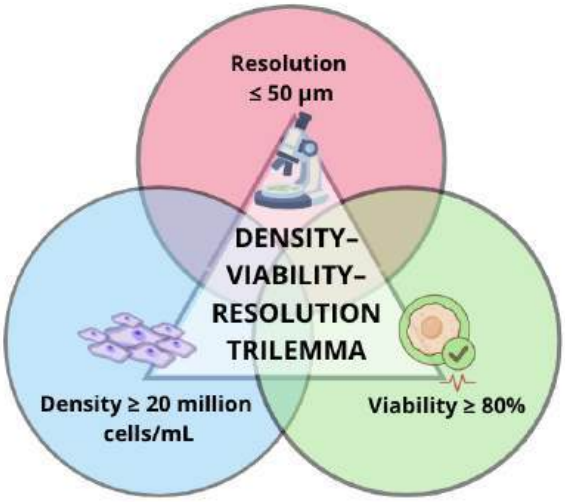

Fig. 1: The density–viability–resolution trilemma in 3D bioprinting, with the three thresholds being ≥20 million cells/mL for density, ≥80% for viability, and ≤50 µm for resolution. These thresholds are used in this review, and the trilemma diagram is based on the work of You et al. ("High Cell Density" 3).

This review primarily deals with the resolution perspective of the trilemma, in which resolution refers to the smallest feature that a bioprinter can accurately create. In the case of light-based bioprinting, closely packed cells can often scatter the light used to form the structure, which leads to a reduction in the precision of the printed features (You et al., "High Cell Density" 2). To mitigate this issue, materials are added to reduce the optical difference between the cell and surrounding bioink, which is a printable mixture that combines a material matrix with cells or other biological components (Moroni et al. 385). This approach is known as refractive-index matching (You et al., "High Cell Density" 2). Nevertheless, improving the resolution does not resolve the other issues concerning cell survival, structural stability, and scale.

Refractive-index matching is currently the most effective route to improve the resolution aspect of the density–viability–resolution trilemma. To date, no study has demonstrated a single method that maintains high density, high viability and fine resolution at the same time, which is why a transplantable organ does not yet exist in spite of bioprinted livers, hearts and spinal-cord scaffolds. The next sections begin first explaining the trilemma and the optical reason for loss of resolution, then proceed to compare refractive-index matching with other methods, evaluate the tissue constructs that have been produced, and finally address what is still not resolved.

Background

Bioprinting matters because it lets scientists build three-dimensional tissue models and control where cells and supporting materials go. Researchers use these structures in regenerative medicine, such as repairing or replacing damaged tissues, and as lab models for drug testing and toxicology (Murphy and Atala 773). The goal is not just to copy the outer shape of a tissue; the printed tissue also needs to match key features of the biological environment marked by different cell types and extracellular materials so it can function as intended (Murphy and Atala 773–74).

To make these tissue models, researchers start by creating a digital design that defines the intended shape and layout. They then convert this design into a series of layers that guide the printer in building the object. Next, researchers select the right cells and materials, and then build the construct for later use or lab study (Murphy and Atala 774–75). The properties of the printable material reflect how different bioprinting methods work. The medium that contains the cells during printing is commonly a hydrogel, which can also help the deposited material keep its form. The qualities needed for the hydrogel are method-dependent. For example, extrusion processes require the material to pass through a nozzle and stay stable once deposited, whereas inkjet processes often require lower-viscosity materials (Moroni et al. 386, 391). Inkjet bioprinting has been applied to the printing of skin and cartilage, while microextrusion has created structures including aortic valves and branched vascular tissues (Murphy and Atala 776–77).

Thus, bioink design must balance printer requirements with cell needs. Bioprinting materials can come from natural polymers like gelatin and collagen, or from synthetic substances such as polyethylene glycol (PEG) (Murphy and Atala 778). Synthetic materials can be more precisely adjusted to achieve certain physical properties, while natural materials can provide features that mimic the extracellular environment (Murphy and Atala 778, box 1). For instance, You et al. employed gelatin methacrylate (GelMA) as the hydrogel material in one cell-containing bioink formulation (“High Cell Density” 3). Another useful example is polyethylene glycol diacrylate (PEGDA), a synthetic PEG-based component that can be added into cross-linked hydrogel formulations. Unagolla and Jayasuriya reported that adding unmodified gelatin and hyaluronic acid to a PEGDA-containing formulation improved the extrusion behavior of the material and created a more uniform final structure (fig. 3A).

Light-based bioprinting differs from techniques that rely on a nozzle to physically lay down material (You et al., “High Cell Density” 2). In some photosensitive bioinks, light exposure solidifies sections, forming the desired structure. This general idea is used in techniques such as stereolithography, DLP (digital light processing), two-photon polymerization, and volumetric printing (You et al., “High Cell Density” 1–2). When adjusted properly, these techniques can generate finer features than extrusion methods since light can be manipulated optically. DLP's appeal is that it can combine precision patterning with reasonably fast fabrication. However, the resolution observed under ideal conditions is not necessarily preserved when cells are added to the bioink. This puts a significant constraint on the printing of densely packed tissues (You et al., “High Cell Density” 2) and leads to the trade-offs among density, viability, and resolution covered in the following section.

The Density–Viability–Resolution Trilemma

Achieving optimal levels of cell density, cell viability, and resolution in a 3D-printed tissue simultaneously is challenging because of their biological and mechanical requirements since improving one can make it harder to maintain the others (You et al., "High Cell Density" 2-3).

High cell density in bioprinting is beneficial because tissues require sufficient cells to act and operate more closely like real biological tissue (You et al., "High Cell Density" 1). Native human tissues can contain approximately 1–3 billion cells/mL, demonstrating how densely packed functional tissues typically are (Enderle and Bronzino 287). Although a sufficient number of cells is necessary for synthetic tissue to mature and function, many bioprinting studies report cell densities of millions rather than billions of cells per milliliter (Zhang et al., 17). Thus, there is a considerable disparity between the cellular density of many synthetic and biological tissues, which in turn drives modern efforts to build tissues that are substantially denser. Instead of suspending individual cells in a large volume of hydrogel, Skylar-Scott and colleagues compacted small tissue units produced from stem cells into a dense matrix using the SWIFT (Sacrificial Writing Into Functional Tissue) approach, then introduced channels that could support perfusion (2). This suggests that higher cell densities are achievable, but it also introduces a new problem of transferring efficient oxygen and nutrient delivery to tightly packed tissues.

At the same time, the printing procedure needs to keep the cells alive and healthy during and after manufacturing (Murphy and Atala 773). The physical parameters applied during printing directly affect cell viability. Billiet et al., for instance, showed that HepG2 cell viability decreased as printing pressure increased. This indicates that the conditions required to extrude a bioink via a nozzle can also harm the cells within it (Billiet et al., 59). In addition, another problem is keeping cells alive after printing. Oxygen and nutrients are only able to travel efficiently over restricted lengths in tissue, roughly 200–300 μm. Therefore, cells deep inside thicker constructs may not receive adequate supply without vascular networks (You et al., "High Cell Density" 1). This challenge is particularly relevant since the goal is to create thick and highly cellular tissues rather than thin laboratory constructs.

Resolution is also essential since bioprinters must be capable of generating small and accurate structures, especially features such as vascular networks (You et al., "High Cell Density" 1). Both extrusion-based printing and light-based printing have distinct resolution problems. In extrusion printing, higher cell densities can require larger nozzles because smaller nozzles create more shear stress on the cells. For example, at cell densities ≥10 million cells/mL, nozzle sizes of 200 μm or bigger may be required, limiting printing resolutions to approximately 200–500 μm (You et al., "High Cell Density" 2). So, maintaining cell viability by utilizing a larger nozzle leads to less intricate structures. In light-based printing, dense cells scatter the light, which blurs the projected shape and reduces printing resolution. This phenomenon is observed in Figure 1E of "High Cell Density" when an acellular (no cells) bioink produces detailed spoke and snowflake shaped features, but a bioink with 100 million cells/mL loses much of its

complexity. If the refractive index is matched, several small features become visible again (You et al., “High Cell Density” 3). In light-based printing, some additives used to limit unwanted curing from scattered light can also be toxic to cells, and using thinner layers to improve resolution increases printing time, which can also harm cell viability (You et al., “High Cell Density” 2, 8).

Achieving high resolution becomes more difficult as the cell density increases (You et al., “High Cell Density” 2-3). For this reason, this review follows resolution as the main corner of the trilemma since increased cell density can directly reduce printing accuracy. Improving resolution could overcome one of the main challenges of printing densely cellular structures, although the viability problem would still be present.

How Cells Reduce Printing Resolution

In light-based bioprinting, the cells present in the bioink are able to alter the light used to produce the pattern. Cytoplasmic refractive index values are about 1.36 to 1.39, the exact value depending on the type of cell. By contrast, hydrogel bioinks generally have a refractive index of around 1.33 (You et al., “High Cell Density” 2). As these values are different, when light reaches a cell it is diverted instead of proceeding in its original direction. Consequently, a single cell can act like a small optical lens, bending part of the light that enters it (You et al., “High Cell Density” 2-3). The greater the number of cells packed into the bioink, the more of the projected light is diverted, and so the intended pattern becomes less accurately defined and the printing resolution decreases (You et al., “High Cell Density” 2-3).

A method of understanding how this problem might be addressed is based on the research carried out by Boothe and colleagues, who looked at refractive-index matching in the context of live microscopy rather than in bioprinting. They found that the difference between a biological sample and the liquid surrounding it can cause distortion of light and result in poorer image quality. Boothe and colleagues used iodixanol, a water-soluble substance which has a relatively high refractive-index value, to modify the surrounding liquid to bring its optical value nearer to that of the sample. Their experiments proved that such matching could improve image resolution (Boothe et al., 3). Even though their study did not include bioinks, it did serve as an earlier example of the physical principle that Chen’s group later applied to cell-containing hydrogels.

You and colleagues applied the same principle in light-based bioprinting by incorporating iodixanol into the hydrogel to make its optical properties more similar to those of the cell cytoplasm. As a result, light scattering caused by the difference between the cells and the surrounding material was reduced (You et al., “High Cell Density” 2-3). This was particularly beneficial when a large number of cells were involved. At a cell density of 100 million cells per milliliter, the unmatched bioink lost most of the detail in the intended spoke and snowflake patterns. After the optical mismatch had been reduced, features of about 50 μm could once again be distinguished, and the researchers found that scattering had decreased by roughly a factor of ten (You et al., “High Cell Density” 2–3). Nevertheless, this approach cannot eliminate scattering

from structures inside the cell, such as the nucleus and the organelles, so it only reduces the problem rather than getting rid of it entirely (You et al., “High Cell Density” 8).

There is another chemical method that illustrates how improving the optics can lead to a biological issue. You et al. refer to TEMPO, a chemical additive that can prevent unwanted polymerization caused by scattered light. Polymerization is the reaction which causes the light-sensitive bioink to solidify. If it takes place outside the desired area, then the edges of the print may become less accurate. TEMPO can help to mitigate this effect, but You et al. point out that chemicals used in this way are usually cytotoxic, that is to say they may damage or kill living cells (You et al., “High Cell Density” 2). This is a clear instance of the trilemma, since a modification that improves resolution can reduce viability. Iodixanol is particularly useful in this context since You et al. found it to have little effect on viability, cell growth, or phenotype, even though its optical correction still has limitations (You et al., “High Cell Density” 8).

Improving Resolution: Refractive-Index Matching and Computational Approaches

A way of dealing with the resolution problem is to alter the optical characteristics of the bioink. As explained in the section before, You and colleagues employed iodixanol (IDX) to reduce the difference in refractive index between the bioink and the cell cytoplasm. Iodixanol thus represents the most effective chemical example discussed here, even though it does not eliminate all the scattering sources within the cells.

Iohexol (IHX) also offers a similar chemical approach. Xiang and colleagues employed IHX to increase the refractive index of GelMA bioinks containing 50 million cells per milliliter and discovered that higher concentrations of IHX decreased cell-induced scattering and enhanced printing fidelity (Xiang et al., 3960). Furthermore, they were able to print a perfusable microtube with a wall thickness of approximately 100 μm at the same cell density (Xiang et al., 3964). The technique thus functioned in a dense, hollow structure and not just in a simple test pattern. Nevertheless, IHX also caused the hydrogel to soften and it was necessary to modify the printing conditions (Xiang et al. 3959-3960), indicating that an optical improvement can give rise to a new material issue.

Another approach is to adjust the projected light pattern rather than altering the bioink. You et al. employed machine learning to produce modified grayscale masks that compensated for scattering in light-based printing (You et al., “Mitigating Scattering Effects” 081002-4). Guan and colleagues then used deep learning, which is a type of machine learning that learns patterns from training data, with cell-containing bioinks. Their method created corrected light patterns that enhanced the accuracy of each printed layer (Guan et al., figs. 5–6). However, the experiments were carried out using 10 million cells per milliliter, which is below the 20 million cells per milliliter benchmark referred to in this review, and the extension to 3D multilayer structures was suggested rather than actually demonstrated (Guan et al. 4–5). While the computational method shows promise, the evidence for its effectiveness at high cell density is still more limited than that for the chemical methods.

Bernal et al. also found that iodixanol reduced scattering in volumetric bioprinting and achieved features of about 50 μm, even though their cell density was below the 20 million cells/mL threshold. This lends support to You and colleagues' reading of Bernal's work as showing that optical tuning can reduce the issue rather than merely leaving it unchanged (Bernal et al. 5). Mierke reports that Bernal's work acknowledges that light scattering still represents a limitation in volumetric bioprinting (37). Yet the results obtained by Bernal et al. indicate that optical tuning can reduce the problem and enhance printing resolution (fig. 3), even though it does not completely eliminate scattering (Bernal et al. 5). Refractive-index matching currently has the more substantial direct evidence regarding high-density printing. That said, none of these studies demonstrates a single method that can maintain high density, high viability, and fine resolution all at the same time. The various available methods do alleviate the resolution problem, but they have not yet eliminated the broader trilemma. The resolution corner of the trilemma is the one that has moved while density and viability have not.

Current Innovations

Five constructs from the field are assessed below in relation to the three criteria established in the Introduction, and the results are compiled in Table 1.

Ma and colleagues developed a liver-like model by using a two-step DLP bioprinting method in which different types of cells were arranged in a controlled pattern. The researchers first arranged the immature liver-lineage cells which had been derived from human iPSCs, and then added a complementary pattern consisting of HUVECs (blood-vessel-lining cells) and ADSCs (supporting stem cells from fat tissue) to mimic the structure of a liver lobule (Ma et al., 2207). Approximately 76% of the cells remained viable two hours after printing, although this figure fell to about 65% after seven days (Ma et al., 2208). The viability reported was below the 80% target used in this review, and neither the final overall cell density nor the minimum feature size was provided. For these reasons, the model shows useful spatial control without fully resolving the density–viability–resolution problem.

Lu and colleagues then raised the density of liver cells substantially by using DLP printing together with primary hepatocytes, which are liver cells that were isolated directly from mice. They compared three different densities of 5, 20, and 80 million cells per milliliter, the 80 million cells per milliliter being the high-density group (Lu et al., fig. 1b). Iodixanol was added to reduce the optical problems caused by the large number of cells in the bioink such as cell-induced light scattering and to improve printing quality at high cell density (Lu et al., fig. 1a). At the highest cell density, viability was about 80% through day 7, but at low density, the viability dropped to 50% by day 7 (Lu et al., figs. 1e–f). At the highest cell density, this construct thus met both the density threshold and viability threshold for seven days. This result strengthens the case for printing hepatocytes at high cell density: the cells survived better when there were more of them, which is consistent with the need for cell-to-cell contact described above. Since the study did not specify a minimum feature size for this liver construct, its resolution cannot be evaluated in relation to the 50 μm benchmark.

Koffler and colleagues printed a 2-mm scaffold with 200 μm guidance channels meant to direct nerve growth. They then introduced neural progenitor cells (NPCs), which are immature cells capable of developing into different types of cells in the nervous system (Koffler et al., fig. 1). The scaffolds were inserted into rats after the spinal cord had been fully cut at the T3 level. The channels that contained the cells enabled axons, the long outgrowths of nerve cells that transmit signals, to grow across the area of injury (Koffler et al., figs. 2–4). Although their printing method was said to have a resolution of 1 μm, the actual guidance channels in the scaffold were 200 μm wide (Koffler et al., fig. 1). The first figure shows the optical resolution of the printer and the second gives the feature size that the authors decided to use. Yet the paper did not give the cell density in units of cells/mL or state the cell viability as a percentage, so the density and viability aspects of the trilemma cannot be directly compared using the same standards.

Another trade-off is found in the way Lee et al. use FRESH (Freeform Reversible Embedding of Suspended Hydrogels), a technique which holds soft materials in a temporary bath during the printing process. The researchers used collagen, a structural protein present in tissues, to print heart-related structures such as vessels, a valve, and a cell-containing ventricle (Lee et al., 2–4). The ventricle included cardiomyocytes (heart muscle cells) as well as supporting cardiac fibroblasts (Lee et al., fig. S11C–H). A cardiac test print kept approximately 96% viability and collagen filaments as thin as about 20 μm were obtained (Lee et al., figs. S11B and 1I–J). Nevertheless, the cell density was not given in units of cells per mL, and therefore all three of the trilemma targets could not be shown together.

Skylar-Scott and colleagues applied SWIFT (Sacrificial Writing Into Functional Tissue) to form channels within highly dense living tissue. Their matrix attained a density of about 200 million cells/mL, and the cardiac version reached about 240 million cells/mL (Skylar-Scott et al., 3). This was clearly above the density threshold. Yet the narrowest channels that could be reliably printed were about 400 μm in width, which is much wider than the 50 μm resolution requirement (Skylar-Scott et al., 3). The study also failed to provide a single overall viability figure. Thus, although SWIFT achieved a very high cell density it did so at the expense of finer resolution.

These five examples show a reporting pattern in the field as well as differences in what has actually been constructed. In four of the five constructs, at least one of the three measurements is not reported, meaning that none of the constructs fully achieves all three objectives (Table 1). Some studies attained high precision, and others achieved higher cell densities or kept strong cell viability, but no study clearly succeeded in meeting all three objectives at the same time. Even though certain constructs were tested in animals, none of the studies resulted in a fully transplantable bioprinted organ. The major challenge is still to combine high density, high viability, and fine resolution in a single functional tissue construct.

Table 1: Comparison of Selected Bioprinting Studies

Paper	Method	Density	Smallest feature	Viability	Main trade-off / limitation
Ma et al. (2016)	DLP liver model	Final density not reported	Not reported	~76% initially; ~65% at day 7	Precise patterning, but viability decreased
Lu et al. (2025)	DLP + iodixanol	80 million cells/mL	Not reported	~80% through day 7; ~70% at day 10	High density and week-long viability, but resolution not reported
Koffler et al. (2019)	μCPP spinal scaffold	Not reported	1 μm printer resolution; 200 μm channels	Not reported	Fine printing capability, but functional channels were larger
Lee et al. (2019)	FRESH collagen printing	Not reported	~20 μm	~96%	Fine resolution and viability, but density not reported
Skylar-Scott et al. (2019)	SWIFT	~200–240 million cells/mL	~400 μm channels	Not reported as %	Very high density, but coarse channel resolution

Limits of Refractive-Index Matching

Although the review focused on resolution, improving resolution alone does not solve every part of the trilemma. Refractive-index matching can reduce the scattering caused by closely packed cells, but it is still necessary for a printed tissue to keep the cells alive as the construct increases in size. Although Lu and colleagues got closest, their construct was small, and its viability was assessed over a period of about a week. The tissue also has to retain its shape for a sufficient time to develop into functional tissue. Thus, these problems are especially significant when going from the small structures mentioned earlier to human-sized tissues. The aspect of the trilemma concerning viability on a large scale and over time is the one that remains standing.

Kang and colleagues tackled a part of the problem with regard to scale by using an integrated tissue organ-printer. They not only relied on a soft hydrogel containing cells but also printed the cell-laden hydrogel together with a biodegradable polymer, which provided additional mechanical support. Also, they added small internal channels to help the delivery of nutrients to the cells located deeper within the printed tissue (Kang et al., 316). This system allowed the researchers to create examples of bone, cartilage and skeletal muscle on human scale (Kang et al., 314-317. The results indicate that when the size of a construct increases new requirements arise that go beyond simple resolution. Although a hydrogel can provide a proper environment for cells other materials are still needed to ensure that a large printed structure remains stable.

Another approach can be to decrease the reliance on a conventional hydrogel bioink. For instance, Ayan and colleagues created a technique called aspiration-assisted bioprinting, in which a small pipette is used to pick up and place the spheroids, which are three-dimensional clusters of

cells. They showed examples of both scaffold-based and scaffold-free arrangements. In the case of the scaffold-free method, the spheroids were temporarily supported during printing and allowed to join together afterward when the supporting material had been removed (Ayan et al., 2). Moreover, they were able to stack the spheroids without using permanent gel support (Ayan et al., 6). This method indicated that it is not necessary for cells to stay suspended within a traditional hydrogel ink in order to form a three-dimensional structure. The two studies do not give the three measurements referred to in Table 1. Instead they focus on scale and structural support rather than the trilemma itself.

Still, the trilemma is not entirely eliminated. Although avoiding a conventional bioink removes certain material constraints, it does so at the cost of raising new questions concerning the accuracy with which a large number of cell clusters can be positioned, supported, and kept alive as the tissue develops. At this stage, improving resolution is an important step, but it does not solve the problem of constructing a living organ. Chen's five-year ARPA-H project on the liver is one example of an effort to achieve this aim by applying bioprinting on a scale corresponding to a patient's particular organ (Labios). The question that this review cannot currently answer is whether any of these methods can result in a human-scale, highly cellular tissue that maintains fine spatial organization and high cell viability over long periods.

Conclusion

The studies mentioned in this review article indicate that it is possible to improve the resolution of bioprinting, but solving the resolution issue alone does not resolve the density–viability–resolution trilemma. Refractive-index matching provides strong evidence that fine features can be produced at high cell densities by reducing light scattering. Still, this improvement covers only one part of the problem.

In the five constructs that are scored in Table 1, none of them fulfills all three targets, and in four of the five at least one target is not given. This fact itself is a finding about how the field assesses its own progress. As the constructs grow in size, further challenges arise, such as nutrient delivery, the need for mechanical support, and long-term cell survival. Therefore, the remaining challenge is not just about improving one aspect of the trilemma, but about keeping the density, viability, and resolution all together as the size and complexity of the tissue increase. The trilemma will not be solved by a printer which can resolve finer details, but by the first construct to report all three measurements at the same time and to clear all three thresholds.

Even though refractive-index matching is an important step in achieving high density and high resolution, a significant advance will take place only when these improvements can be maintained without damaging cell viability or long-term tissue function. For now, the patients counted at the beginning of this review are still waiting for a donor, not for a printer.

Works Cited

Ayan, Bugra, et al. “Aspiration-Assisted Bioprinting for Precise Positioning of Biologics.” Science Advances, vol. 6, no. 10, 2020, eaaw5111. https://doi.org/10.1126/sciadv.aaw5111.

Bernal, Paulina Nuñez, et al. “Volumetric Bioprinting of Organoids and Optically Tuned Hydrogels to Build Liver-Like Metabolic Biofactories.” Advanced Materials, vol. 34, no. 15, 2022, art. 2110054. https://doi.org/10.1002/adma.202110054.

Billiet, Thomas, et al. “The 3D Printing of Gelatin Methacrylamide Cell-Laden Tissue-Engineered Constructs with High Cell Viability.” Biomaterials, vol. 35, no. 1, 2014, pp. 49–62. https://doi.org/10.1016/j.biomaterials.2013.09.078.

Boothe, Tobias, et al. “A Tunable Refractive Index Matching Medium for Live Imaging Cells, Tissues and Model Organisms.” eLife, vol. 6, 2017, e27240. https://doi.org/10.7554/eLife.27240.

Enderle, John D., and Joseph D. Bronzino. Introduction to Biomedical Engineering. 3rd ed., Academic Press, 2012.

Guan, Jiaao, et al. “Compensating the Cell-Induced Light Scattering Effect in Light-Based Bioprinting Using Deep Learning.” Biofabrication, vol. 14, no. 1, 2022, art. 015011. https://doi.org/10.1088/1758-5090/ac3b92.

Kang, Hyun-Wook, et al. “A 3D Bioprinting System to Produce Human-Scale Tissue Constructs with Structural Integrity.” Nature Biotechnology, vol. 34, no. 3, 2016, pp. 312–319. https://doi.org/10.1038/nbt.3413.

Koffler, Jacob, et al. “Biomimetic 3D-Printed Scaffolds for Spinal Cord Injury Repair.” Nature Medicine, vol. 25, no. 2, 2019, pp. 263–269. https://doi.org/10.1038/s41591-018-0296-z.

Labios, Liezel. “ARPA-H Project Awarded at UC San Diego Aims to End Liver Transplant Shortage with 3D Bioprinting.” UC San Diego Today, 14 Jan. 2026, https://today.ucsd.edu/story/arpa-h-project-awarded-at-uc-san-diego-aims-to-end-liver-transplant-shortage-with-3d-bioprinting. Accessed 21 Sept. 2026.

Lee, Andrew, et al. “3D Bioprinting of Collagen to Rebuild Components of the Human Heart.” Science, vol. 365, no. 6452, 2019, pp. 482–487. https://doi.org/10.1126/science.aav9051.

Lu, Ting-Yu, et al. “Bioprinted High Cell Density Liver Model with Improved Hepatic Metabolic Functions.” Biomaterials, vol. 320, 2025, art. 123256. https://doi.org/10.1016/j.biomaterials.2025.123256.

Ma, Xuanyi, et al. “Deterministically Patterned Biomimetic Human iPSC-Derived Hepatic Model via Rapid 3D Bioprinting.” Proceedings of the National Academy of Sciences, vol. 113, no. 8, 2016, pp. 2206–2211. https://doi.org/10.1073/pnas.1524510113.

Mierke, Claudia Tanja. “Bioprinting of Cells, Organoids and Organs-on-a-Chip Together with Hydrogels Improves Structural and Mechanical Cues.” Cells, vol. 13, no. 19, 2024, art. 1638. https://doi.org/10.3390/cells13191638.

Moroni, Lorenzo, et al. “Biofabrication: A Guide to Technology and Terminology.” Trends in Biotechnology, vol. 36, no. 4, 2018, pp. 384–402. https://doi.org/10.1016/j.tibtech.2017.10.015.

Murphy, Sean V., and Anthony Atala. “3D Bioprinting of Tissues and Organs.” Nature Biotechnology, vol. 32, no. 8, 2014, pp. 773–785. https://doi.org/10.1038/nbt.2958.

Skylar-Scott, Mark A., et al. “Biomanufacturing of Organ-Specific Tissues with High Cellular Density and Embedded Vascular Channels.” Science Advances, vol. 5, no. 9, 2019, eaaw2459. https://doi.org/10.1126/sciadv.aaw2459.

T.C. Sağlık Bakanlığı Konya Şehir Hastanesi. “Organ Bağışlayın, Başka Bedenlere Hayat Armağan Edin.” 4 June 2026, https://konyasehir.saglik.gov.tr/TR-522113/organ-bagislayin-baska-bedenlere-hayat-armagan-edin-.html. Accessed 4 Sep. 2026.

Unagolla, Janitha M., and Ambalangodage C. Jayasuriya. “Hydrogel-Based 3D Bioprinting: A Comprehensive Review on Cell-Laden Hydrogels, Bioink Formulations, and Future Perspectives.” Applied Materials Today, vol. 18, 2020, art. 100479. https://doi.org/10.1016/j.apmt.2019.100479.

United Network for Organ Sharing. “Frequently Asked Questions.” UNOS, https://unos.org/transplant/frequently-asked-questions/. Accessed 4 Sept. 2026.

Xiang, Yi, et al. “Iohexol as a Refractive Index Tuning Agent for Bioinks in High Cell Density Bioprinting.” Biomaterials Science, vol. 13, no. 14, 2025, pp. 3958–3971. https://doi.org/10.1039/D5BM00585J.

You, Shangting, et al. “High Cell Density and High-Resolution 3D Bioprinting for Fabricating Vascularized Tissues.” Science Advances, vol. 9, no. 8, 2023, eade7923. https://doi.org/10.1126/sciadv.ade7923.

You, Shangting, et al. “Mitigating Scattering Effects in Light-Based Three-Dimensional Printing Using Machine Learning.” Journal of Manufacturing Science and Engineering, vol. 142, no. 8, 2020, art. 081002. https://doi.org/10.1115/1.4046986.

Zhang, Zhengyi, et al. “Evaluation of Bioink Printability for Bioprinting Applications.” Applied Physics Reviews, vol. 5, no. 4, 2018, art. 041304. https://doi.org/10.1063/1.5053979.

The Role of TRP Channels in Alzheimer's Disease Pathogenesis and Potential Therapeutic Implications By Seoyoon Chung

Abstract

Alzheimer's Disease (AD) is a progressive neurodegenerative disorder that causes pathological alterations in the brain by impairing its functions and resulting in permanent cognitive impairment. Its main pathological features include amyloid-beta (Aβ) plaques and tau protein tangles, which impair brain cell function and cause neurodegeneration. While these have been the focus of AD research for decades, recent evidence suggests that other cellular pathways may also contribute to AD development. One of the most promising areas involves Transient Receptor Potential (TRP) channels, which are a diverse family of ion channels that regulate the movement of calcium and other ions. TRP channels are found in various tissues, including the brain, where they help detect and respond to stimuli such as oxidative stress and inflammation. These functions are important, especially in AD, as calcium imbalance, oxidative damage, and chronic inflammation are known to worsen neuronal damage. This review examines the current understanding of TRP channels' role in AD pathogenesis. TRPM2 has been implicated in AD pathogenesis through oxidative stress-induced calcium dysregulation and neuronal damage. TRPV1 has been associated with both neuronal injury and protective microglial responses, while TRPC1 is involved in calcium signaling and neuronal survival, although its direct role in AD remains less established. Modulating these channels has shown potential in preclinical studies to reduce pathological processes associated with AD, but further development of selective TRP channel modulators and additional investigations are needed to fully assess their therapeutic feasibility. Further research is therefore needed to clarify the contribution of TRP channel dysfunction to AD pathogenesis and evaluate its therapeutic potential.

Introduction

Alzheimer's Disease (AD) is a progressive, irreversible brain disorder that leads to memory loss, confusion, and eventual behavioral changes. As the most common cause of dementia, AD mainly affects older adults and currently has no cure. Its core pathological hallmarks include Aβ plaques and tau protein tangles, both of which impair brain cell function and cause neurodegeneration (Selkoe). While these features have been the focus of AD research for decades, growing evidence suggests that other cellular pathways may also play important roles in the disease's development (Yamamoto et al.).

One of the most promising areas of recent research involves Transient Receptor Potential (TRP) channels, a diverse family of ion channels that regulate the movement of calcium and other ions across the cell membrane (Yang and Kim). TRP channels are found in many tissues, including the brain, where they help detect and respond to physical and chemical stimuli such as temperature, oxidative stress, and inflammation.

Specific TRP subfamilies have been linked to pathological processes relevant to AD. For example, TRPM2 is activated by oxidative stress and contributes to calcium overload, leading to

neuronal death in AD models (Malko and Jiang). Likewise, TRPV1 activity has been associated with both neuronal damage and protective microglial responses, while TRPC1 is involved in calcium signaling and neuronal survival, although its direct role in AD remains less established (Hong et al.; Lu et al.).

This review explores the current understanding of TRP channel roles in AD development and progression and discusses the therapeutic implications of targeting TRP channels as an emerging approach to the disease.

General Overview of TRP Channels

TRP channels are a group of ion channels located mostly on the plasma membrane of numerous animal cells, frequently neurons (Figure 1). The term "Transient Receptor Potential" originates from their initial discovery in Drosophila melanogaster, where mutations in these channels led to abnormal light responses (Venkatachalam and Montell). TRP channels play a key role in the signal transduction of cells as a medium for cations, commonly Ca^{2+}, Mg^{2+}, Na^{+}, and K^{+}, to transport in and out of the cells (Minke and Cook). Among these ions, calcium plays a particularly important role because it functions as a secondary messenger in numerous intracellular signaling pathways (Venkatachalam and Montell). When activated by external stimuli, TRP channels become permeable to ions, allowing them to enter the neuron's cytoplasm. This ion influx can depolarize the membrane, initiating downstream cellular responses such as neurotransmitter release or gene expression (Venkatachalam and Montell). This mechanism varies with the trigger, since TRP channel activation depends on the type of external stimulation they receive (Numata et al.). The factors can be categorized into five: cellular signaling pathway, cellular chemical environment, type of external stimulation, direct action of chemicals, and interaction of proteins (Numata et al.). This variety in activation mechanisms allows TRP channels to function as multimodal sensors within the nervous system (Venkatachalam and Montell).

In particular, stimulation of G protein-coupled receptors (GPCRs) or receptor tyrosine kinases (RTKs) activates phospholipase C (PLC), which hydrolyzes phosphatidylinositol 4,5-bisphosphate (PIP_2) into inositol 1,4,5-trisphosphate (IP_3) and diacylglycerol (DAG) (Venkatachalam and Montell). IP_3 induces calcium ion (Ca^{2+}) release from the endoplasmic reticulum (ER), indirectly modulating TRP channel activity, whereas DAG can directly activate certain TRP channels (Venkatachalam and Montell). Additionally, receptor-induced recruitment of TRP channels into signalosomes can relocate channels to the plasma membrane or alter their functional properties (Numata et al.). Collectively, these signaling cascades act as molecular relays, translating extracellular stimuli into TRP channel activation. Such pathways highlight how TRP channels bridge surface receptor activation with intracellular calcium signaling (Venkatachalam and Montell).

In addition, TRP channels are highly sensitive to the cellular chemical environment, including changes in redox status and pH (Venkatachalam and Montell). Oxidation or reduction of TRP cysteine residues can modulate channel gating, whereas acidic or alkaline shifts in intra-

or extracellular pH levels can alter channel responsiveness (Numata et al.). These chemical parameters fine-tune TRP activity, allowing the channels to integrate intracellular and extracellular cues into appropriate cellular responses. This sensitivity is especially relevant in pathological conditions, such as inflammation or oxidative stress, where cellular chemistry is altered (Venkatachalam and Montell).

TRP channels act as versatile detectors of environmental changes, integrating thermal, electrical, and mechanical information. Meanwhile, certain channels, such as TRPV1, function as heat sensors, whereas TRPM8 channels respond to cold stimuli. TRPV1 is also activated by chemical ligands such as capsaicin, linking thermal and chemical sensation to pain perception (Venkatachalam and Montell). In addition, changes in plasma membrane voltage can influence channel gating, and mechanical forces such as pressure, stretch, or vibration can modulate the activity of certain TRP channels through mechanostimulation. This multimodal capability makes TRP channels essential for sensory neurons that must rapidly adapt to oscillating environments (Venkatachalam and Montell).

These channels can also be directly activated by chemical ligands through noncovalent interactions. For example, capsaicin, the pungent compound in chili peppers, directly activates TRPV1, producing a sensation of heat, whereas menthol activates TRPM8, inducing a cooling sensation. This mechanism demonstrates that TRP channels can transduce specific chemical signals into electrical or biochemical cellular responses without requiring intermediate signaling cascades. Such direct activation enables rapid sensory perception, critical for survival (Venkatachalam and Montell). Finally, protein–protein interactions modulate TRP channel function by forming signaling complexes with other cellular proteins. These interactions might regulate channel localization, gating properties, and integration toward broader signaling networks. Through such cooperative mechanisms, TRP channels coordinate with partner proteins to fine-tune cellular responsiveness to multiple stimuli (Numata et al.).

In sensory neurons, activation of certain TRP channels can depolarize the cell and contribute to the generation of an action potential. The impulse travels to the brain or spinal cord via an axon. It enables the perception of sensations such as heat, pain, pressure, and chemicals. Overall, TRP channels act as critical molecular interfaces between environmental stimuli and neuronal signaling, underscoring their importance in both normal sensory processing and neurological disease (Venkatachalam and Montell).

General Overview of AD

AD is the most common form of dementia, which is a brain disorder that profoundly affects one's performance in memory, thinking, and social abilities (Figure 2). First described by Dr. Alois Alzheimer in 1906, it was reported as a strange, severe disease of the cerebral cortex (Hippius and Neundörfer). This initial observation provided the groundwork for modern Alzheimer's research, which has since identified specific molecular and cellular abnormalities associated with the disease. As research evolved, amyloid-beta (Aβ) plaques and tau tangles

were identified as central pathological hallmarks of AD, commonly used to support neuropathological diagnosis (Querfurth and LaFerla).

Aβ refers to the type of peptides made up of 36 to 43 amino acids that constitute the amyloid plaques in the brains of AD patients. Among these, Aβ42 is considered particularly toxic because it readily aggregates and forms insoluble plaques (Selkoe and Hardy). These peptides derive from the Aβ precursor protein (APP), when the protein is cleaved by two proteins: beta-secretase (β-secretase) and gamma-secretase (γ-secretase). This amyloidogenic processing pathway contrasts with the non-amyloidogenic pathway, which does not produce toxic Aβ fragments and is more common in normal brains. In this process, β-secretase primarily cleaves the protein into smaller segments, which γ-secretase then further cleaves. The resulting Aβ peptides aggregate extracellularly, forming plaques that disrupt neuronal communication, activate immune responses, and contribute to synaptic dysfunction (Querfurth and LaFerla).

Neurofibrillary tangles are formed when tau accumulates aberrantly inside neurons (Grundke-Iqbal et al.). Normally, tau proteins bind to microtubules in neurons and help them organize. Microtubules play an important role in maintaining neuronal structure and transporting essential molecules along the axon (Wang and Mandelkow). However, in AD, abnormal chemical reactions cause tau proteins to detach from microtubules and aggregate with other tau molecules (Wang and Mandelkow). These eventually form threads that join together to form tangles inside neurons, blocking their transport system and disrupting synaptic communication. As a result, affected neurons experience impaired signaling and are more vulnerable to degeneration and cell death (Iqbal et al.). This accumulation occurs in specific cerebral regions responsible for memory, such as the entorhinal cortex and hippocampus, which are among the first areas affected in AD. The progressive spread of tau pathology from these regions to other areas of the brain strongly correlates with the severity of cognitive decline observed in patients with AD (Braak and Braak).

In addition to these, other biological processes, including neuroinflammation and oxidative stress, have also proven to play a considerable role in AD progression. Chronic activation of microglia and astrocytes releases inflammatory cytokines and reactive oxygen species, further damaging neurons and accelerating disease progression (Heneka et al.). Despite extensive research into these mechanisms, there is currently no cure for AD. This highlights the need for continued research into alternative molecular pathways and therapeutic targets to better address the underlying causes of AD.

The Role of TRP Channels in AD

TRP channels have recently emerged as important contributors to AD pathophysiology, particularly through their roles in calcium signaling, neuroinflammation, and oxidative stress (Yamamoto et al.). Increasing evidence suggests that dysregulation of specific TRP channel subtypes may worsen neurodegenerative processes in the central nervous system (Hong et al.).

TRPM2

Within the TRP family, TRPM2 has been recognized as a possible molecular mechanism that contributes to oxidative damage and neuronal injury. Oxidative stress can trigger TRPM2 channel activation, causing abnormal accumulation of Ca^{2+} ions in cells, which may eventually lead to cell death, including in neurons. This mechanism has been linked to various neurological conditions, including AD, stroke, and Parkinson's disease (Belrose and Jackson).

In AD, oxidative stress and Aβ-associated cellular damage may contribute to TRPM2 activation and subsequent disruption of calcium homeostasis. Malko and Jiang describe TRPM2-mediated cell death as an important mechanism linking oxidative stress to pathological conditions. These findings suggest that TRPM2 may contribute to AD progression by increasing neuronal vulnerability to oxidative damage. However, further research is needed to determine the extent of its contribution to AD in humans.

TRPV1

Within the TRP family, TRPV1 has received particular attention due to its role in inflammatory signaling and neuronal damage. TRPV1 is a non-selective cation channel activated by capsaicin, vanilloids, acidic conditions, and high temperatures, and it is expressed in various cell types in the brain, including neurons and microglia. Activation of TRPV1 increases intracellular calcium levels, which can influence mitochondrial function, oxidative stress, and inflammatory responses, all of which are strongly implicated in AD pathology (Venkatachalam and Montell; Hong et al.).

Experimental studies suggest that TRPV1 activation can contribute to neuronal damage under certain conditions. Kim et al., in their study of mesencephalic dopaminergic neurons, exposed rat neurons to the TRPV1 agonist capsaicin and observed significant neuronal cell death. This effect was inhibited by the TRPV1 antagonists capsazepine and iodo-resiniferatoxin, indicating that TRPV1 activation mediated the observed neurotoxicity (Kim et al., "Transient Receptor Potential Vanilloid Subtype 1 Mediates Cell Death"). The study also showed that capsaicin-induced TRPV1 activation elevated intracellular Ca^{2+} levels and caused mitochondrial damage, accompanied by increased cytochrome c release and cleaved caspase-3 immunoreactivity, both associated with apoptosis (Kim et al., "Transient Receptor Potential Vanilloid Subtype 1 Mediates Cell Death"). These results indicate that excessive TRPV1 activation promotes neuronal injury through calcium-dependent mitochondrial dysfunction. Although this experiment did not directly investigate AD, it demonstrates a mechanism through which TRPV1 dysregulation may contribute to neuronal damage associated with neurodegenerative diseases.

Similar mechanisms were observed in a subsequent study by Kim et al., who found that TRPV1 activation increased intracellular calcium levels and induced mitochondrial damage and cell death in rat microglia and a human microglial cell line (Kim et al., "Transient Receptor Potential Vanilloid Subtype 1 Mediates Microglial Cell Death"). These effects were reduced by TRPV1 antagonists, further supporting the involvement of TRPV1 in calcium-mediated cellular injury (Kim et al., "Transient Receptor Potential Vanilloid Subtype 1 Mediates Microglial Cell

Death"). Given microglia's central role in amyloid-beta–induced neuroinflammation, these findings suggest a possible connection between TRPV1 activity and cellular processes relevant to AD, although the experiment did not directly investigate AD pathology.

Beyond these cellular studies, broader evidence supports TRPV1 involvement in AD pathology. Lu et al. found that pharmacological TRPV1 activation improved metabolic impairments in Aβ-tolerant microglia and restored their phagocytic and autophagic functions. In contrast, microglia-specific TRPV1 deletion accelerated amyloid pathology and memory impairment in APP/PS1 mice. Capsaicin and vanillin, both TRPV1 agonists, have been shown to attenuate memory impairment (Jiang et al.; Jayant et al.), while vanillin has also been shown to reduce oxidative and nitrosative stress, as well as acetylcholinesterase activity in experimental AD models (Jayant et al.). Furthermore, modulation of TRPV1 activity has been associated with changes in tau phosphorylation (Jiang et al.) and pro-inflammatory cytokine expression, which suggests that TRPV1 may influence multiple downstream pathways involved in AD progression.

TRPC1

TRPC1 is another member of the TRP family involved in calcium signaling and neuronal function. As calcium homeostasis is essential for maintaining normal neuronal activity, disruption of calcium regulation may contribute to neuronal damage associated with neurodegenerative diseases (Hong et al.).

Hong et al. discuss the role of TRPC1 in regulating calcium entry and neuronal survival, particularly in experimental models of Parkinson's disease and Huntington's disease. These findings suggest that TRPC1 may influence cellular processes relevant to neurodegeneration. More specifically, Li et al. investigated the role of TRPC1 in Aβ-induced neuronal damage and found that TRPC1 deletion aggravated Aβ-induced learning and memory deficits and increased cell death in the mouse hippocampus. Conversely, increased TRPC1 expression reduced Aβ production and apoptosis in cells overexpressing amyloid precursor protein (APP). These findings suggest that TRPC1 may have a protective role against Aβ-associated neuronal damage.

However, other studies suggest that TRPC1 may also contribute to AD pathology under different conditions. Ye et al. found that increased intracellular tau levels upregulated TRPC1 expression, enhancing store-operated calcium entry (SOCE), intracellular calcium accumulation, and endoplasmic reticulum stress. Furthermore, TRPC1 overexpression aggravated tau hyperphosphorylation and cognitive impairment in experimental models, whereas TRPC1 inhibition or deletion attenuated tau pathology and improved synaptic and cognitive functions. These findings suggest that TRPC1 dysregulation may contribute to AD pathology under certain experimental conditions.

Together, these findings indicate that TRPC1 may have both protective and damaging effects depending on the pathological conditions. However, its precise role in human AD pathogenesis remains less established. Therefore, further human-based studies are needed to determine how TRPC1 dysfunction contributes to Aβ- and tau-associated pathology in AD.

Therapeutic Implications and Limitations

Melastatin-related TRP channels (TRPM channels) may be therapeutic targets for degenerative diseases, including AD (Belrose and Jackson). Hong et al. emphasize the potential of targeting TRP channels in neurodegenerative diseases, particularly through the development of selective channel modulators. Yet, therapeutic approaches targeting Aβ aggregation alone have not fully addressed the progression of AD, highlighting the importance of investigating alternative molecular mechanisms (Selkoe and Hardy).

TRPM2 represents a potential therapeutic target due to its involvement in oxidative stress-induced calcium dysregulation and neuronal damage. Inhibiting excessive TRPM2 activity may help reduce calcium overload and protect neurons from oxidative injury (Belrose and Jackson; Malko and Jiang). However, further investigations are needed to determine whether this approach can produce effective therapeutic outcomes in human AD.

TRPV1 also presents therapeutic possibilities, although its effects appear to depend on the experimental conditions. While TRPV1 antagonists have been shown to reduce neuronal and microglial damage in experimental studies (Kim et al., "Transient Receptor Potential Vanilloid Subtype 1 Mediates Cell Death"; Kim et al., "Transient Receptor Potential Vanilloid Subtype 1 Mediates Microglial Cell Death"), Lu et al. demonstrated that TRPV1 activation could improve microglial function and reduce amyloid pathology in AD mouse models. These contrasting findings suggest that selective regulation of TRPV1 activity may be necessary to achieve therapeutic benefits.

TRPC1 may also be relevant to therapeutic research due to its involvement in calcium signaling and neuronal survival (Hong et al.). Although experimental studies suggest that TRPC1 influences Aβ- and tau-associated pathology, its effects vary across disease models. Further research is needed to determine whether targeting TRPC1 could provide therapeutic benefits in AD.

Recent findings suggest that modifying TRP channels may help reduce certain pathological processes associated with AD. Among its members, TRPA1, TRPC6, TRPM2, and TRPV1 have shown potential therapeutic relevance in preclinical research, although the evidence varies between channel subtypes (Hong et al.). Nevertheless, most findings remain limited to cellular and animal models, and the widespread physiological functions of TRP channels make selective targeting challenging. Further research is needed to develop selective TRP channel modulators and determine their safety and effectiveness in human AD.

Conclusion

TRP channels may contribute to AD pathogenesis, providing further insight into the mechanisms and potential solutions for AD (Hong et al.). This review found that TRPM2 may contribute to oxidative stress-induced calcium dysregulation and neuronal damage, while TRPV1 can have both damaging and protective effects depending on the cellular conditions. TRPC1 is involved in calcium signaling and neuronal survival, although its direct role in AD remains less established. This connection may play a critical role in future research on treatments for

neurodegenerative diseases, further expanding the areas of focus for therapeutic interventions for AD. By uncovering the role of TRP channels in AD, these findings may spur more thorough research into ion channel-based treatments, supplementing current progress in treating neurodegenerative diseases, especially those associated with amyloid plaques and tau tangles.

Although the relationship between TRP channels and AD has been investigated, further research is needed to establish the extent to which TRP channel dysfunction contributes to AD development and progression. Clarification is needed on the functional pathways by which specific TRP channel subfamilies influence AD pathology, as well as on translating preclinical findings into clinical investigations (Hong et al.). Longitudinal and human-based studies are particularly needed to determine the therapeutic viability of modulating TRP channels. By addressing these gaps, research can more decisively evaluate whether TRP-targeted interventions would represent an achievable and effective strategy for treating AD.

Works Cited

Belrose, Jillian Corinne, and Michael Frederick Jackson. “TRPM2: A Candidate Therapeutic Target for Treating Neurological Diseases.” Acta Pharmacologica Sinica, vol. 39, no. 5, 2018, pp. 722–32, https://doi.org/10.1038/aps.2018.31.

Braak, H., and E. Braak. “Neuropathological Stageing of Alzheimer-Related Changes.” Acta Neuropathologica, vol. 82, no. 4, 1991, pp. 239–59, https://doi.org/10.1007/BF00308809.

Grundke-Iqbal, I., et al. “Abnormal Phosphorylation of the Microtubule-Associated Protein Tau (Tau) in Alzheimer Cytoskeletal Pathology.” Proceedings of the National Academy of Sciences, vol. 83, no. 13, 1986, pp. 4913–17, https://doi.org/10.1073/pnas.83.13.4913.

Heneka, Michael T., et al. “Neuroinflammation in Alzheimer’s Disease.” The Lancet Neurology, vol. 14, no. 4, 2015, pp. 388–405, https://doi.org/10.1016/S1474-4422(15)70016-5.

Hippius, Hanns, and Gabriele Neundörfer. “The Discovery of Alzheimer’s Disease.” Dialogues in Clinical Neuroscience, vol. 5, no. 1, 2003, pp. 101–08, https://doi.org/10.31887/DCNS.2003.5.1/hhippius.

Hong, Chansik, et al. “TRP Channels as Emerging Therapeutic Targets for Neurodegenerative Diseases.” Frontiers in Physiology, vol. 11, 2020, article 238, https://doi.org/10.3389/fphys.2020.00238.

Iqbal, Khalid, et al. “Tau Pathology in Alzheimer Disease and Other Tauopathies.” Biochimica et Biophysica Acta (BBA) – Molecular Basis of Disease, vol. 1739, nos. 2–3, 2005, pp. 198–210, https://doi.org/10.1016/j.bbadis.2004.09.008.

Jayant, Shalini, et al. “Protective Effect of Transient Receptor Potential Vanilloid Subtype 1 (TRPV1) Modulator, against Behavioral, Biochemical and Structural Damage in Experimental Models of Alzheimer's Disease.” Brain Research, vol. 1642, 2016, pp. 397–408, https://doi.org/10.1016/j.brainres.2016.04.022.

Jiang, Xia, et al. “Capsaicin Ameliorates Stress-Induced Alzheimer's Disease-Like Pathological and Cognitive Impairments in Rats.” Journal of Alzheimer's Disease, vol. 35, no. 1, 2013, pp. 91–105, https://doi.org/10.3233/JAD-121837.

Kim, Sang R., et al. “Transient Receptor Potential Vanilloid Subtype 1 Mediates Cell Death of Mesencephalic Dopaminergic Neurons In Vivo and In Vitro.” The Journal of Neuroscience, vol. 25, no. 3, 2005, pp. 662–71, https://doi.org/10.1523/JNEUROSCI.4166-04.2005.

Kim, Sang R., et al. “Transient Receptor Potential Vanilloid Subtype 1 Mediates Microglial Cell Death In Vivo and In Vitro via Ca^{2+}-Mediated Mitochondrial Damage and Cytochrome c Release.” The Journal of Immunology, vol. 177, no. 7, 2006, pp. 4322–29, https://doi.org/10.4049/jimmunol.177.7.4322.

Li, Mengzhu, et al. “TRPC1 Null Exacerbates Memory Deficit and Apoptosis Induced by Amyloid-β.” Journal of Alzheimer's Disease, vol. 63, no. 2, 2018, pp. 761–72, https://doi.org/10.3233/JAD-180077.

Lu, Jia, et al. “TRPV1 Sustains Microglial Metabolic Reprogramming in Alzheimer’s Disease.” EMBO Reports, vol. 22, no. 6, 2021, article e52013, https://doi.org/10.15252/embr.202052013.

Malko, Philippa, and Lin-Hua Jiang. “TRPM2 Channel-Mediated Cell Death: An Important Mechanism Linking Oxidative Stress-Inducing Pathological Factors to Associated Pathological Conditions.” Redox Biology, vol. 37, 2020, article 101755, https://doi.org/10.1016/j.redox.2020.101755.

Minke, Baruch, and Boaz Cook. “TRP Channel Proteins and Signal Transduction.” Physiological Reviews, vol. 82, no. 2, 2002, pp. 429–72, https://doi.org/10.1152/physrev.00001.2002.

Numata, Tomohiro, et al. “Activation of TRP Channels in Mammalian Systems.” TRP Channels, edited by Michael X. Zhu, CRC Press/Taylor & Francis, 2011, ch. 3, https://www.ncbi.nlm.nih.gov/books/NBK92819/.

Querfurth, Henry W., and Frank M. LaFerla. “Alzheimer’s Disease.” The New England Journal of Medicine, vol. 362, no. 4, 2010, pp. 329–44, https://doi.org/10.1056/NEJMra0909142.

Selkoe, Dennis J. “Alzheimer’s Disease: Genes, Proteins, and Therapy.” Physiological Reviews, vol. 81, no. 2, 2001, pp. 741–66, https://doi.org/10.1152/physrev.2001.81.2.741.

Selkoe, Dennis J., and John Hardy. “The Amyloid Hypothesis of Alzheimer’s Disease at 25 Years.” EMBO Molecular Medicine, vol. 8, no. 6, 2016, pp. 595–608, https://doi.org/10.15252/emmm.201606210.

Venkatachalam, Kartik, and Craig Montell. “TRP Channels.” Annual Review of Biochemistry, vol. 76, 2007, pp. 387–417, https://doi.org/10.1146/annurev.biochem.75.103004.142819.

Wang, Yipeng, and Eckhard Mandelkow. “Tau in Physiology and Pathology.” Nature Reviews Neuroscience, vol. 17, no. 1, 2016, pp. 22–35, https://doi.org/10.1038/nrn.2015.1.

Yamamoto, Shinichiro, et al. “Transient Receptor Potential Channels in Alzheimer’s Disease.” Biochimica et Biophysica Acta (BBA) – Molecular Basis of Disease, vol. 1772, no. 8, 2007, pp. 958–67, https://doi.org/10.1016/j.bbadis.2007.03.006.

Yang, Dongki, and Jaehong Kim. “Emerging Role of Transient Receptor Potential (TRP) Channels in Cancer Progression.” BMB Reports, vol. 53, no. 3, 2020, pp. 125–32, https://doi.org/10.5483/BMBRep.2020.53.3.016.

Ye, Jinwang, et al. “Tau-Induced Upregulation of C/EBPβ-TRPC1-SOCE Signaling Aggravates Tauopathies: A Vicious Cycle in Alzheimer Neurodegeneration.” Aging Cell, vol. 19, no. 9, 2020, article e13209, https://doi.org/10.1111/acel.13209.

Formation Flight Applied to Commercial Aviation: A Review of the Aerodynamic, Operational, and Safety Constraints By Giacomo Guenzi

Abstract

In commercial aviation, aircraft flying in formation consume less fuel compared to the same aircraft flying solo because the trailing aircraft can fly in the wake of the leading aircraft reducing induced drag. Many aircraft travel trans-Atlantic flights due to popular destinations. Due to this, it has been thought that the aircraft can fly together to complete the route in formation which reduces fuel consumption for the trailing aircraft. This paper will investigate the several factors associated with affecting the efficiency of fuel consumption from the aerodynamics, the operational aspects, and the current gaps in research that exist. The extent to which fuel consumption can be reduced during formation flight depends on the shape of the wing because the wake flown on the trailing aircraft is created by the initial wing. Many other factors affect efficiency as well, such as distance between aircraft, flight routes, etc. Applying this to commercial aviation, most commercial jets have winglets which reduce the wake turbulence hence the results provided in the paper for efficiency may be optimistic compared to reality.

Introduction

An airline's greatest tool is also its greatest weight: fuel. Lately the price of fuel makes up nearly 30% of an airline's total spending [IATA]. Some resistance such as induced drag that an aircraft encounters in flight increases fuel burn: the air leaves the wing tips creating vortices which apply resistance to the aircraft's forward motion. Since there is resistance, the engines need to burn more fuel to generate more thrust to maintain the same speed and altitude. An aircraft's design can fix this problem only to a certain extent, but another proposed solution considers having aircraft fly in a specific formation. Indeed, this reduces the costs of flying not only for the airline but also for the passenger.

Around 72% of an airline's income is covered by passenger revenue [IATA]. Hence, a large problem emerges, where inefficiency isn't only a problem for airlines but also for the consumer.

This extra cost on the consumer is unfortunate not only for the consumer who must pay more but also for the airline who might receive less customers. A proposed solution is the formation flight. This idea of formation flight includes a trailing aircraft flying in the upward moving air created by the vortexes at a wingtip. The upward air provides extra lift to the trailing aircraft which in turn decreases the power of its engines to maintain the same speed and altitude. This is like how birds fly in V formation [Portugal]. In decreasing the power of the engines, less fuel is burnt per unit time which in the end decreases total costs for the airline.

However, whether flying formation can be applied to commercial aviation remains unsolved due to regulations forcing aircraft to fly at large distances outside of the leading wake, the difficulty with route and schedule optimizations, and a general contrary public opinion.

Furthermore, wingtips, which are very common on commercial aircraft, reduce the vortex that the trailing aircraft need to fly in forbidding formation flight [Pinto] & [Sarkar].

The extent to which formation flight can reduce inefficiencies during flight depends on factors such as the aircraft's wing design since a wing with a winglet reduces the vortices that the trailing aircraft can fly on. A prime example of this is how winglets affect the vortexes created by an aircraft. The calculation required to predict a vortex includes Betz's method. While Betz's method is very helpful at predicting the vortex of a wing with an elliptical lift distribution, its limitation resides in wings without an elliptical lift distribution. Hence, Ning points out that extended methods of the Betz method are helpful to predict vortexes from a wingtip. He investigates that even with the extended methods it still isn't clear if vortexes created will merge or disperse or stay together, and what this suggests about the efficiency is that it is harder to fully understand the benefit of wing tips for vortexes [Ning].

This paper starts off by investigating the physical mechanism behind the vortex reducing fuel consumption before explaining how the shape of a wing, both standard and non-standard, can affect the vortex by changing its strength or size. Following this, the paper will look at the safe distance between aircraft flying in formation and then how aircraft are paired up with specific routes and schedules to reduce fuel consumption for transatlantic flights. Since all of this is theoretical, these predictions will be tested against real life flight data. After that the paper turns to the safety requirements affecting formation flight, which builds up to a conflict between wingtip devices designed for solo flight efficiency and the strength of vortex a trailing aircraft needs to benefit from formation. Finally, the paper shows how all the values from the different sources are consistent with the ideas they represent but also outlines the gaps in research for flying in formation with aircraft that have winglets.

Section 1: The Physical Mechanism of Vortex Generation and Its Effect on Fuel Consumption

Vortex generation, created by the changes in lift generated across different sections of a wing, is the process allowing the fuel savings in formation flight. Understanding the physical mechanism behind the vortex generation with mathematics, ranging from Betz's standard model to the extensions for non-standard wings, is central to calculating how much fuel saving can be completed by modern aircraft.

In formation flight, the trailing aircraft can fly on the wake of the leading aircraft created by a vortex which reduces the drag for the trailing aircraft allowing for fuel saving [Durango]. The magnitude and shape of the vortex depends on the shape of the wing because of how lift is distributed across the wing's span [Ning]. This section will look at the mathematics used to calculate the lift distribution and how the lift distribution creates different vortexes with different types of wings (standard and non-standard wings) using several examples throughout. The reduction in costs of fuel will also be considered while going through the mathematics.

1a) Standard wings

When observing the amount of lift created by a wing across its span it is observed that the wing doesn't produce uniform lift. The lift produced by a wing is largest near the root of the wing and it decreases as we move closer to the tip of the wing. This uneven distribution of lift is what creates the vortex that swirls from the wingtip [Ning]. The full structure of this vortex can be predicted with the standard mathematical approach involving Betz's method which takes the wing's lift distribution as an input and predicts the final shape of the vortex as an output [Ning]. How the Betz method works is that it tracks four given quantities of a vortex as it rolls up into the larger wake at the tip of the wing. The four quantities are total circulation, the position of the centre point of the vortex, and how spread the vortex is around that centre point. Below are the equations for each quantity which are treated as constants.

$$\Gamma = \int \gamma dA$$

$$\Gamma_y = \int y\gamma dA$$

$$\Gamma_z = \int z\gamma dA$$

$$\Gamma_r = \int (y^2 + z^2)\gamma dA$$

Taking the time derivative of the four quantities above gives four equations which are all constant apart from the last one which can be seen below.

$$\frac{d\Gamma_r}{dt} = -\int_{-\infty}^{\infty} (z - \overline{z}) \frac{w^2}{2} |_{y=0} \, dz$$

The rate of change of the spread of the vortex depends on the vertical velocity which is part of the prediction given by the method so it can't be solved for directly, and if the rate of change fluctuated as the vortex rolled up it would be impossible to predict the final shape. Hence Betz had to come up with an assumption to treat the spread as a constant so the vortex could be calculated. Betz found that if the vortex distribution was vertically symmetric the rate would be zero. Since the vortex is vertically symmetric at the start when it is flat and at the end when it is rolled up, Betz assumed that the quantity of the spread remained roughly constant throughout the whole process. The reason that this assumption is only used with elliptical distributions is that the vortex unfolds smoothly across the wing with a smooth lift distribution without any sudden disruptions, and hence it can be assumed that there are no major changes in spread of the vortex [Ning].

However, Betz's method is limited only to lift distributions that are elliptic or close to elliptic because otherwise the method's predictions become incorrect [Ning]. We can categorize wings with an elliptic curve as standard and wings with different curves as non-standard. The limitation of the method is what separates these two wings which will be investigated in the next

section. Since a large part of this review is optimizing how different wings produce larger vortices, the mechanism of the vortex generation must be applied to different wings. This requires taking a single result out of Betz's method which is the vortex radius as the suitable comparison metric between the vortexes of different wings to identify which one is larger. A common misconception is that the vortex radius can be calculated as a fixed percentage of a wing's span but Ning shows this is wrong because the lift distribution of a wing can change while keeping the same span and it is changes in the lift distribution which create the vortexes not the changes in span. This distinction is important because when comparing wings the vortex radius obtained from the specific lift distribution must be used instead of the span. The result of elliptical wings out of Betz's method can only be used with elliptical wings and in the next section the extended methods used for non-standard wings will be walked through.

1b) Non-standard wings

As previously discussed, there is a distinction between standard and non-standard wings due to their different lift distribution shapes ranging from elliptical to non-elliptical. Non-elliptical shaped wings fall outside the predictions and calculations of the Betz method which is why several extensions of the method have been developed for non-standard wings. An example of this is Donaldson and Rossow extending these methods to lift distributions generated from a flaps-down configuration. These extensions are highly important to this review paper because Ning claims that these extended methods may be necessary to study wings with winglets [Ning]. The reason Ning claims this is because a winglet completely changes the lift distribution from standard to non-standard away from the elliptical case. Hence, rather than using the Betz method these methods seem most suitable for studying winglet equipped wings that many commercial airliners use today. These methods however can't provide precise results for whether vortices would remain separated or merge. In practice winglets are able to change the vortex, as shown in a computational study of the Airbus A380 wing. Blended winglets at four different cant angles of 0, 15, 45 and 80 degrees provided the smallest vortex as the cant angle increased, with the smallest winglet vortex corresponding to the 80 degree cant angle [Pinto].

An example showing that this research isn't purely theoretical is the strut-braced wing being investigated by Durango which is a non-standard wing [Durango]. Now that the paper has gone over the mathematical approaches to model vortexes with both standard and non-standard wings, the next step is to understand how the vortex can reduce drag and lower fuel consumption for the trailing
aircraft in formation.

1c) The effect of the vortex on the cost of fuel

The leading aircraft is cruising and is creating vortexes from its wingtips. Once these vortexes have formed a trailing aircraft can fly into the air moving upwards created by the vortexes and benefit from it with a reduction in drag. The precise mechanism is that the vortexes manipulate the air in two directions: upwards and downwards, where the upwards movement of

air is considered the upwash and the downwards movement the downwash [Durango].

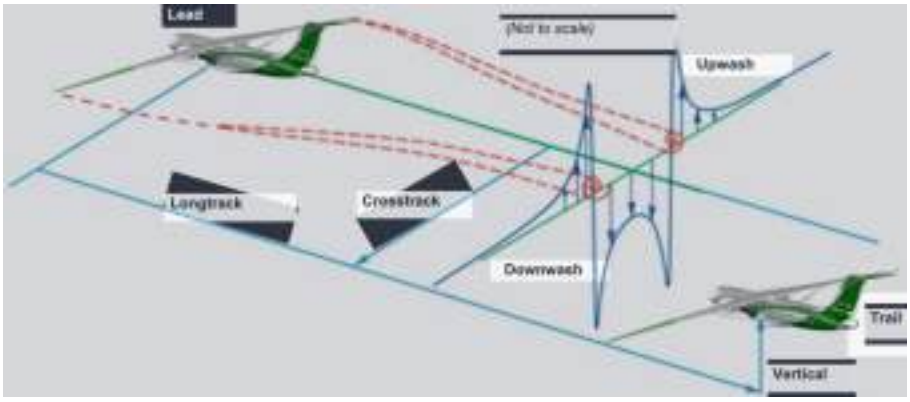

Figure 1. Vertical induced velocity field due to aircraft wake. Adapted from [Durango, Lawson, Shahneh].

If the aircraft flies in the upcoming upwash, the incoming air onto the wing will have changed direction and since the lift vector of an aircraft is perpendicular to the incoming flow of air, a shift from a horizontal flow to a diagonal flow will lead to the lift vector rotating from being perfectly vertical to tilting to the left. The lift vector tilting to the left cancels with a part of the drag vector reducing overall drag on the aircraft.

With this reduction in drag the aircraft can afford to use less fuel while staying at the same altitude and speed [Durango]. The reduction in fuel consumption shows up in real flight testing with the investigation of two C-17s where the maximum average fuel reduction on both wings varied close to 7% [Pahle]. The limitation here is that there is no instrument to directly measure drag and it can only be inferred and calculated using changes in thrust and fuel consumption and here it is safe to say that there has been a decrease in drag.

Section 2: Safe distance for formation flying and vortex generation

Flying in formation to save the most fuel requires the aircraft to fly with a specific distance between them [Bower], as well as with a good flight geometry and efficient flight paths [Kent].

This section will cover the best distances between aircraft so that the benefit is fully received during flight, which formation shapes are the best and why, and how all these factors affect real flight across real trans-Atlantic routes because of the focus on commercial flight.

2a) Formation Design

Formation design decides how much of an aircraft's wing can sit in an upwash, and it can be changed to try to increase fuel savings. Below is a table relating formation, aircraft, spacing, and fuel savings.

Formation	Aircraft	Spacing	Fuel savings
Echelon	2	10%	4%
Echelon	2	Overlapping	>10%
Inverted-V	3	10%	7%
Inverted-V	3	Overlapping	16%

Figure 2: Effect on fuel savings while changing formation, aircraft and spacing. Data adapted from [Bower, Flanzer, and Kroo]

Flying in two aircraft echelon formation leads with wingtip spacing being 10% of the wingspan leads to fuel savings of 4% [2]. As that spacing decreases to an overlap of the two wingtips, the overall fuel savings increase past 10%. If we change the formation to an inverted V-formation and add an aircraft so we have three aircraft, the fuel savings range from 7% to 16% while keeping the same wingtip spacing and overlapping. These results were consistent with other real flight tests such as Airbus's fellowship programme which reported 5% reductions in fuel and emissions for a long-haul flight [Airbus fello'fly]. Hence, the number of aircraft flying in is just as important as the formation the aircraft fly in. There is a trade-off with this formation however [Ning] : when flying in inverted V formation, most of the benefit received from fuel-saving is directed towards the trailing aircraft which makes the fuel-saving more sensitive to small changes such as errors in GPS or winds which can easily set the trailing aircraft off centre and since that aircraft is the only aircraft benefiting from the formation flight, the whole point of the formation flight breaks down. What all this shows is that formation design can only be used to a certain extent to maximise how much fuel an aircraft can save in formation flight due to different trade-offs for different types of formation. The findings presented by Bower, Flanzer, and Kroo as well as Ning show the reader the benefits and trade-offs of formation design, but the next question is if each formation with its own specific spacing and shape can be maintained in a real transatlantic route and if the benefits remain for the aircraft.

2b) Flight Path Cost-Benefit Analysis

Formation design only matters if two aircraft can successfully meet mid-air while taking the most fuel-efficient paths because before the formation flight actually happens the two aircraft have to leave from two different airports and the route taken to meet a single point requires fuel. One way to solve this problem in the most effective way possible is to treat it as a weighted Fermat problem [Kent]: minimising the weighted distance between three points (two departures and one destination), where the weighting is a coefficient representing the rate of fuel burn. With A, B being two different departure points, P being the point of intersection and C the final destination the total fuel cost depending on P can be seen below:

$$w_A \left|\left|\overrightarrow{AP}\right|\right| + w_B \left|\left|\overrightarrow{BP}\right|\right| + w_C \left|\left|\overrightarrow{CP}\right|\right|$$

In this equation, the w coefficients are the weighted quantities for each individual rate of fuel burn. The A and B quantities are weighted the same and the C quantity less because the C quantity represents the joint part of flight where the two aircraft from A and B fly together and burn less fuel per km. Due to the difference in weighting the aim of this approach isn't to find the shortest path but to change the length of the individual routes so the aircraft fly longer in formation therefore saving fuel, although it may mean extra distance travelled. Geometrically this means that the point of intersection will move further away from the destination.

There are three paths which converge into one point: two paths from A and B converging to P and then the final path from P to C. This creates three angles due to three lines, and to minimise the total weighted distance the angles reduce or increase depending on how the lines rotate around the intersection P as shown below. The angle where A and B meet is provided by using the law of cosines with the weighted quantities giving the following expression:

$$\theta_c = \cos\left(\frac{-w_A^2 - w_B^2 + w_C^2}{2w_A w_B}\right)^{-1}$$

The angle c is the angle APB. If the paths were of equal length, the angle would be 120 degrees however the length PC is weighted lower hence its magnitude will increase leading to an increase in the angle as well [Kent].

The good thing about this approach by Kent and Richards is that they've created a mathematical approach that can be solved for different flight scenarios and repeated rather than testing all possible scenarios and checking which one is the best [Kent]. A limitation of this approach is that there are some assumptions which are made, such as no airspace restrictions, constant altitude and constant rate of fuel burn. Taking into account the amount of fuel spent on take-off and landing, the best case for fuel saving for two aircraft in formation is 9.3% [Kent]. To conclude, realistically flying in formation is achievable using certain mathematical approaches to calculate the most efficient flight routes and meeting points.

Section 3: Safety requirements and effects on formation flight

After taking into consideration the safety requirements involved in formation flight so it can be safely applied to commercial flights it was realised that the aircraft might have to fly in a weaker part of the wake reducing the potential to reduce fuel consumption.

In the previous sections spacing between aircraft in different formations was investigated as an aerodynamic benefit to analyse at which point or in what formation could aircraft fully benefit from the upwash and not consume too much fuel. However, in commercial aviation this freedom to fly in whatever position at whatever distance is taken away due to various safety measures and general regulations or restrictions to keep the passengers safe. An example of this is distinguishing close formation from extended formation and finding that extended formation is safer than the close formations but the trailing aircraft in an extended formation receive a

more dispersed and decayed vortex, hence receiving less benefit [Durango]. Summing all of this together, it can be established that the distance at which aircraft fly in formation is not a priority of saving more fuel but rather keeping the passengers safe.

Formation flight is controlled by an operational constraint, not an aerodynamic one. The distance between aircraft are measured as wingspans where close formations are under 10 wingspans and extended formations are up to 40 wingspans with the extended formation being the safer option [Durango].

Figure 3: Vortex lattice model showing three aircraft in extended inverted-V formation at 10 and 20 wingspans separation. Adapted from [Durango, Lawson, and Shahneh].

The formation shown above is at the shorter end of the extended formation range because Durango, Lawson and Shahneh say that the extended formation reaches 40 wingspans. The safety constraint does subtract a benefit from the formation flight: at a distance of under 10 wingspans the vortexes don't change as much when they meet the trailing wing so the trailing aircraft feels less drag, contrary to an extended distance where over that larger distance and period of time before the vortex reaches the trailing aircraft the vortex has experienced decay and propagation meaning the trailing aircraft does not receive such a decrease in drag. Another problem arises from this decay and propagation, and it is that locating the vortex is much harder and may require more complex sensors and avionics which as well raises manufacturing costs and maintenance costs for the airline [Durango]. On top of all these technical problems there are safety regulations which simply don't allow formation flight in commercial aviation. This is clearly stated in FAR §91.111 [Durango], and even though other organisations such as the EASA and the FAA are examining this problem and pushing for applications in civil and military formation flight no real regulations exist [Ewert]. Furthermore, Durango claims that there is a "psychological factor" which comes into play reducing the chance that passenger aircraft flying close together will be accepted by the public. This is why Durango suggests that this will be first introduced in the air cargo sector due there being no passengers and hence much less risk or fear. To conclude, regulations pile up to restrict formation flight in commercial aviation and even if it would be introduced the aircraft would be flying at large distances and since larger distances mean weaker vortex the overall benefit of formation flight would be heavily decreased anyways.

Section 4: Summaries and gaps in research

Returning to the main research question, this section asks if the results provided by the sources support whether formation flight can be applied to commercial aviation due to the

common wingtip shape which reduces the vortex required for formation flight and the existing safety regulations prohibiting formation flight.

There is a lot of formation flight data provided by the 5 sources in my bibliography but none of them publish any data on how much a winglet affects the trailing aircraft because there is no study linking flights with a winglet versus without a winglet.

This section revises to what extent the five sources cover data about formation flight, the five sources being Ning, Bower, Durango, Kent, and Pahle. Even though these papers cover aerodynamic mechanisms, the effect on fuel savings, paths and flight geometry, and finally the safety associated with formation flight, they don't cover these mechanisms and their effect on fuel savings with modern airliners with winglets. Hence, the concept of formation flight is not yet applicable to commercial aviation.

A trailing aircraft flying in the upwash of a leading aircraft's wingtip vortex saves fuel [Durango], and the strength of the vortex that the trailing aircraft flies in depends on the shape of the wing [Ning]. This fuel saving depends on several conditions, as shown by the sources. Different geometries for formation and wingtip spacing sets a range. These savings ranged from 4% for two aircraft in echelon to 16% for three aircraft to 16% for three aircraft in an overlapping inverted V formation [Bower]. The second condition is routing. The meeting point for both aircrafts must be a meeting point which overall minimises total fuel burn. The best case for two aircraft is 9.3% with take-off and landing [Kent]. The third condition is safety measures. The extended spacing between aircraft to keep passengers safe means that the trailing aircraft flies on a decayed and weaker vortex [Durango]. On top of this, FAR §91.111 and the public consensus don't allow this [Durango]. The conditions above reduce the overall benefit from formation flight in commercial aviation, and they are explored differently in each source; however no source explores the individual cost for an aircraft with winglets which I describe below.

Having gone through the structure of the review and what the papers provide, it can be seen that the cost that winglets have on a trailing aircraft is never measured. The C-17 test is the only measurement taken for an aircraft with winglets hence being the most accurate, but since there is no flight test with a C-17 without winglets the cost of the winglets on fuel saving and drag reduction couldn't be measured. In terms of mentioning the cost of the winglets and that it exists Ning did that by talking about the extended methods to calculate a vortex generated by a wing with winglets, but he doesn't take any tests or measurements. By comparing blended winglets against the A380's original wingtip and exploring how the wingtip vortex varies by changing cant angles Pinto et al. isolates the effect the most. However, the study is computational, there is no further investigation about a trailing aircraft, and the baseline in the study is a wing with a wingtip device hence there is no investigation about how changing the shape affects the fuel savings of the trailing aircraft. Durango completes a study on a specific aircraft called "ATUGA" where he claims that because it has a drag efficient design it creates weaker vortexes, so the result isn't extended to winglets but only to the fact that it has a long wingspan even though the winglets are a factor of the plane being drag-efficient. To conclude,

none of the five sources present verifiable data investigating how winglets directly affect the fuel savings on a trailing aircraft because none of them have completed tests of aircraft with winglets against aircraft without winglets.

On top of the gap in research of the winglets, gaps for the safety section exist in the research. While Durango mentions that more advanced avionics might be necessary when applying the formation flight to commercial routes, he doesn't mention how advanced those avionics might be for civil use. All the organisations preventing formation flight were useful to know but no timeline was given by Durango of when new regulations may be introduced so that formation flight can be used in commercial aviation. Hence, the concept of formation flight hasn't yet been demonstrated to enter commercial aviation.

Works Cited

[1] Airbus fello'fly, 2021: reducing fuel consumption up to 10%.

[2] Bower, Flanzer, Kroo, 2009: Formation Geometries and Route Optimization for Commercial Formation Flight.

[3] Durango, Lawson, Shahneh, 2016: Formation flight investigation for highly efficient future civil transport aircraft.

[4] Ewert, Mäurer, 2023: Safety and Security Considerations on the Airbus Wake Energy Retrieval Program "fello'fly".

[5] IATA, 2026: Global Outlook for Air Transport.

[6] IATA, 2026: Industry Statistics Sheet.

[7] Kent, Richards, 2015: Analytic Approach to Optimal Routing for Commercial Formation Flight. [8] Ning, Flanzer, Kroo, 2011: Aerodynamic performance of Extended Formation Flight.

[9] Pahle, Berger, Venti, Duggan, Faber, Cardinal, 2012: An Initial Flight Investigation of Formation Flight for Drag Reduction on the C-17 Aircraft.

[10] Pinto, Marimuthu, Rajendran, Natarajan, Murugesan, 2025: Computational Analysis of Blended Winglet Designs to Reduce the Wake Turbulence on the Airbus A380 Wingtip.

[11] Portugal et al, 2014: Upwash exploitation and downwash avoidance by flap phasing in ibis formation flight.

[12] Sarkar, 2025: Winglets, what are they, and what is their role.

The Impact of Simultaneous Animated Movie Consumption on High School Students' Academic Recollection By Isabella Neill

Literature Review

Throughout the past decade, media-multitasking among adolescents has been steadily increasing as "the average 8- to 18-year-old American spends 29% of their media consumption time multitasking, an increase of 16% a decade ago" (Rideout et al.). Media multitasking is defined in two different aspects: the first in relation to multiple uses of media while the primary task is media-focused, and the second uses media while performing non-media tasks. Presently, academic media multitasking fits into both definitions as some work is computer-based while other forms are traditionally paper and pencil. As media multitasking popularizes among the younger generation, students within academic settings are more prone to media multitasking while completing their classwork. Frequent usage of media multitasking often prompts "the depletion of the attentional resource" (May & Elder), leading to decreased working memory capacity and cognitive impairment. Information retention is typically the main objective of required academic coursework, and as media multitasking begins to interfere with this aspect, then it is likely that students will face decreases in academic performance. In 2007, one group of adults and one group of fourth grade students were assessed on their performance during reading comprehension and mathematical assessments while being presented with either television or music stimuli, with a control group in a silent room. In both groups, the television stimulated group had the lowest scoring on nearly all the assessments compared to music and silent conditions (Broussard). Therefore, audiovisual media seemingly impacts academic performance more adversely than any other form of stimuli due to the increase in cognitive capacity required to process the various stimuli (Balıkçı et al.).

Cognitive Capacity & Media Multitasking Success

Cognitive capacity is the "limited capacity in which necessary information is temporarily stored and manipulated at the same time" (Guler & Aydin), while cognitive flexibility is "the ability to think flexibly and adjust to new demands" (Murphy et al.). Moreover, cognitive capacity relates to the limit of stored information an individual can hold while multitasking, and cognitive flexibility refers to the effectiveness of multitasking in general. Cognitive capacity includes the processes of attention, memory, perception, and processing speed. Attentional capacity becomes highly demanded when an individual media multitasks, "resulting in deficits in cognitive control" (May & Elder) and increased attention impulsivity, leading to more frequent task switches. Especially within tasks that require heavy concentration, like learning, background media most adversely affects the individual's cognitive ability due to this increase in attention impulsivity. However, media multitasking builds upon an individual's foundation for necessary attention switching and ability to filter irrelevant from relevant information within stimulated environments (Amez et al.). On average, an individual with increased cognitive flexibility is often able to effectively multitask "in terms of retention of information from a media

multitasking situation" (Seddon et al.). Furthermore, although media multitasking holds adverse effects over cognitive capacity and select academic situations, valuable skills derive from media multitasking from exercising cognitive flexibility and foundationally improving attention switching dexterity.

Working Memory Productivity

Working memory is defined as the "way in which we retain information over short periods simultaneously manipulating it" (Seddon et al.), meaning that this skill aids individuals in learning, reasoning, and problem-solving. In an academic environment, working memory is crucial to retaining and learning to apply information to coursework. The consumption of audiovisual material "places demand on working memory in that it requires the retention and manipulation of information from multiple media streams" (Seddon et al.) when media multitasking. If a student is media multitasking while completing schoolwork, this therefore concludes that their learning is impaired due to the excess strain on their working memory skills. An additional study states that "working memory load might increase task-irrelevant information interference" (Balıkçı et al.), meaning that the cognitive function of separating irrelevant information from relevant may also become impaired; and thus promoting distractions for a scholar amidst learning and completing coursework. Chronic multitasking depletes the working memory, and heavy media multitaskers were even "found to perform poorly in certain cognitive tasks involving task switching, selective attention, and working memory" (Lui & Wong). Moreover, if a student consistently media multitasks while completing their schoolwork, their learning retention will decrease and reduce their academic productivity.

Attention Impulsivity

The distribution of focus among two or more stimuli often leads to attention impulsivity as one study found that "attentional impulsivity positively related to the degree to which participants multitasked with media" (Uncapher et al.). The bottleneck theory of attention states that "attention can be allocated to only one task at a time, thus, multitasking is a myth" (May & Elder), as the mind instead switches between tasks rather than dividing attention. This theory proposes that one primary task can only be focused on rather than a split between secondary or tertiary tasks. From this theory, two separate hypotheses branch out in response to incorporate multitasking into attention distribution. The scattered attention hypothesis states that the executive system allocates mental resources, like attention, to carry out operations; and the trained attention hypothesis believes that media multitasking can train the brain into developing mental flexibility to adapt to filtering irrelevant stimuli out (May & Elder). This means that the first hypothesis suggests a depletion of performance quality of the primary task, while the second proposes a development in the mental capabilities of the individual performing the tasks. However, most literature supports the scattered attention hypothesis as most studies that analyzed dual-task performance found a deterioration in productivity and increased attention division when media multitasking (van der Schuur et al.). The relevance and individual-based interest of

the secondary stimulus impacts the dual-task performance as the preferred task typically garners more attention, therefore increasing the number of task switches. Media engagement requires an immersion of cognitive, behavioral, or affective aspects for an individual to immerse themselves fully to a stimulus (Perks). As an individual secondarily tunes into a form of audiovisual media, like a movie on television, their interest may deviate from their primary task due to the promoted attention impulsivity. Furthermore, media multitasking fosters attention impulsivity through increased task switches and deteriorates an individual's performance on the main objective, especially in academic settings.

Gap

Amidst the Internet Revolution, students are more prone to deviate their attention to secondary activities on their own personal devices due to this unrestricted access. This promotes the aforementioned attentional impulsivity and buffers their recollection abilities related to their primary task. In academic settings, this equates to learned material or disruptions in study methods involving active recall. Many studies uncovered the significance of background music on academic performance, yet very few studies address the relationship between academic performance while audiovisual media is present. Therefore, audiovisual media still requires thorough investigation. Additionally, the studies that do address simultaneous audiovisual media often only measure the general academic performance of students amidst media multitasking or their broad cognitive abilities throughout their multitasking habits. A specification of one cognitive ability, recall, is crucial to gather evidence of how the working memory is impacted in academic settings by several simultaneous media stimuli. Thus, a gap emerges within this field as this study will specifically address the audiovisual media of animated classic movies to incorporate imaginative visuals, musical compositions, and dialogue as potentials for affecting a student's active recollection of academic material. This study aims to improve the academic habits of modern-day high school students by providing evidence to base how their academic recollection is impacted by simultaneous animated movie stimuli.

Methodology

This study used a mix of quantitative and qualitative approaches in order to answer the research question: How does simultaneous consumption of an animated classic movie impact a high school student's academic recollection abilities? Qualitative data was incorporated so a student's prior media-multitasking habits were gauged before participation in the experiment, and quantitative was utilized so their performance on an academic task could be quantified. Similar to Broussard's 2007 study, I wanted to utilize an experimental design by exposing my participants to simultaneous movie media when completing a reading comprehension test. However, my method expands upon this study because although an experimental approach could have been solely used, it was necessary to acknowledge prior exposure to media-multitasking by surveying to ensure that the participants begin the experiment as fairly and equally as possible.

Survey Method Overview

A survey study often determines qualitative research by being applied to a larger audience with intentions of gathering mass data. Surveys involve a variety of questions including but not limited to multiple choice, rating scales, or open-ended questions. My survey was given only to the participants of my experiment so that I could understand their exposure to media-multitasking in tangent to academics prior to the experiment, and their own perception of their academic success when media-multitasking. To compare their answers with their experimental test score results, the participants were asked to provide their name and grade, however their demographic information will not be referenced throughout this report to maintain confidentiality. The survey begins by asking the participants how frequently they multitask during the week, and to name their most frequent form of background media if applicable. As it progresses to the next section, the survey requests that students assess the frequency and effectiveness of their academic media-multitasking habits using multiple choice questions. To ensure their understanding of the practice, the definition of academic media-multitasking is supplied as "when one or more simultaneous tasks being completed are related to schoolwork." The final section requests that the participants address the perceived impact of media-multitasking on their academics by incorporating a rating scale and various open-ended questions. This study produced qualitative data that aided the numerical test score data to explain why the participants performed how they did on the assessment.

Experiment Method Overview

An experimental study involves a significant sample size of participants that are randomly assigned to a conditional group. In an average experiment, there is one control group and at least one experimental group in which the researcher manipulates one select variable at a time. The independent variable is the manipulated factor of the experiment, and the dependent variable is the measurable data that directly corresponds to the manipulated variable. In my experiment, the manipulated variable is the silent condition versus the background movie condition in which the students will complete their assessment. The constants of the experiment are the article that the students will read and attempt to memorize in the allotted time and the reading comprehension assessment questions, and the amount of time each group is allotted for each task. This study produced quantitative data in the form of the comparative test scores between the silent and movie groups.

Data Source

My study was presented to all students taking the psychology courses at a southeastern Pennsylvania high school because this class is offered to all grade levels, so I was able to measure the recollections of students aged 14-18 years old. This age range is crucial in order to answer my research question by including all high school ages in my study. This tactic is called stratified sampling, using a previous existing group, selecting participants, and forming a new group based off of these participants. The participants were chosen through randomly selecting

around four students from a variety of grade levels per class period. However, this sample may limit the significance of the data if applicable to high schoolers outside of this specific high school, but the data can aid in producing a significant and considerable trend of data in future large-scale national or global studies. Additionally, I wanted the participants to have previous exposure with experimental practices and procedures, and because psychological experiments are part of the psychology course, the participants were informed of their rights and the outline of a traditional psychological experiment. The participants were then split into two groups and assigned to be in a silent condition or a movie condition. I was able to gather twenty-four student participants with twelve participants in the movie condition, and twelve participants in the silent condition.

Materials

I constructed the survey questionnaire and the reading comprehension assessment questions with complete originality based upon the desired metrics I wanted to measure. Attached in Appendix A below are the full online survey questions and the reading comprehension assessment questions and answer key. The survey was constructed on Google Forms with intentions of measuring qualitative data based on the participants' opinions on the efficiency of media-multitasking alongside academic settings. The reading comprehension questions focused on main ideas and takeaways from the reading comprehension article to gather quantitative data in order to measure their recollection abilities. The reading comprehension article, "Why the Emerald Swallowtail Butterfly Changes Color" was derived from the American Museum of Natural History due to its 1350L, suitable for grades 9-10. This article will additionally be attached below in Appendix B. The reading comprehension questions were created by myself in a multiple choice format that addressed the main topics of the article, additionally attached in Appendix C.

Procedure & Data Collection

Before participation in the experiment, student participants were instructed to participate in a preliminary survey with an estimated completion time of ten minutes. Following, students within both groups were given ten minutes to read and memorize the reading comprehension article's material. Students within the experimental group read while the "Circle of Life" scene from *The Lion King* played concurrently. However, the control group students read their article in complete silence. After the ten minute timer rang out, the students were instructed to stop reading and the article was taken. The students were then given the assessment and five minutes was provided to answer the eight questions based on their memory and recollection of the article's material. The assessments between the two groups were graded and separately input into two online spreadsheets in order to organize the data for later comparison and analysis.

Ethical Considerations

Informed consent was obtained through both student assent and parental consent forms that addressed the layout of the experiment, the usage of participant data, and the experiment's objective. Topics related to the survey were also addressed in the consent forms, and participants were ensured of their participant rights, and that they could skip any questions they felt they cannot answer effectively. Prior to the study, participants were briefed on the purpose of their data and what would be expected of them as participants. Although names of the participants were collected on both the survey and the assessments, the participants were assured that their names will not be referenced in this report, and that their data will be kept confidential only to the experimenter. There were no foreseeable risks within the study, but participants were allowed to withdraw at any point if they felt uncomfortable or that they could not effectively participate. This study was approved by the Institutional Review Board as it ensured correct confidentiality measures, detailed participant rights, and an informative layout of both the experiment and survey.

Data Analysis

After the survey was completed, participant survey responses gathered via online spreadsheet to identify commonalities and trends among students regarding academic media multitasking. Their perceptions of productivity and personal challenges regarding simultaneous media forms were compared through table organization. When the assessments were completed, I graded the students by correct answers out of the eight available questions using my constructed answer key as shown in Appendix C. The experimental versus control group scores were then inputted into separate online spreadsheets and rounded to the nearest whole percentages. The two groups' mean scores were then able to be compared through the calculation of group averages. Additionally, standard deviations of the two groups were calculated in order to understand if the score results were an average representation of student performance. Finally, to determine the extent to which the test score results were significant, an unpaired two-tailed T-test was conducted with an additional P-value calculation.

Results

This section presents the results of the study, which investigated the impact of an animated classic movie stimulus on a student's academic recollection abilities, measured by a reading comprehension assessment. Data was collected from twenty-four psychology student participants, including a preliminary survey and a reading comprehension experimental test to assess their academic recollection. The previous proposed hypothesis was that students presented with the animated movie background condition will perform worse than the students taking the reading comprehension assessment in silence because their academic recollection will face greater difficulty.

Survey Results

Before the experiment, students were prompted to identify their general media multitasking habits to understand their previous identifiable habits within academic media multitasking and to identify any trends across the participant sample. Full survey questions are accessible within the Appendix A, however the below reported statistics are most integral to answering the question: How does simultaneous consumption of an animated classic movie impact a high school student's academic recollection abilities?

Figure 1

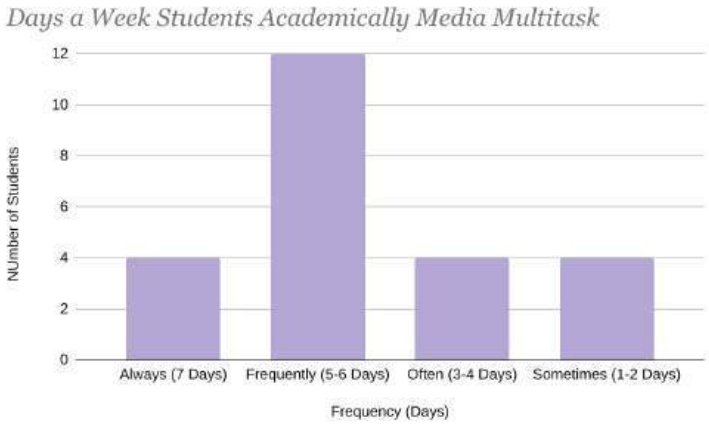

Figure 1 illustrates that the majority of students media multitask 5-6 days out of the week, and no students report that they do not academically media multitask throughout the week.

Figure 2

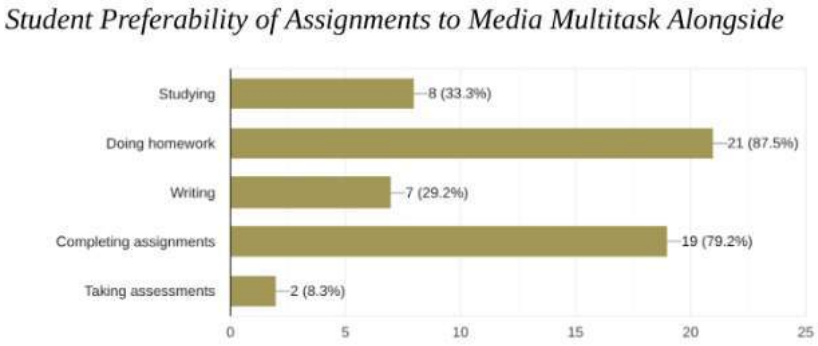

As shown in Figure 2, students were prompted to select which academic tasks were preferred to media multitask alongside. The majority of the students preferred to multitask while doing homework and completing assignments.

Figure 3

Student Perception of Most Difficult Activities to Media Multitask Alongside

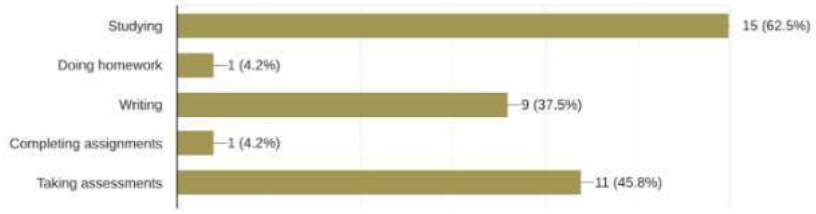

Students were then asked to identify the most difficult tasks to media multitask while completing. Figure 3 portrays that the majority of students perceive that studying and completing assessments are most difficult to media multitask alongside.

Experiment Results

During the experiment, students completed an eight-question assessment within their conditioned group. The test results are compared below, with their scores reflected as percentages. Percentages are calculated by the number of questions right out of eight questions, with the result rounded to the nearest whole number.

Table 1

Comparative Score Results: Silent vs Movie Conditioned Groups

Participant #	Silent Group	Movie Group
1	100	100
2	100	88
3	100	88
4	100	88
5	88	88
6	88	88
7	88	75
8	88	75
9	88	63
10	88	63
11	75	50
12	63	50

As shown in Table 1, the movie group scored lower than the silent group on the reading comprehension assessment.

Experiment Statistics

In order to determine the relationship between the two groups' test scores, the mean test score and standard deviation of both groups were independently calculated. The null hypothesis states that animated classic movie background media does not impact a student's academic recollection abilities, thus the reading comprehension test scores between the two groups will not significantly differ. An alternative hypothesis suggests that animated classic movie background media will impact a student's academic recollection abilities, and therefore the movie condition's test scores will be lower than the silent condition's scores.

Figure 4

Silent Group Calculations

$$\text{Avg Score} = \frac{Sum\ of\ Scores}{Number\ of\ Participants} = \frac{88+88+100+88+88+75+88+88+63+100+100+100}{12} \approx 88.83\%$$

$$\text{Standard Deviation} = \sqrt{\frac{\Sigma(xi-\mu)^2}{N}} = \sqrt{\frac{(88-88.83)^2+\ldots(100-88.83)^2}{12}} = \sqrt{\frac{1361.67}{12}} = \sqrt{113.472} \approx 10.65$$

Figure 5

Movie Group Calculations

$$\text{Avg Score} = \frac{Sum\ of\ Scores}{Number\ of\ Participants} = \frac{88+50+63+63+88+100+75+88+88+88+50+75}{12} \approx 76.33\%$$

$$\text{Standard Deviation} = \sqrt{\frac{\Sigma(xi-\mu)^2}{N}} = \sqrt{\frac{(88-76.33)^2+\ldots(75-76.33)^2}{12}} = \sqrt{\frac{2986.67}{12}} = \sqrt{248.889} \approx 15.78$$

Table 2

Comparative Statistics Between the Silent vs Movie Group

Silent Group	Movie Group
N = 12 Mean = 88.83 SS = 1361.67 Sample Variance = 123.79	N = 12 Mean = 76.33 SS = 2986.67 Sample Variance = 271.52

Figure 6

Comparing the Reading Comprehension Tests with a Two-Tailed Unpaired T-Test

$$\text{Pooled Variance} = \frac{123.79+271.52}{2} = 197.65$$

T-Test = 2.1779

P-Value = 0.040416

As shown in Figures 4 and 5, the mean average test score is higher for the silent group versus the movie group. Additionally, the individual standard deviation expanded formulas as portrayed in Figure 4 and 5 prove that the variances between the two groups are unequal as the movie group's standard deviation exceeds the standard deviation of the silent group. Figure 6

proves that rejection of the null hypothesis is possible because the p-value (0.040) is less than 0.05.

Discussion of Results

Qualitative survey data and experimental statistics were integral processes in order to effectively answer the question: How does simultaneous consumption of an animated classic movie impact a high school student's academic recollection abilities? Through survey data, the baseline of the average student's media multitasking habits was able to be identified in order to understand the preexisting consumption of media throughout academic settings. This survey also tested the student participants' perception of their own personal academic productivity during media multitasking. An experiment was then conducted to test the hypothesis that an animated classic movie will impact a student's academic recollection when simultaneously consumed. This hypothesis was supported by Broussard's 2007 study in which a television-stimulated group performed poorly on all assessments, and Seddon et al. 's conclusion that media multitasking creates a large demand on working memory productivity.

Survey Data Interpretation

Every student participant reported that they media-multitask throughout the week at least once, suggesting the confirmation of the hypothesis that academic media multitasking is on the rise within the new generation of students. The majority of students reported that they media-multitask 5-7 days out of the week, therefore this habit is increasing in commonality and frequency. The measurement of this frequency was crucial in order to test if any students developed cognitive flexibility through the practice of academic media multitasking. The cognitive flexibility theory was suggested by Seddon et al. because the average media multitasker often has increased cognitive flexibility, and therefore should, theoretically, be able to adapt and perform sufficiently within simultaneous media-stimulated settings. By comparing the movie condition's self-reported frequency of media multitasking to their test score results, it is suggested that the majority of students cannot develop this cognitive flexibility in academic settings, as the media-multitasking group performed with an average of 76.33% in comparison to the control group's average of 88.83%. If the deviation between the two averages was smaller, then this could suggest that cognitive flexibility was adapted by the student participants within the movie condition, however this large of a deviation suggests that the participants within the movie condition could not thrive under the simultaneous audiovisual media.

Experimental Data Interpretation

As shown in Table 2, the silent group vastly outperformed the movie-conditioned group as their mean score on the reading comprehension assessment was 88.8% in contrast to the movie group's mean of 76.33%. Because the p-value of the unpaired T-test was calculated to be 0.04, this suggests that the alternative hypothesis that simultaneous animated movie media does impact a student's recollection within academic settings can be accepted. This implies that

simultaneous audiovisual media reduces a student's academic performance abilities. Additionally, only one student in the experimental group received a perfect score of 100% on the reading comprehension assessment, compared to the four students who received perfect scores in the silent group. Clearly, the simultaneous background audiovisual media hindered the students' scores and performance capabilities, therefore this form of animated classic movie media is a disruption to student recollection, as proven by the comparison of the two groups' scores.

Implications

As found within both the survey and the experiment, the student participants' academic recollection abilities were significantly impacted by the animated audiovisual media when simultaneously performed. Thus, the demand of audiovisual media on a student's cognition proves to mend the gap of how this form of media impacts a student's recollection of academic material. By modeling a similar experiment to Broussard's 2007 study in which adults and children were tested on their reading comprehension and mathematical capabilities while watching television simultaneously, this study aimed to narrow the parameters to specifically measure recollection through reading comprehension amidst animated movie media. Although some studies found that music could improve focus, the imposition of both a musical soundtrack paired with visual animations proved to be a distraction while performing in a cognitively demanding task, such as a reading comprehension assessment. While this study tested the impact of a specific form of audiovisual media, animated classic movies, future studies can model a similar experiment to test the impact of other forms of media on student academic recollection by replacing the movie media with the desired form of background media.

Limitations

However, this study does include limitations to its generalizability. While the survey questionnaire aimed to measure the approximate habits of a student media-multitasker within academic settings, the potential of a self-reporting bias is possible. Students could have over-approximated their productivity when answering questions about their performance while media-multitasking. Selective recall also could have prevented student participants from answering questions about the frequency of their media-multitasking habits fully accurately. Additionally, the measurement of recollection abilities is subjective, and the assessment process used in this study was based on its ability to quantitatively measure academic recollection. Though this study could suggest a trend with overall recollection abilities, the interpretation of this study focuses solely on academic recollection due to its assessment-type. Finally, there are limited tangible strategies to measure a student's existing recollection hindrances and thus this study does not account for students with presently existing memory disabilities or otherwise adversity.

Conclusion

The goal of this study was to answer if simultaneous animated classic movie media would impact student recollection in academic tasks. It was found that simultaneous animated classic movie media negatively impairs a student's academic recollection abilities. In order to quantify recollection as a measurable unit, students were tasked with a memory-based academic reading comprehension assessment. To measure their recollection abilities, this test was scored based on accuracy of the information they read and attempted to memorize. As shown in this study, students within the animated movie conditional group performed lower on a standard reading comprehension assessment test than the students that took the test in silence. Although students acknowledge media multitasking as a hindrance to academic work, students prefer to maintain participation in academic media multitasking habits. Likely, this correlates to the increased access of online media due to the Internet, where students are more easily to deviate from their preferred tasks in order to increase their satisfaction while completing mundane tasks, like schoolwork. Thus, the younger generation is more susceptible to media multitasking within academic settings. The negative impact on academic recollection of this practice should be acknowledged and recognized within academic communities in order to ensure student awareness of the implications of media multitasking. Future research should take away the hindrance of simultaneous audiovisual media into consideration when observing the other emerging academic habits of students in tangent to device usage. In order to obtain stronger results in accuracy, future research could divide the experimental versus control groups based on the self-reported frequency of multitasking habits. This could help to better understand the cognitive flexibility theory by comparing more experienced media multitasking performances against weaker media multitaskers. Furthermore, the use of simultaneous classic movie audiovisual media impairs a student's academic recollection amidst the performance of scholarly tasks, and thus students should become aware of this hindrance in order to reassess and maximize their academic performance potential.

Works Cited

Amez, S., Baert, S., Heydencamp, E., & Wuyts, J. (2021). Does Multitasking Affect Students' Academic Performance?: Evidence from a Longitudinal Study. IZA - Institute of Labor Economics. http://www.jstor.org/stable/resrep63845

Balıkçı, İ., Tok, S., & Binboğa, E. (2023). The effect of background audio and audiovisual stimuli on students' autonomic responses during and after an experimental academic examination. Brain and Behavior, 13, e3153. https://doi.org/10.1002/brb3.3153

Beentjes, J. W. J., Koolstra, C. M., & van der Voort, T. H. A. (1996). Combining background media with doing homework: Incidence of background media use and perceived effects. Communication Education, 45(1), 59–72. https://doi.org/10.1080/03634529609379032

Broussard, K., (2007) "The Effects of Television, Music, and Silence Conditions on Performance on Reading Comprehension and Math Word Problem Tests: A Developmental Study" Honors Theses. 32. https://repository.lsu.edu/honors_etd/32

Guler K, Aydin A. Comparing the Working Memory Capacity with Cognitive Flexibility, Cognitive Emotion Regulation, and Learning Styles of University Students: A Domain General View. Dement Geriatr Cogn Dis Extra. 2022 Nov 14;12(3):131-149. doi: 10.1159/000526226. PMID: 36465839; PMCID: PMC9710454.

Lui, K. F., & Wong, A. C. N. (2012). Does media multitasking always hurt? A positive correlation between multitasking and multisensory integration. Psychonomic bulletin & review, 19(4), 647-653.

May, K. E., & Elder, A. D. (2018). Efficient, helpful, or distracting? A literature review of media multitasking in relation to academic performance. International journal of educational technology in higher education, 15(1), 1-17. https://doi.org/10.1186/s41239-018-0096-z

Murphy, K., McLauchlan, S., & Lee, M. (2017). Is there a link between media-multitasking and the executive functions of filtering and response inhibition?. Computers in Human Behavior, 75, 667-677. https://doi.org/10.1016/j.chb.2017.06.001

Perks, L. G. (2021). Binge-Watching Conditions and Multitasking: The Enjoyable Ephemeral. In M. Jenner (Ed.), Binge-Watching and Contemporary Television Research (pp. 82–97). Edinburgh University Press. http://www.jstor.org/stable/10.3366/j.ctv27qzrrt.10

Rideout, Victoria & Foehr, Ulla & Roberts, Donald. (2010). Generation M[superscript 2]: Media in the Lives of 8- to 18-Year-Olds. Henry J. Kaiser Family Foundation.

Seddon, A. L., Law, A. S., Adams, A. M., & Simmons, F. R. (2021). Individual differences in media multitasking ability: The importance of cognitive flexibility. Computers in human behavior reports, 3, 100068. https://www.sciencedirect.com/science/article/pii/S2451958821000166

Uncapher, M.R., K. Thieu, M. & Wagner, A.D. (2016). Media multitasking and memory: Differences in working memory and long-term memory. Psychon Bull Rev 23, 483–490 https://doi.org/10.3758/s13423-015-0907-3

van der Schuur, Winneke & Baumgartner, Susanne & Sumter, Sindy & Valkenburg, Patti. (2015). The consequences of media multitasking for youth: A review. Computers in Human Behavior. 53. 204-215. 10.1016/j.chb.2015.06.035.

The Baseline of One: A Safety-First Framework for Longitudinal Smartphone Monitoring of the Diabetic Foot

By Eshaan Naik[1] and David G. Armstrong, DPM, MD, PhD[2]

Abstract

Diabetes-related foot complications are often managed through episodic clinical encounters even though clinically meaningful changes can occur between visits. This paper proposes a smartphone-based framework for longitudinal foot monitoring organized around a single design concept: a "baseline of one." Rather than treating each photograph primarily as a comparison with a population average, the framework treats the same patient at an earlier time as the principal reference, and it treats explicit abstention and clinician escalation, rather than autonomous classification, as the primary output. The framework is illustrated through SoleIQ Health, a software-only platform that uses a patient's existing smartphone rather than dedicated hardware. The proposed system combines baseline risk context, guided multi-angle image capture, pre-analysis image-quality checks, patient-reported symptom mapping, a computer-vision layer, a separate narrative layer, and a longitudinal timeline. The proposed system is therefore designed to learn and display how the same patient's feet change over time, to declare when an image cannot support a conclusion, and to route uncertainty to a person rather than resolve it internally. This distinguishes the framework from population-trained, single-image severity classifiers, which ask which category a photograph most resembles rather than what has changed in this foot since the last capture. The approach is intended to complement, not replace, established clinical assessment. The paper also develops a safety-first validation framework emphasizing image quality, reproducibility, escalation sensitivity, false-reassurance risk, subgroup performance, usability, privacy, and clinician interpretability. No patient data, prospective outcomes, or model-performance results are reported; the contribution is conceptual and methodological. Some components described here are implemented in the current build; others, including the versioned patient-specific baseline, registration of serial images, change classification, and the explicit abstention and escalation pathway, remain proposed rather than built. No component has been clinically validated, and implementation of a software function should not be read as demonstrated clinical performance. The proposed architecture suggests that the most useful role for smartphone artificial intelligence in diabetic foot care may not be autonomous diagnosis, but structured longitudinal observation that helps patients and clinicians notice change earlier and communicate it more consistently.

Keywords diabetic foot; artificial intelligence; smartphone imaging; remote monitoring; longitudinal analysis; computer vision; digital health

Introduction

Diabetic foot disease is a persistent challenge because risk develops over time while

[1] SoleIQ Health, Cary, North Carolina, United States
[2] Keck School of Medicine of USC, University of Southern California, Los Angeles, California, United States
Conceptual Review and Proposed Technical Framework

clinical observation is intermittent. Peripheral neuropathy can reduce a person's awareness of minor trauma, and diabetes-related foot complications can progress from seemingly small changes to ulceration, infection, hospitalization, and amputation. Current standards therefore emphasize regular foot evaluation, risk stratification, preventive education, and more frequent inspection for people with sensory loss or previous ulceration or amputation (American Diabetes Association Professional Practice Committee for Diabetes).

Deciding how often a foot should be examined has long depended on identifying who is most likely to ulcerate, with neuropathy, deformity, and prior ulceration among the factors that carry the most practical weight (Armstrong, Lavery, and Harkless). The International Working Group on the Diabetic Foot (IWGDF) similarly places prevention, structured risk assessment, and repeated surveillance at the center of diabetic foot care (International Working Group on the Diabetic Foot, “Practical Guidelines”; International Working Group on the Diabetic Foot, “Prevention Guideline”).

The clinical problem is not simply whether a patient receives a foot examination. It is what happens between examinations. Armstrong, Boulton, and Bus describe diabetic foot disease as a condition in which recurrence and continued risk remain important even after a wound has healed. This supports a broader view of prevention: a healed or apparently normal foot is not the end of monitoring but the beginning of continued surveillance. Home monitoring research using temperature has also shown that repeated patient-generated signals can contribute to an early-warning strategy, particularly when an abnormal signal leads to an appropriate behavioral or clinical response (Lavery et al.; Bus et al.). Sensors, wearables, and telehealth are increasingly described as a connected ecosystem for remote management of diabetes-related foot disease rather than as isolated devices (Golledge et al.).

Smartphones create a different pathway for generating such signals. A camera is already present in a device many patients use every day, and software can guide capture, standardize views, preserve a longitudinal record, and route structured information to a care team. Yet a camera alone does not create a clinically useful monitoring system. A photograph can be blurred, poorly illuminated, incompletely framed, or disconnected from the patient's risk history. A model can also produce an apparently precise output that exceeds what the image can support. For smartphone monitoring to be useful, the system must therefore be designed around observation quality, change over time, uncertainty, human escalation, privacy, and clinical context.

This paper proposes a design framework for that purpose using SoleIQ Health as the working system. SoleIQ Health is a smartphone-based, software-only platform focused on early recognition of diabetic foot change and amputation prevention. Its distribution model is physician-first rather than direct-to-consumer, and its outputs are intended to support investigational longitudinal monitoring and triage support rather than function as autonomous diagnosis. The core research question is: How should a smartphone-based artificial intelligence system be structured so that patient-generated foot images become safer, more interpretable, and more clinically useful between visits?

The answer developed here is that the system should be treated less as a one-time classifier and more as a longitudinal monitoring pipeline. That pipeline begins before an image is analyzed, continues through model inference and patient context, and ends only when information is presented in a form that supports an appropriate next step. The proposed framework centers on five linked ideas, of which the patient-specific baseline is the organizing one: standardized guided capture, a versioned baseline of one, separation of visual and narrative reasoning, dual patient-and-clinician outputs with explicit abstention, and prospective validation against safety-oriented endpoints rather than headline accuracy alone.

Figure 1. Proposed clinical framework for smartphone self-monitoring of diabetic foot.

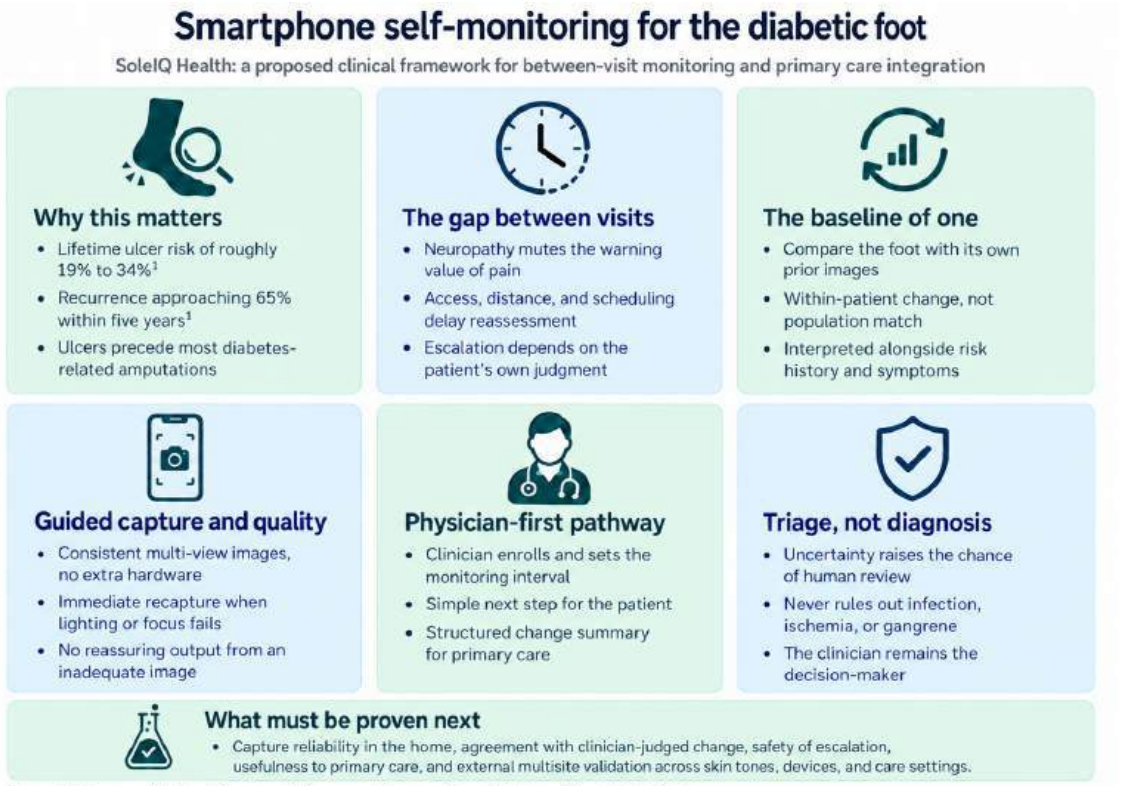

Figure 1. Proposed clinical framework for smartphone self-monitoring of the diabetic foot.
Overview of key components and considerations for between-visit monitoring and integration with primary care using SoleIQ Health. [1]Armstrong DG, Boulton AJM, Bus SA. Diabetic Foot Ulcers and Their Recurrence. N Engl J Med. 2017;376(24):2367–2375. doi:10.1056/NEJMra1615439.

Related Work

Several groups have established that patients can generate usable foot images at home and that those images can be reviewed serially. The Foot Selfie system demonstrated the feasibility of routine home plantar-foot photography with serial within-person review in a high-risk population (Swerdlow et al.). Patient-owned wound surveillance has been evaluated for feasibility, adherence, and the role of clinician feedback in sustaining engagement (Lo et al.). The MyFootCare application examined usability and engagement in longitudinal photo-based self-monitoring, including the practical difficulty of maintaining a photographic routine over time (Ploderer et al.). On the analytic side, smartphone wound-imaging applications have been

clinically validated for wound measurement in diabetic foot ulcers (Chan et al.), and automated ulcer detection has been evaluated in real-world conditions on low-cost handsets, which bears directly on device access (Cassidy et al.). Broader reviews place these tools within a wider ecosystem of sensors, wearables, and telehealth for remote management of diabetes-related foot disease (Golledge et al.).

Most closely related is recent work developing and validating a predictive artificial-intelligence framework for diabetic foot ulcer monitoring and severity assessment, positioned toward self-monitoring and primary care integration (Satheesh et al.). That work trains on population data and assigns a severity classification from a single image, and it reports model-performance results for that task. The present manuscript is deliberately different in target and in kind. It reports no predictive model, no prediction horizon, and no performance data; it proposes longitudinal within-person comparison against a versioned patient-specific baseline rather than population-referenced classification of an isolated image; and it treats abstention and clinician escalation, rather than a severity label, as the primary output. The two lines of work are complementary. A population-trained classifier and a within-person change detector answer different questions, and a hybrid of the two is one of the comparisons this paper argues should be tested directly.

Taken together, prior studies have generally emphasized image acquisition, wound-size estimation, population-based classification, or remote clinician review. The framework proposed here takes a different primary target: longitudinal detection of within-person change using a versioned patient-specific baseline, explicit model abstention when the image cannot support a conclusion, and a predefined pathway to clinician review. What this paper adds is not a better classifier but a safety architecture for the monitoring problem, together with a validation program capable of testing whether the baseline-of-one premise holds.

Conceptual Approach

This article is a conceptual review and system-design paper. It does not report a clinical trial, patient cohort, diagnostic-accuracy study, or model-development experiment. No new patient data were collected, no clinical images were analyzed for this manuscript, and no claims of clinical effectiveness are made.

The framework was developed from three inputs. First, established diabetic foot guidance was used to define the clinical context in which risk stratification, repeated inspection, referral, and recognized classification systems matter. Second, published evidence on home monitoring was used at a high level to support the principle that patient-generated longitudinal signals can be meaningful when they trigger appropriate action. Third, the existing SoleIQ Health product architecture was mapped into a safety-oriented monitoring pipeline.

The analysis is therefore design-oriented rather than empirical. Each component of the proposed system is evaluated according to a simple question: what failure can occur at this stage, and what design requirement could reduce that failure? The paper then translates those requirements into a staged validation plan. The resulting framework is intended to be testable.

Where a feature is described as potentially useful, that statement is a hypothesis for future validation rather than a statement of demonstrated performance.

A Smartphone-Only Monitoring Model

SoleIQ Health begins from the premise that the highest-value function of a smartphone may be repeated observation rather than one-time classification. In standard care, a clinician can inspect the skin, test protective sensation, evaluate pulses, review ulcer history, and integrate many findings that a photograph cannot capture. The smartphone should not attempt to reproduce the entire examination. It can, however, help preserve a visual record of what the patient sees between visits and organize that information so changes are easier to recognize.

Monitoring and diagnosis answer different questions: diagnosis asks what a condition is, while monitoring asks whether something has changed and whether the person should seek further evaluation. The narrower target is the more defensible one for a patient-facing system, and it matches the reality that a patient may have access to a phone far more frequently than to a specialist.

Figure 2. Comparison of episodic office assessment under current diabetic foot ulcer (DFU) care standards with the proposed SoleIQ framework.

Characteristic	Episodic office assessment (standard care) Intermittent, in-clinic evaluations as part of routine care (may include risk stratification, self-inspection, footwear/offloading, temperature monitoring, and caregiver participation)[1]	Proposed SoleIQ framework Structured, frequent capture at clinician-set intervals complementing standard care
Observation window	Intermittent visits (e.g., routine appointments) with variable timing	Structured, frequent capture at clinician-defined intervals
Basis for assessment/ comparison	Provider examination along with existing clinical information (risk assessment, patient history, self-reports, etc.)	Patient-specific baseline with longitudinal comparison over time
Detection of change	May rely on patient reporting, clinical examination, and available data (e.g., prior images, temperature, symptoms)	Designed to highlight serial visual changes; clinical benefit remains to be tested
Image quality	Variable, often opportunistic photos taken during clinic visits	Guided image capture with framing and lighting assistance to improve consistency
Handling uncertainty	Clinical judgment; additional follow-up or testing as needed	Explicit uncertainty and clinician review; effect on false-alert burden requires evaluation
Clinician output	Integration with overall clinical assessment, informing ongoing care plan (e.g., education, footwear/offloading, referrals)	Single, organized view of image timeline, trends, symptoms, and risk indicators to support clinical decision-making
Evidence required	Established standard of care supported by guidelines and clinical practice	Prospective, diverse, multisite validation required

Figure 2. Comparison of episodic office assessment under current diabetic foot ulcer (DFU) care standards with the proposed SoleIQ framework.

The proposed SoleIQ workflow therefore starts with context rather than image interpretation. During onboarding, the patient completes a risk questionnaire that records elements such as diabetes history, HbA1c, neuropathy, and prior ulcer or amputation history. These factors do not replace a clinical risk examination, but they provide background for how strongly a new visual change should be treated. In a physician-first deployment, this baseline context can also be confirmed or supplemented by the care team.

Next, the patient is guided through a consistent multi-angle capture sequence. The purpose is not merely convenience. Repeating similar views reduces one source of variation in

longitudinal comparison. If the first image shows the plantar surface from one angle and a later image shows only the dorsal surface, the model may be asked to compare information that is not actually comparable. A guided sequence turns image acquisition into a reproducible measurement procedure.

Standardization is what makes that comparison trustworthy, and it has to be specified rather than assumed. A longitudinal capture protocol should define the required views, including planar forefoot, midfoot, and heel, dorsal, medial and lateral borders, the interdigital spaces, and the posterior heel, and should record which views are missing rather than silently proceeding without them. Distance and angle should be constrained by on-screen guidance, ideally including a translucent overlay of the patient's own prior capture so that the current image is registered against the baseline at the moment of acquisition rather than only afterward. A physical scale reference and a color reference, such as a printed calibration card placed in frame, allow size measurements and color comparisons to be normalized across sessions and lighting conditions. Lighting guidance and a consistent flash policy reduce a further source of variation that is easily mistaken for erythema or pallor.

The system should also record the technical conditions of each capture: device model, camera and lens used, operating-system version, and any on-device image processing applied, since computational photography pipelines apply scene-dependent tone mapping and color correction that differ across manufacturers and change with software updates. When a patient changes phones or the operating system materially alters image processing, the platform should flag the discontinuity, treat subsequent comparisons against the prior baseline with reduced confidence, and prompt a re-baselining capture rather than reporting the resulting difference as clinical change. Where the platform permits access to minimally processed image data, retaining it alongside the displayed image preserves the option of reanalysis under a consistent pipeline. In a longitudinal system, technical drift can look like biological change, and the design should make that failure mode detectable rather than invisible.

Before any analysis occurs, the software performs on-device quality validation. Images that are blurred, poorly framed, or otherwise inadequate are rejected and recaptured. This gating is implemented in the current build; whether it yields more interpretable images without driving abandonment has not been tested. This is a central safety control because downstream intelligence cannot recover visual information that was never captured. A system that gives no answer when input is unreliable can be safer than a system that always gives an answer.

The patient can then identify tender, painful, or otherwise concerning regions using an interactive foot map. Neuropathy means absence of pain cannot be treated as reassurance, but self-reported symptoms remain useful context when present. Combining structured self-report with image data also prevents the photograph from becoming the only source of information.

Finally, SoleIQ is designed to generate two outputs from one analysis. This dual-output structure is only partially implemented, and several of the elements described here depend on components that remain proposed (Table 2). The patient would receive a plain-language summary and next-step guidance, together with a change indicator that depends on the proposed

change-classification component. The clinician would receive a structured snapshot intended to include observed visual features, change from prior images, image-quality status, relevant patient-reported context, an explicit uncertainty or indeterminate status, and a suggested level of review or escalation; the change, abstention, and escalation elements of that snapshot are proposed and unvalidated. That snapshot deliberately excludes differential considerations and any image-derived ulcer grade. A photograph alone cannot establish depth, ischemia, infection, bone involvement, protective sensation, or perfusion, and presenting inferences that depend on those findings would misrepresent what the image supports. The purpose of the dual-output design is not to hide information from the patient; it is to translate the same underlying event into language appropriate for two different decisions. A patient needs clarity about what to do next. A clinician needs enough structure to decide whether the finding warrants review, escalation, or in-person assessment.

Figure 3. Workflow for smartphone-based self-monitoring of the diabetic foot.

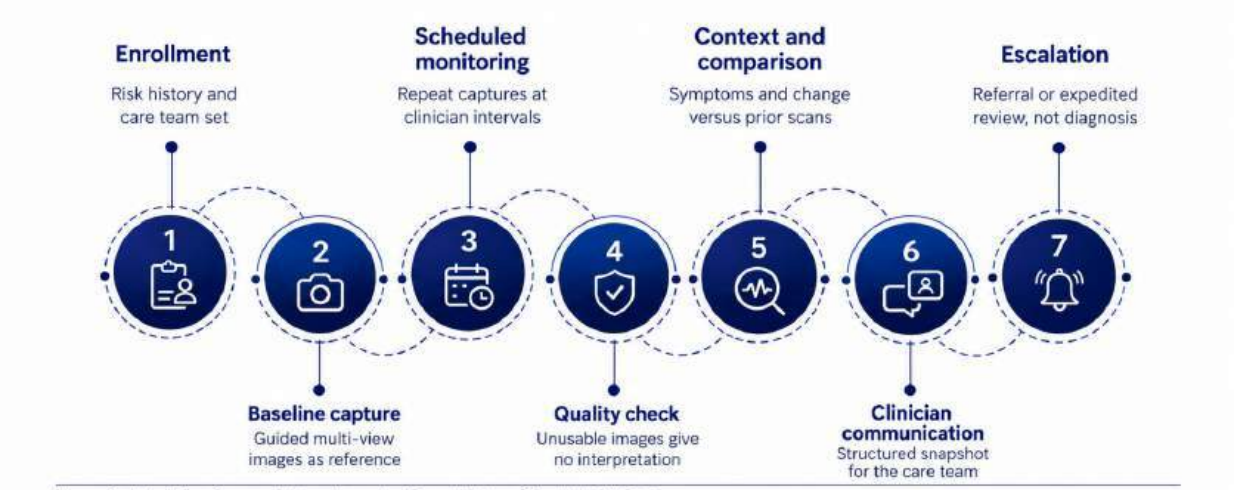

Figure 3. Workflow for smartphone-based self-monitoring of the diabetic foot.
Seven key steps from enrollment through escalation: (1) enrollment; (2) baseline capture; (3) scheduled monitoring; (4) quality check; (5) context and comparison; (6) clinician communication; and (7) escalation.

Source: Adapted from Armstrong DG, Boulton AJM, Bus SA. Diabetic Foot Ulcers and Their Recurrence. *N Engl J Med.* 2017;376(24):2367–2375. doi:10.1056/NEJMra1615439.

Table 2. Implementation status of the proposed components

The framework described in this paper mixes components that exist in the current SoleIQ build with components that are proposed but not yet built, and none of them has been clinically validated. Table 2 separates the two so that a conceptual architecture is not mistaken for a validated clinical system. Nothing in the third column has been established; it describes the evidence that would be required before any corresponding claim could be made.

SoleIQ System Elements: Implementation Status and Evidence Requirements

System version/status date: August 2026 | Development period: October 2025 – August 2026

System element	Current status (as of August 2026)	Evidence still required
Guided multi-angle capture sequence	**Implemented** Confirmed in current software build	Capture success and view-completeness rates in patients at home, including those requiring caregiver assistance.
On-device image-quality gating and recapture	**Implemented** Confirmed in current software build	Whether rejection and recapture yield more interpretable images without driving abandonment.
Baseline risk questionnaire and patient context	**Implemented** Confirmed in current software build	Agreement between patient self-report and the clinical record; incremental value over the image alone.
Interactive foot map for symptom self-report	**Implemented** Confirmed in current software build	Whether symptom mapping adds information beyond the image in neuropathic patients.
Server-side analysis service, authenticated access, and patient-level data isolation	**Implemented** Confirmed in current software build	Independent security review; completeness and immutability of audit logs.
Separation of vision layer from narrative layer	**Partially implemented** Confirmed in current software build	Faithfulness of narrative output to structured visual features; separate attribution of perception versus communication error.
Dual patient and clinician outputs	**Partially implemented** Confirmed in current software build	Patient comprehension testing; clinician-rated sufficiency of the snapshot for a triage decision.
Versioned patient-specific baseline ("baseline of one")	**Proposed**	Prospective comparison of population-only, baseline-only, and hybrid models on the same dataset.
Registration and alignment of serial images	**Proposed**	Alignment error under realistic variation in angle, distance, lighting, and device.
Change classification as stable, improving, worsening, or indeterminate	**Proposed**	Reproducibility across repeated captures; agreement with blinded expert assessment of change.
Explicit abstention and defined escalation pathway	**Proposed**	Abstention rate and appropriateness; false alerts per patient-month; time from signal to clinician review.
Device and technique drift detection	**Proposed**	Whether drift is separable from biological change across phone and operating-system upgrades.

Note: Software development and implementation occurred from October 2025 through August 2026. Entries marked "Implemented" indicate functionality present in the current software build, while "Partially implemented" indicates functionality that is present but not yet complete relative to the intended architecture. These classifications do not constitute evidence of clinical validation or performance.

The "Baseline of One"

The conceptual differentiator of SoleIQ is the "baseline of one." Most image-classification systems are naturally organized around population labels: the image is compared with patterns learned from many other people, and the output describes which category it most resembles. Population learning remains necessary, but longitudinal monitoring creates a second reference point that may be equally important: the same patient at an earlier time.

A patient-specific baseline asks a different question. Instead of only asking, "Does this image resemble a concerning foot image?" the system can also ask, "What is different about this foot compared with this patient's own prior scans?" For visible features such as skin breakdown, callus pattern, color change, swelling, wound dimensions, or distribution of lesions, intra-patient

comparison may reduce the importance of normal variation between people. Skin tone, foot shape, scars, chronic deformities, and other stable characteristics can otherwise look unusual when judged only against a population template. A repeated personal baseline can help distinguish stable individual features from newly emerging change.

This concept is not a claim that longitudinal comparison is automatically more accurate. It creates its own technical problems. Images must be aligned despite differences in angle, distance, lighting, device, and pose. A patient's true baseline can also change after surgery, wound healing, edema, or other events. The system must therefore version the baseline rather than treating the first image ever captured as permanent ground truth.A practical implementation would maintain a rolling visual history.

Figure 4. Illustrative ulcer-healing progression. The serial wound images shown are synthetic, computer-generated illustrations presented solely to demonstrate a hypothetical pattern of wound-area reduction over time and are not derived from photographs of actual patients or clinical cases.

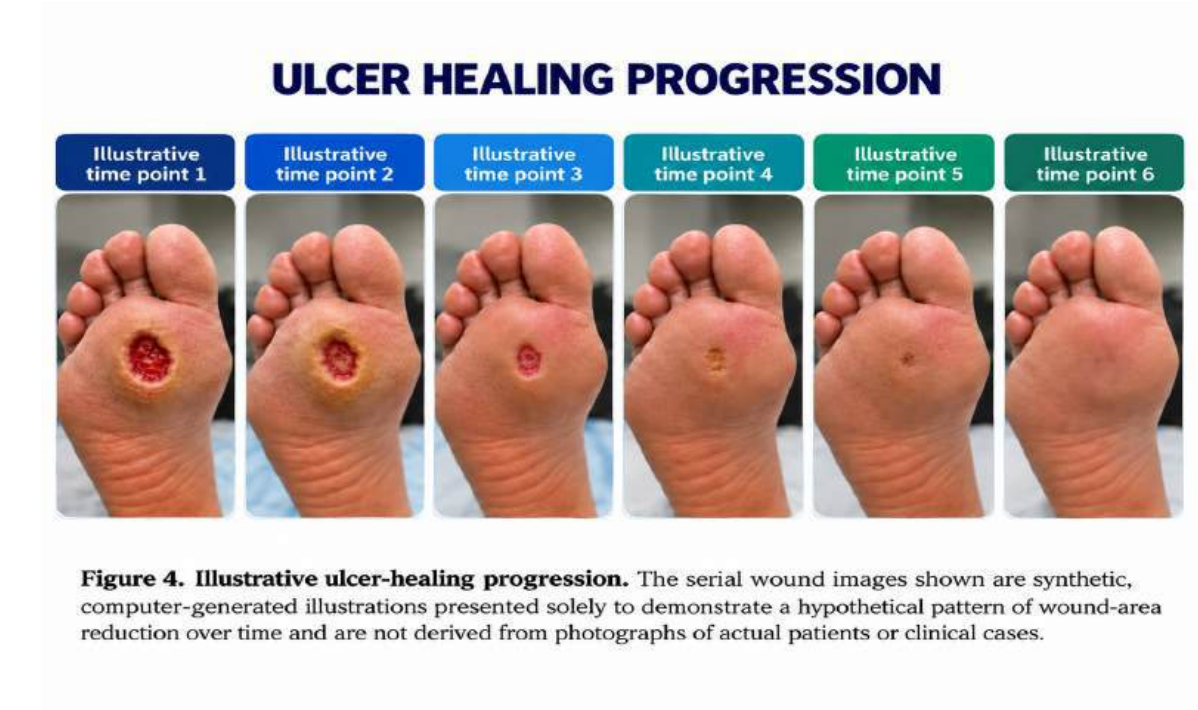

Figure 4. Illustrative ulcer-healing progression. The serial wound images shown are synthetic, computer-generated illustrations presented solely to demonstrate a hypothetical pattern of wound-area reduction over time and are not derived from photographs of actual patients or clinical cases.

New scans could be compared with both recent images and a longer-term reference set. The system could represent change as direction and magnitude rather than as a single disease label: stable, improving, worsening, or indeterminate. A clinician could then review the underlying images rather than relying on the label alone. This is especially important because the absence of a visible difference does not rule out vascular, neurologic, infectious, or biomechanical risk that a camera cannot measure.

A patient-specific baseline may already contain clinically important abnormalities; stability relative to that baseline must not be interpreted as evidence of a healthy or clinically safe foot. Concerning findings or symptoms should trigger the defined clinical review pathway regardless of longitudinal stability.

For the same reason, re-baselining must be additive rather than corrective. When a new reference is established after a device or operating-system change, or after a clinical event such as surgery, debridement, or wound healing, the prior visual history and any findings associated with it are retained and remain reviewable. Establishing a new reference point resets the comparison window; it does not resolve, close, or normalize an unresolved concerning finding, which should remain flagged until a clinician documents its resolution.

The baseline-of-one model also changes the research agenda. Validation should not only ask whether an algorithm agrees with expert labels on isolated images. It should ask whether the system can detect clinically meaningful change earlier or more consistently than unaided patient observation, whether the detected change is reproducible across repeated captures, and whether clinician review of the longitudinal record changes management. These are prospective questions and should be measured directly.

A Two-Layer AI Architecture

The proposed SoleIQ architecture separates visual analysis from narrative interpretation. This separation is partially implemented, and the comparison and change functions described below rest on components that are proposed rather than built (Table 2). The first layer is a dedicated computer-vision service that operates independently of the patient-facing application. Its role is deliberately constrained: find, measure, compare, and return structured features. The raw output of this vision layer is not intended to be shown directly to the patient.

Separating the vision layer serves two purposes. Technically, it allows detection, segmentation, measurement, or similarity components to be improved without rebuilding the patient-facing reasoning layer. Clinically, it creates a boundary between what the image model actually observed and how that observation is explained. This makes it easier to audit whether a narrative statement is supported by visual evidence.

The second layer is a secure server-side analysis service. It assembles the available images and patient context, calls the analysis model, and returns both the patient and clinician outputs in a structured pass. Model credentials are held server-side rather than on the device. The platform uses authenticated access, private storage, and patient-level data separation so that each patient's record is isolated at the data layer. The system is delivered as a web application that requires no dedicated hardware.

This architecture is useful only if the boundary between the two layers remains explicit. The vision service should report observable or measurable features with uncertainty. The narrative layer should not invent visual findings that the vision layer did not produce. Likewise, patient context should modify interpretation and urgency without being allowed to fabricate image evidence. The system should preserve the intermediate structured record so a clinician or auditor can reconstruct how the final output was generated.

The architecture must also define what happens when the primary narrative model declines to analyze an image. The safe default in that situation is abstention and human review, not automatic rerouting to a second model in order to obtain an answer. A refusal is information: it indicates that the system, as configured, was unwilling to interpret the input, and treating that

as a technical obstacle to route around inverts the safety logic of the rest of the design. A safety refusal is also never equivalent to a negative finding, and it should never be rendered to the patient as reassurance.

If a fallback model is retained for availability reasons, it should be governed as a distinct clinical component rather than as a substitute path. That means a separately locked and versioned model, its own validation against the same endpoints used for the primary model, its own audit trail recording every invocation and the reason for it, and a constraint that it cannot alter the intended clinical meaning of the output. Fallback use should be visible to the clinician on the record itself, monitored as a rate rather than an exception, and reviewed if it rises. Where the primary and fallback routes disagree, the more conservative result should stand and the case should be escalated. If neither route can produce a grounded result, the system should state plainly that the image could not be analyzed and direct the patient toward human review.

Clinical Framework Anchoring

SoleIQ is designed to anchor outputs to recognized clinical frameworks rather than informal color-coded labels. The IWGDF risk system provides a structured language for risk stratification and screening frequency, while the Wagner and University of Texas systems are commonly used to describe ulcer characteristics and severity. The purpose of incorporating these frameworks is not to claim that a smartphone image can reproduce a complete clinical classification. Some classification elements require examination findings, depth assessment, perfusion information, infection assessment, or other data that may not be available from a photograph alone.

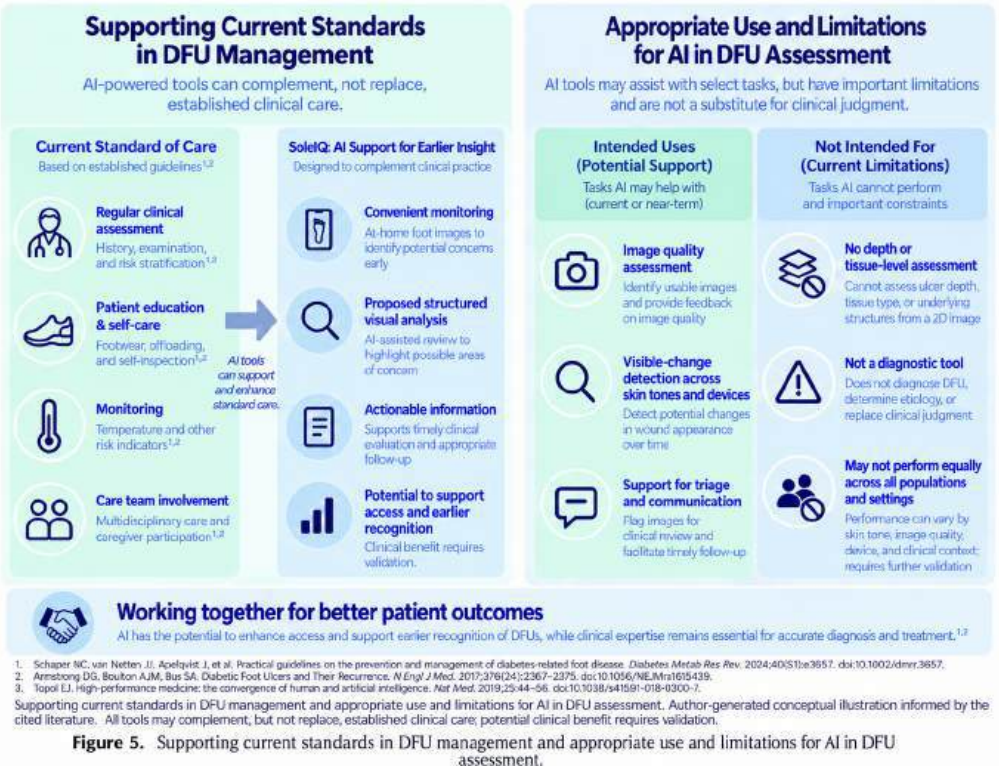

1. Schaper NC, van Netten JJ, Apelqvist J, et al. Practical guidelines on the prevention and management of diabetes-related foot disease. *Diabetes Metab Res Rev.* 2024;40(S1):e3657. doi:10.1002/dmrr.3657.
2. Armstrong DG, Boulton AJM, Bus SA. Diabetic Foot Ulcers and Their Recurrence. *N Engl J Med.* 2017;376(24):2367–2375. doi:10.1056/NEJMra1615439.
3. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med.* 2019;25:44–56. doi:10.1038/s41591-018-0300-7.

Supporting current standards in DFU management and appropriate use and limitations for AI in DFU assessment. Author-generated conceptual illustration informed by the cited literature. All tools may complement, but not replace, established clinical care; potential clinical benefit requires validation.

Figure 5. Supporting current standards in DFU management and appropriate use and limitations for AI in DFU assessment.

For this reason, the system should not generate a Wagner or University of Texas grade from a photograph alone. Depth, probing to bone, perfusion, infection status, and protective sensation are not recoverable from a visible-light image, and a grade derived without them would carry an unearned appearance of clinical completeness even if labeled an estimate. The safer design presents the visual evidence and leaves classification to the clinician, who can assign a grade after integrating the examination findings the photograph cannot supply. Recognized frameworks are therefore used here as shared vocabulary for reporting, risk context, and escalation thresholds, not as an output the model produces. The clinician snapshot should distinguish observed visual features from patient-reported context and from inferred interpretation. For example, "visible skin break with surrounding erythema" is different from "possible infection," and both are different from a confirmed diagnosis. Clear separation of observation and inference makes the output easier to review and safer to challenge.

The ADA recommends comprehensive foot evaluation at least annually for people with diabetes, more frequent inspection for people at higher risk, and urgent specialist referral for open ulceration or unexplained swelling, erythema, or increased skin temperature (American Diabetes Association Professional Practice Committee for Diabetes). A smartphone system should therefore be designed to fit around those clinical pathways. Its job is to strengthen surveillance and communication between visits, not to create a parallel standard of care.

The physician-first distribution model is important in this context. Enrollment can connect the patient to an existing clinical relationship, define escalation instructions, and create a destination for concerning scans. A monitoring system without a response pathway risks identifying change without ensuring that anyone can act on it. The technical product and the clinical workflow should therefore be validated together.

Safety-First Design Requirements

A safety-first monitoring system should be designed around failure rather than average performance. Several classes of failure are foreseeable.

The first is acquisition failure. A low-quality image can conceal the very feature the model is asked to evaluate. Quality gating should therefore assess at minimum blur, exposure, framing, and whether required views are present. The validation endpoint is not simply whether the quality model detects bad photographs; it is whether rejection and recapture lead to more interpretable downstream images without causing users to abandon the process.

The second is inference failure. An image may be technically clear but still ambiguous. The system should be allowed to return "indeterminate" or request human review. This is consistent with broader AI risk-management principles that emphasize validity, reliability, safety, transparency, and ongoing evaluation across the lifecycle (Tabassi). Coverage should not be maximized at the expense of false reassurance.

The third is communication failure. A clinically cautious model can still become unsafe if the language presented to a patient sounds definitive. Patient-facing outputs should avoid diagnostic certainty, clearly distinguish monitoring from medical evaluation, and use escalation

language that is understandable without specialized training. When the system cannot determine whether a change is significant, the correct output may be uncertainty plus a request for professional review.

The fourth is workflow failure. A high-risk flag has little value if it does not reach the clinician, reaches the wrong clinician, or arrives without enough information to act. The escalation pathway should therefore be specified in the protocol before deployment rather than left to local improvisation. That specification should name the reviewer responsible for each alert tier, define separate response windows for urgent and routine findings, require explicit acknowledgment by the named reviewer rather than passive confirmation that a message was delivered, and define a backup route that activates when an alert is not acknowledged within its window, including who is contacted next and what the patient is told in the interim. The appropriate windows are for the implementing service to set against its own staffing and case mix; the requirement is that they be defined in advance, documented, and measured rather than assumed. Validation must therefore include alert delivery, review time, acknowledgment, and escalation completion. Technical accuracy and operational reliability are part of the same safety system.

The fifth is distributional failure. A model may work differently across skin tones, camera types, lighting environments, age groups, deformities, or wound locations. The platform should report subgroup performance and should not treat a pooled accuracy statistic as sufficient evidence of safety. This is particularly important for a product whose stated mission includes underserved and rural populations.

The sixth is temporal failure. Longitudinal systems can create false stability if capture technique changes or if a new device produces different color and image-processing characteristics. The system should detect major changes in device or capture conditions and avoid interpreting them as biological change without confirmation.

A Staged Validation Framework

A credible validation program for SoleIQ should proceed in stages, with each stage answering a narrower question before clinical claims expand.

Stage 1 is technical verification. The goal is to establish that the software performs its intended mechanical tasks consistently. Tests should measure capture completion, image-quality rejection, view coverage, secure upload, patient-level data isolation, deterministic handling of missing inputs, and reproducibility of structured outputs. Vision components should be evaluated with locked test sets and clearly defined metrics. Narrative components should be tested for factual faithfulness to the structured visual features they receive.

Stage 2 is retrospective clinical validation. A de-identified image set with expert annotations can be used to test whether observed visual features, image-derived measurements where supported, and longitudinal change assessments agree with clinician assessment. Importantly, the benchmark should contain repeated images from the same individuals if the

baseline-of-one concept is being evaluated. The unit of analysis should not be only the individual image; it should also include change between time points.

Stage 3 is prospective observational validation. Patients would use the system in realistic home conditions while clinicians independently evaluate the resulting images or conduct standard assessments. Primary endpoints should include successful capture rate, proportion of scans requiring recapture, clinician agreement on whether change is present, sensitivity for changes clinicians believe require follow-up, false-reassurance rate, and time from concerning scan to clinician review. Patient usability and clinician workload should be treated as safety outcomes, not secondary marketing measures.

Stage 4 is clinical-utility evaluation. Once the system can reliably capture and communicate change, the next question is whether it improves care. Relevant outcomes could include time to clinical contact after a concerning change, proportion of escalations judged appropriate, unplanned urgent visits, ulcer-related events, and patient adherence to follow-up. Such studies would need prospective protocols, appropriate oversight, and predefined statistical analysis.

Stage 5 is deployment monitoring. Performance can drift after release because phone cameras, operating systems, patient populations, and models change. Versioned models, audit logs, subgroup dashboards, and post-deployment error review should therefore be built into the product. NIST's AI Risk Management Framework emphasizes risk management across design, deployment, measurement, and ongoing management rather than as a one-time pre-release exercise (Tabassi).

Several methodological requirements apply across the retrospective and prospective stages and should be stated in advance rather than settled during analysis. Training, validation, and test sets must be separated at the patient level, not the image level, because repeated images of the same foot otherwise leak across splits and inflate apparent performance in exactly the longitudinal setting this framework targets. Because each patient contributes repeated images and repeated scan sessions, analyses and uncertainty estimates should account for within-patient correlation, for example through mixed-effects models or patient-level cluster-robust standard errors, rather than treating individual images or sessions as independent observations. The model should be locked before evaluation, with the version recorded. Reference assessments should be made by blinded expert reviewers with inter-rater agreement reported, since a single unblinded annotator cannot establish whether a detected change is real. External validation across multiple sites, clinical settings, and patient populations should follow single-site work rather than replace it.

The endpoints themselves should extend beyond discrimination. Calibration should be reported, not only sensitivity and specificity, because a monitoring system communicates degrees of concern rather than binary labels. Abstention and indeterminate outputs should be measured as a performance characteristic in their own right: how often the system declines, whether it declines on the cases it should, and what happens to those cases afterward. Operational burden should be quantified through false alerts per patient-month, time from signal to clinician review,

clinician workload, and evidence of alert fatigue. Patient-side behavior should be measured through adherence, missing captures, and abandonment following recapture requests, since a system that improves image quality by exhausting its users has not improved monitoring. Performance should be reported by subgroup, including skin tone, phone model and camera class, age, mobility, and whether capture required caregiver assistance.

Most importantly, the studies should pre-specify a direct comparison of population-only, baseline-only, and hybrid models evaluated on the same prospective dataset with the same endpoints. Without that comparison, the baseline of one remains a design preference. With it, the central claim of this framework becomes a testable hypothesis that can be supported or rejected on evidence.

The most important principle across all stages is that the chosen metric must match the intended use. If the purpose is monitoring and triage, then conventional multiclass accuracy alone is insufficient. A system can achieve respectable average accuracy while still failing at the one task that matters most: recognizing a concerning change and avoiding false reassurance. Sensitivity for escalation, calibration, abstention behavior, and human follow-through are therefore central endpoints.

Table 3. Proposed validation sequence

Stage	Primary question	Example endpoints
Technical verification	Does the software behave as designed?	Capture completion, quality rejection, isolation, reproducibility
Retrospective validation	Does image analysis agree with expert review?	Feature agreement, measurement error, change detection
Prospective observation	Does the system work under home conditions?	Usable capture, recapture, escalation sensitivity, false reassurance
Clinical utility	Does information change care in a useful way?	Time to review, appropriate escalation, follow-through
Deployment monitoring	Does performance remain reliable after release?	Drift, subgroup performance, model-version error review

Equity and Real-World Deployment

A software-only smartphone approach has an important theoretical advantage: it does not require the patient to purchase a dedicated sensor, mat, or wearable. That can reduce one barrier to scale and makes deployment possible wherever a compatible phone and connectivity are available. A progressive web application can also simplify distribution because users do not necessarily need a specialized hardware setup.

However, software-only should not be confused with access-neutral. Patients may have older cameras, limited data plans, unreliable broadband, low digital literacy, visual impairment, limited mobility, or difficulty photographing the plantar surface. Evaluations of ulcer-detection models running on low-end handsets indicate that device class is a performance variable rather than a deployment detail (Cassidy et al.). Some may require a caregiver. Others may be unable to produce consistent images despite repeated prompting. These limitations should be measured rather than assumed away.

An equity-oriented validation plan should therefore stratify performance by device class, skin tone, age, mobility, caregiver assistance, and connectivity environment. The product should log when a scan fails because of image quality or incomplete views, since systematic failure to capture a subgroup is itself a performance disparity. Interface testing should include large text, clear prompts, minimal steps, and language accessibility.

The physician-first distribution model can help with inclusion if clinical teams identify patients who would benefit and provide onboarding support. It can also create a new access problem if participating practices are concentrated in well-resourced settings. The goal should be to test the platform in community and rural environments early enough that deployment constraints shape the product rather than appear only after development is complete.

Privacy, Security, and AI Governance

Longitudinal foot monitoring produces sensitive health information and repeated images linked to a patient's identity. Privacy architecture is therefore part of clinical safety. SoleIQ's current architecture enforces authenticated access, private storage, and patient-level data separation. API credentials are kept in server-side functions rather than distributed to the client. These are useful design choices, but they are only the beginning of governance.

The platform should minimize the data collected, document why each field is necessary, define retention periods, and record access to patient data. Model-development datasets should be separated from production records and should not be assumed available for training merely because they were collected during care. Consent for monitoring and consent for secondary model development are distinct purposes.

Auditability is also essential. Each analysis should retain the model version, input version, quality status, structured visual features, patient context used, final output, and whether a fallback model was invoked. If a result is later questioned, the team should be able to reconstruct the pathway that produced it.

Responsible AI guidance also supports continuous evaluation for reliability, security, transparency, explainability, privacy, and harmful bias (Tabassi). For SoleIQ, explainability should be practical rather than decorative. A clinician should be able to see what visible features changed, compare current and prior images, and determine whether the narrative output is grounded in those changes.

Figure 6. Strategies to address real-world deployment restrictions in AI models.

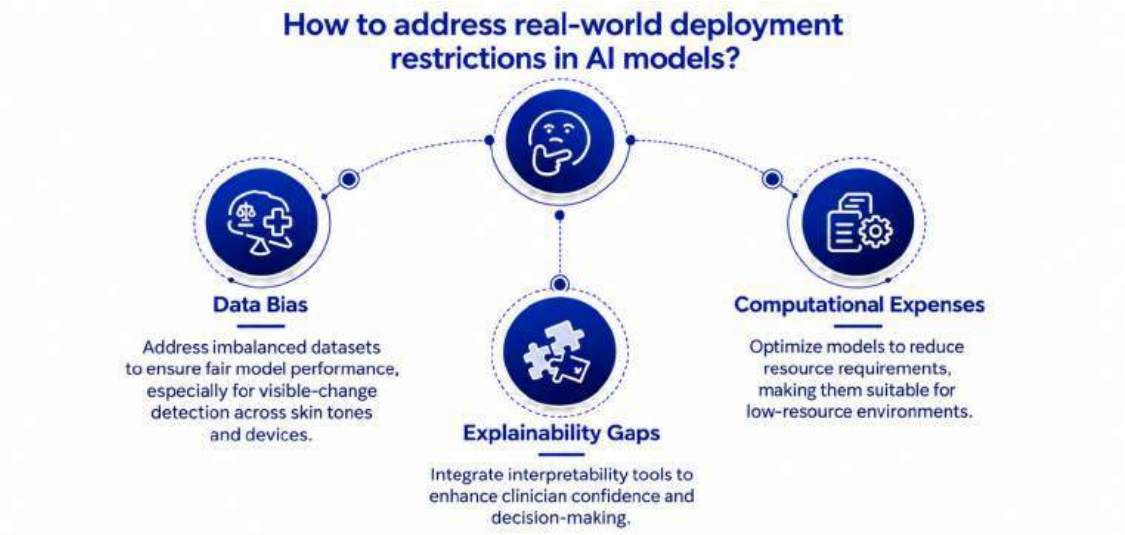

Figure 6. Strategies to address real-world deployment restrictions in AI models.
Key considerations include addressing data bias, bridging explainability gaps, and reducing computational expenses to facilitate broader, real-world implementation.

Regulatory Positioning

Describing a system as wellness monitoring does not determine its regulatory status. The functions proposed here include detection of change in the diabetic foot, triage support, disease-related outputs, and guidance on next steps. Regulatory characterization requires formal assessment of the intended use, claims, users, and functionality, and each of those can shift as a product develops. FDA guidance distinguishes general-wellness software from functions related to diagnosis, cure, mitigation, prevention, or treatment of disease, and separately addresses clinical decision-support software and patient- or caregiver-facing functions (U.S. Food and Drug Administration, “General Wellness”; U.S. Food and Drug Administration, “Clinical Decision Support Software”).

For that reason, regulatory strategy should evolve alongside validation. If future versions make stronger disease-detection, severity, or treatment claims, the evidence and oversight requirements may change. The safest research posture is to define the current intended use precisely, validate against that use, and avoid allowing marketing language to advance faster than the underlying evidence. Accordingly, the system described in this paper should be understood as investigational. It is not presented as a cleared or approved device, no regulatory determination has been made, and formal characterization will require review as claims, users, and validation evidence evolve.

The same principle applies clinically. A tool that gives next-step guidance should have clear emergency and escalation boundaries. Open wounds, rapidly spreading redness, marked swelling, systemic illness, or other concerning findings may require in-person assessment regardless of what the model says. Human clinical judgment remains the final authority.

Discussion

The central argument of this paper is that smartphone artificial intelligence for diabetic foot care should be evaluated as a longitudinal observation system rather than merely as an image classifier. That shift changes both the technical architecture and the definition of success.

First, it moves quality control upstream. If a photograph is not adequate, the system should identify that before analysis rather than attach a confident interpretation to uncertain input. Second, it makes the patient's own history a core source of signal. Population models remain necessary for recognizing common patterns, but the baseline-of-one approach emphasizes within-person change as a separate and potentially valuable dimension. Third, it separates observation from explanation. A vision model extracts structured features; a narrative model translates those features and the patient's context into outputs designed for different audiences. Fourth, it makes the clinician workflow part of the product rather than an afterthought.

This design is compatible with the broader clinical logic of diabetic foot prevention. Guidelines emphasize risk stratification, repeated surveillance, patient education, and timely referral, while home-monitoring studies show that patient-generated measurements can become useful when paired with an actionable response (International Working Group on the Diabetic Foot, “Prevention Guideline”; Lavery et al.; Bus et al.). SoleIQ extends that logic to a software-only visual monitoring pathway.

The approach also has important limits. A camera cannot measure protective sensation, pedal pulses, tissue perfusion, bone involvement, systemic infection, or many biomechanical forces that contribute to ulcer formation. Color and morphology can be distorted by illumination and camera processing. A longitudinal image history can show visible change while still missing invisible risk. The platform must therefore resist the temptation to turn a convenient data source into a complete representation of disease.

Visible-light photography is also only one channel among several that may eventually contribute to remote foot monitoring. Work on biochemical and biophysical sensing, including instrumented dressings that report on the wound environment directly, suggests a multimodal future in which images supply morphology and change while other modalities supply information the camera cannot reach (Mirani et al.). These approaches are more likely to complement one another than to compete. The strongest use case for the present system may therefore be narrower: helping a patient capture consistent images, notice change, and send a structured snapshot to a clinician sooner than would otherwise occur. That value proposition is clinically meaningful even if the model never becomes an autonomous diagnostic engine; limiting autonomy may make the system more useful by preserving the clinician as the decision-maker while reducing the information gap between visits.

The “baseline of one” is especially promising as a research hypothesis because it can be tested without claiming that every patient should be compared only with themselves. A hybrid model can use population training to understand general diabetic foot features while using patient history to estimate change. Future studies can compare population-only, longitudinal-only, and hybrid approaches on the same prospective dataset. The outcome should

be whether the model identifies clinically relevant change, not whether one architecture wins a benchmark by a small margin.

SoleIQ's two-layer architecture also creates an opportunity for more transparent evaluation. If the visual layer and narrative layer are tested separately, investigators can distinguish a perception error from a reasoning or communication error. This matters because the corrective action differs. A missed lesion may require better vision training or capture quality. An exaggerated patient warning may require narrative constraints. A correct high-risk flag that is never reviewed may require workflow redesign. Treating all of these as one "model error" hides the real failure mechanism.

Ultimately, a safe monitoring system is an interaction among patient behavior, camera quality, computer vision, clinical context, language generation, clinician workflow, and follow-up. No single accuracy number can represent the reliability of that chain. The research program should therefore measure the chain.

Limitations

This paper has several limitations. It is a conceptual design analysis, not a systematic review or empirical validation study. The SoleIQ architecture described here has not been shown in this manuscript to reduce ulceration, hospitalization, amputation, or other clinical outcomes. No performance metrics are reported, and none should be inferred from the proposed design.

The baseline-of-one concept is a hypothesis. It may improve sensitivity to patient-specific change, but it may also be vulnerable to inconsistent image capture, changing devices, baseline drift, or false interpretation of benign variation. Prospective comparison is required.

The smartphone camera is intrinsically limited. Important aspects of diabetic foot risk require physical examination, vascular and neurologic assessment, laboratory or imaging studies, pressure measurement, or other data that cannot be recovered from a conventional photograph. Accordingly, the proposed system cannot substitute for comprehensive foot care.

The technical stack is also not equivalent to a validated medical system. Security features such as private storage and enforced access controls reduce risk but do not by themselves establish regulatory compliance, cybersecurity adequacy, or clinical safety. Those requirements depend on implementation, documentation, testing, intended use, and deployment environment.

Finally, this paper focuses on system design rather than comparative effectiveness. The relevant next step is not to make stronger claims from the same conceptual framework, but to test the framework prospectively with patients and clinicians under realistic conditions.

Conclusion

Diabetic foot prevention depends on noticing risk before a small change becomes a major complication. Clinical examinations remain essential, but they occur intermittently. A smartphone can help fill part of the interval if it is used not as a replacement clinician but as a structured observation tool.

SoleIQ Health provides a model for this approach: baseline risk context, guided multi-angle capture, image-quality gating, patient symptom mapping, a dedicated vision layer, a

separate narrative layer, dual patient-and-clinician outputs, and a longitudinal timeline. The baseline-of-one concept places patient-specific change at the center of monitoring while still allowing population-trained models to contribute general visual knowledge.

The research challenge is now validation. The most important questions are whether patients can capture usable images consistently, whether the system detects meaningful change without false reassurance, whether outputs remain grounded and interpretable, whether performance is equitable across patients and devices, and whether concerning scans lead to timely clinical action. If those questions can be answered prospectively, smartphone-based AI may become a practical bridge between self-monitoring and professional diabetic foot care.

Conflict of Interest

Eshaan Naik is the founder of SoleIQ Health. David G. Armstrong serves in a clinical advisory relationship with SoleIQ Health. The system described in this article is presented as a proposed monitoring and research framework. No clinical performance claims are made.

Ethics Statement

No human participants were enrolled and no patient data were collected or analyzed for this conceptual paper. Institutional review board approval was therefore not applicable to the work reported here.

Works Cited

American Diabetes Association Professional Practice Committee for Diabetes. "12. Retinopathy, Neuropathy, and Foot Care: Standards of Care in Diabetes—2026." Diabetes Care, vol. 49, suppl. 1, 2026, pp. S261-S276. https://doi.org/10.2337/dc26-S012.

Armstrong, David G., Andrew J. M. Boulton, and Sicco A. Bus. "Diabetic Foot Ulcers and Their Recurrence." New England Journal of Medicine, vol. 376, no. 24, 2017, pp. 2367-2375. https://doi.org/10.1056/NEJMra1615439.

Armstrong, David G., Lawrence A. Lavery, and Lawrence B. Harkless. "Who Is at Risk for Diabetic Foot Ulceration?" Clinics in Podiatric Medicine and Surgery, vol. 15, no. 1, 1998, pp. 11-19. https://doi.org/10.1016/S0891-8422(23)01025-X.

Bus, Sicco A., et al. "Effectiveness of At-Home Skin Temperature Monitoring in Reducing the Incidence of Foot Ulcer Recurrence in People with Diabetes: A Multicenter Randomized Controlled Trial (DIATEMP)." BMJ Open Diabetes Research & Care, vol. 9, no. 1, 2021, e002392. https://doi.org/10.1136/bmjdrc-2021-002392.

Cassidy, B., et al. "Artificial Intelligence for Automated Detection of Diabetic Foot Ulcers: A Real-World Proof-of-Concept Clinical Evaluation." Diabetes Research and Clinical Practice, vol. 205, 2023, 110951. https://doi.org/10.1016/j.diabres.2023.110951.

Chan, K. S., et al. "Clinical Validation of an Artificial Intelligence-Enabled Wound Imaging Mobile Application in Diabetic Foot Ulcers." International Wound Journal, vol. 19, no. 1, 2022, pp. 114-124. https://doi.org/10.1111/iwj.13603.

Golledge, J., M. Fernando, P. A. Lazzarini, B. Najafi, and D. G. Armstrong. "The Potential Role of Sensors, Wearables and Telehealth in the Remote Management of Diabetes-Related Foot Disease." Sensors, vol. 20, no. 16, 2020, 4527. https://doi.org/10.3390/s20164527.

International Working Group on the Diabetic Foot. "Guidelines on the Prevention of Foot Ulcers in Persons with Diabetes: IWGDF 2023 Update." IWGDF Guidelines, 2023, https://iwgdfguidelines.org/prevention-guideline-2023/. Accessed 20 Aug. 2026.

International Working Group on the Diabetic Foot. "Practical Guidelines on the Prevention and Management of Diabetes-Related Foot Disease: IWGDF 2023 Update." IWGDF Guidelines, 2023, https://iwgdfguidelines.org/practical-guidelines-2023/. Accessed 20 Aug. 2026.

International Working Group on the Diabetic Foot. "Guideline on Classification of Diabetes-Related Foot Ulcers: IWGDF 2023 Update." IWGDF Guidelines, 2023, https://iwgdfguidelines.org/classification-2023/. Accessed 20 Aug. 2026.

Lavery, Lawrence A., et al. "Preventing Diabetic Foot Ulcer Recurrence in High-Risk Patients: Use of Temperature Monitoring as a Self-Assessment Tool." Diabetes Care, vol. 30, no. 1, 2007, pp. 14-20. https://doi.org/10.2337/dc06-1600.

Lo, Z. J., et al. "A Feasibility Study on the Efficacy of a Patient-Owned Wound Surveillance System for Diabetic Foot Ulcer Care (ePOWS Study)." Digital Health, vol. 9, 2023, 20552076231205747. https://doi.org/10.1177/20552076231205747.

Mirani, B., Z. Hadisi, E. Pagan, S. M. H. Dabiri, A. van Rijt, L. Almutairi, I. Noshadi, D. G. Armstrong, and M. Akbari. “Smart Dual-Sensor Wound Dressing for Monitoring Cutaneous Wounds.” Advanced Healthcare Materials, vol. 12, no. 18, 2023, e2203233. https://doi.org/10.1002/adhm.202203233.

Ploderer, B., D. Clark, R. Brown, J. Harman, P. A. Lazzarini, and J. J. van Netten. “Self-Monitoring Diabetes-Related Foot Ulcers with the MyFootCare App: A Mixed Methods Study.” Sensors, vol. 23, no. 5, 2023, 2547. https://doi.org/10.3390/s23052547.

Satheesh, S. S., A. Rayampalli, A. G. Prabhune, and V. R. Sri Hari. “Development and Validation of a Predictive AI Framework for Diabetic Foot Ulcer Monitoring and Severity Assessment: A Step towards Self-Monitoring and Primary Care Integration.” Healthcare Informatics Research, vol. 32, no. 1, 2026, pp. 69-76. https://doi.org/10.4258/hir.2026.32.1.69.

Swerdlow, M., L. Shin, K. D’Huyvetter, W. J. Mack, and D. G. Armstrong. “Initial Clinical Experience with a Simple, Home System for Early Detection and Monitoring of Diabetic Foot Ulcers: The Foot Selfie.” Journal of Diabetes Science and Technology, vol. 17, no. 1, 2023, pp. 79-88. https://doi.org/10.1177/19322968211053348.

Tabassi, Elham. Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology, 2023. NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1.

U.S. Food and Drug Administration. “Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff.” U.S. Food and Drug Administration, Jan. 2026, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software. Accessed 20 Aug. 2026.

U.S. Food and Drug Administration. “General Wellness: Policy for Low Risk Devices.” U.S. Food and Drug Administration, Jan. 2026, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/general-wellness-policy-low-risk-devices. Accessed 20 Aug. 2026.

Sex Specific Implications of Insular Cortex Degradation During Alzheimer's Disease on Interoceptive Accuracy in Decision-Making By Deepika Srilalita Manikonda

Abstract

Alzheimer's disease is commonly associated with memory loss; however, it also has a profound impact on our intuition, affecting regions responsible for internal bodily signals, regulating emotions, and guiding decision-making, such as the insular cortex. Neurodegenerative diseases, particularly Alzheimer's, target these salience networks, leading to severe disruptions in signal processing and choices. The research question this review paper aims to answer is: Do sex differences in insular cortex involvement in Alzheimer's disease contribute to differences in interoceptive accuracy and intuitive, somatic marker-based decision-making between women and men? Previous investigations primarily utilized fMRI (functional magnetic resonance imaging) during autonomic challenges and risk-taking tasks to map insular cortex activation, including lesions within the region to establish the necessity of the anterior insula in bodily awareness and effects on decision-making processes once disrupted. This review paper integrates structural neuroimaging, individual brain atlases, neuropsychological and behavioral paradigms, and neuropathology to evaluate sex differences in interoception and intuitive decision-making decline in Alzheimer's disease. It is expected to show that female-skewed Alzheimer's pathology in tau accumulation in the anterior insula and sex-based differences in the structure of the insular cortex contribute to differences in interoceptive accuracy and somatic marker-based decision-making. By integrating research on Alzheimer's pathology, interoception, and biological sex differences in brain anatomy, this review argues that the insular cortex is an important link underrepresented in current Alzheimer's research. This review will contribute to the field by providing a sex specific framework that links Alzheimer's pathology directly to everyday decision-making deficits, which will bridge an important gap between neurodegeneration and interoception to see how intuitive processes are affected. Resolving this gap will be crucial for developing diagnostic markers for behavioral symptoms in Alzheimer's patients by understanding the sex based differences more closely.

Keywords Insular cortex, interoceptive accuracy, somatic marker hypothesis, sex differences, Alzheimer's disease

Introduction

Our brains consistently use internal bodily responses to guide everyday decisions, yet the effects of neurodegenerative diseases like Alzheimer's on disrupting those functions differently in women and men remain largely unexplored. The insular cortex, a fold buried in the cerebral cortex between the frontal and temporal lobes, functions as the main region for interoception – the continuous sensing of internal physiological signals such as heart rate, respiration, hunger, thirst, and muscle tension (Wang et al.). The cortex is divided into two divisions: posterior (back) and anterior (front). The posterior insula is the first part to receive raw somatosensory and homeostatic signals from the thalamus, responsible for processing homeostatic information such

as pain, temperature, itch, local oxygen status, and sensual touch. These signals are then processed forward through the middle insula and into the anterior insula (Wang et al.). The anterior insula works alongside the anterior cingulate cortex in the salience network of the brain to translate those signals into subjective emotional awareness and somatic markers. These somatic markers are our physiological "gut feelings" that guide our risk evaluations and decision-making (Damasio,1994,1996; Castelhano et al.; Droutman et al.). Lesion studies in the anterior insula confirm that this region is necessary for interoceptive accuracy (Wang et al.).

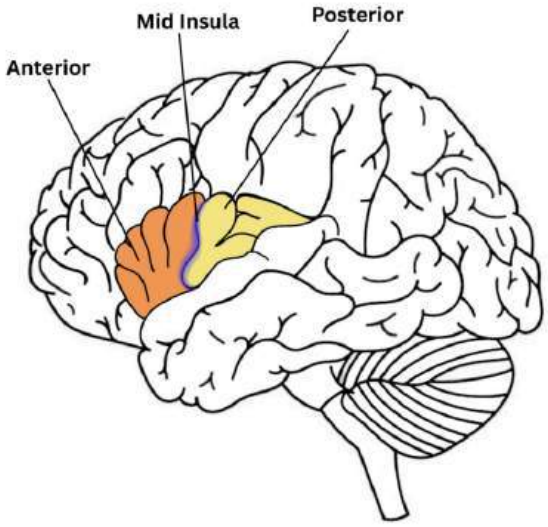

Fig 1: Visual of the insular cortex divided into 3 subregions- anterior, middle, and posterior insula

Importantly, the insular cortex also exhibits some baseline sex related functional differences, including structural disparities. Under physiological stress challenges, women and men showcase different patterns of insular cortex activation, autonomic regulation, and risk-taking behaviors (Lighthall et al.; Macey et al.). Structural and functional differences are also shown in Alzheimer's disease that demonstrate significant sex disparities: women carrying the APOE-4 allele experience rapid tau accumulation and faster disease progression (Laws et al.). While recent and past Alzheimer's models emphasize the hippocampus and entorhinal degeneration (Wenk), studies reveal that Alzheimer's tau pathology also selectively targets the anterior insula (Bonthius et al.).

This paper will investigate the effects of sex differences in the insular cortex during Alzheimer's disease on interoceptive accuracy in decision-making between men and women. To maintain scope, the review will focus specifically on adult human populations with mild cognitive impairment and mild to moderate Alzheimer's. It excludes familial genetic factors causing the disease, individuals with severe end-stage dementia, and psychiatric disorders.

Despite the evidence that Alzheimer's involves non-memory symptoms, including changed perceptions of pain, lack of risk-evaluative behaviors, and apathy, current dementia research focuses on memory networks, especially the entorhinal and hippocampal regions (Moon

et al.; Wenk). While some studies have:

- Mapped insular cortex subregions and their role in decision-making phases (Droutman et al.)
- Established the anterior insula's pivotal role in interoception (Wang et al.) ● Identified the baseline sex differences in insula cortex stress responses (Lighthall et al.; Macey et al.)
- Detailed tau vulnerability in the insular regions within Alzheimer's (Bonthius et al.)

These four areas have remained isolated, creating a gap in the literature. This framework, which connects all four areas into a sex-differentiated model, has not yet been directly evaluated or tested in the existing literature and is proposed for the first time. If Alzheimer's pathology selectively attacks the anterior insular areas (Bonthius et al.) that create pre-decision somatic markers (Castelhano et al.), individuals are likely to experience a comparable degradation in their intuitive abilities and harm-avoidant decision-making. Given that women exhibit a higher baseline reliance on their anterior insula (Iezzi et al., 2024), including an increase in Alzheimer's tau accumulation (Laws et al.), the breakdown of the insular cortex during the progression of this disease may manifest differently in terms of severity between biological sexes. This review argues that Alzheimer's pathology selectively disrupts the anterior insular cortex and that baseline sex differences in the insular cortex cause this degeneration to impair interoceptive accuracy and decision-making in fundamentally different ways between women and men: disrupting contextual intuition in women while hindering autonomic risk guidance in men. It is hypothesized that an acceleration in tau tangle accumulation in the anterior insula of females with Alzheimer's interacting with the baseline female insular region (with right anterior insula dominant during sympathetic stress and heavy reliance on anterior plasticity), results in a qualitatively distinct losses in interoceptive accuracy and somatic marker signals during risk evaluations, including a decrease in advantageous behaviors in decision making when compared to men at equal stages of the disease.

Methods: A systematic literature review was conducted across PubMed, MEDLINE, and Google Scholar to identify peer-reviewed articles that investigated insular cortex anatomy, interception, sexual differences in autonomic control, and Alzheimer's disease pathology. Keywords and search items include: *insular cortex, anterior insula, posterior insula, interocceptive accuracy, somatic marker hypothesis, sex differences, autonomic regulation, tau tangles, Alzheimer's disease, and decision making.* Studies were reviewed if they provided empirical data (fMRI, lesions, slice electrophysiology, ot volumetric MRI) or theories addressing the insular cortex connection or the progression of dementia. Clinical studies focused on mild cognitive impairment (MCI) and mild to moderate Alzheimer's, excluding genetically inherited forms of Alzheimer's and severe psychiatric disorders.

To address this gap and evaluate this hypothesis, this review will go through these independently studied subparts to bring them together into a unified model:

1. Insular anatomy and decision-making

- Establishing the physiological pathway mapping the posterior-to-anterior insula in signals

2. Sex differences in insula anatomy
    - Explores the differences in male and female insular cortex processing and behavioral strategies before the onset of neurodegenerative diseases to understand the baseline differences.
3. Alzheimer's pathology in the insular cortex
    - Examining how and why Alzheimer's specifically affects insular cortex tissue structures, interoceptive errors, and memory consolidation.

Insular Anatomy and Decision Making

The insular cortex is a region of the brain located deep within the cerebral cortex. Surrounded by multiple arteries and veins. It is folded inside the lateral sulcus (Sylvian fissure) between the frontal and temporal lobes (Kortz and Lillehei) in both hemispheres of the brain (right and left). The insula works with other brain regions to transform internal bodily sensations into conscious awareness. It provides emotional context for physiological experiences. The insular cortex is divided into 3 sections: anterior, middle, and posterior, with variability in the number of gyri per section. The anterior insula consists of three main *short* gyri, and the posterior insula is made of two main *long* gyri. These two sections are separated by the central sulcus (middle).

The insula acts as the main region where interoception occurs. Interoception is the ability to feel, understand, and react to signals from inside the body, such as homeostatic conditions like hunger, thirst, a fast heartbeat, or muscle tension, in response to external stimuli, which is an integral part of intuitive processes. Visceral signals travel from the through the spinal and vagus nerve, reaching the brainstem (specifically the nucleus of the solitary tract, NTS). Those signals then pass through the parabrachial nucleus and thalamus to the posterior insula. Then they move towards the mid insula, where contextual information is processed, and finally to the anterior insula, where the physical signals are processed into conscious feelings. (Craig; Wang et al.; Zhang et al.). The anterior insula provides awareness of feelings as a result of those internal bodily signals (Namkung et al.). Here, the anterior insula is a significant part of the insula, working with more regions of the brain to further interpret these signals. The right anterior insula, more precisely, has been found to have a high density of bipolar spindle neurons existing in humans and apes. With their large-diameter axons, they work to transmit information over long distances across brain regions (Gu et al.). These cells are utilized to process complex emotions such as empathy, guilt, trust, and intuition. Neuroimaging studies have shown that the anterior insula has consistent activity during decision-making tasks involving uncertainty or risk (Zhang et al.). When individuals focused on internal bodily cues like breathing and heart rate, the anterior insular cortex showed a spike of positive brain activity in the anterior insular cortex, as seen in fMRI results, with both the right ($T=6.98$, $Z=5.68$) and the left side ($T=5.50$, $Z=4.76$), with a particularly pronounced effect observed in the right hemisphere (Wang et al.). It is also

important to note that the right anterior insula showed statistically significant activity when processing delayed external signals compared to normally timed external signals (Wang et al.). Participants whose left and right anterior insula reacted more strongly to the mismatch of the computer delay scored higher on interoceptive accuracy, with the data showing a positive correlation between the interaction effect of the AIC and relative interoceptive accuracy (Wang et al.). Although there was variation between participants, it demonstrates that people differ in their interoceptive accuracy, where individuals who have higher sensitivity in regions like the anterior insula cortex often pick up on bodily cues faster, which allows them to use subconscious bodily "hints" to make quick, intuitive decisions.

However, the insular cortex does not function alone for interoceptive processes. The cortex works alongside many different regions of the brain, including the limbic system's amygdala, hypothalamus, anterior cingulate cortex, the olfactory cortex, as well as sensorimotor regions like the parietal lobe. These three sections of the insula are individually connected to these regions differently, as found by in vivo diffusion-weighted imaging studies. For example, the anterior insula has connections to the anterior cingulate, frontal, orbitofrontal, and anterior temporal cortices. The posterior insula has connections to the cingulate and parietal cortices. The middle insula has also shown this same overlap (Uddin et al.). These connections form networks essential for sensory perception, such as touch, pain, taste, sound, and smell, interoception, and emotions to guide decision-making and the regulation of behavior.

The input received from the posterior insula to the anterior insula then gets allocated to these different regions, such as the striatum in the dorsal region, basolateral/central amygdala, and the motor and sensory cortices (Zhang et al.). This indicates the insular cortex's heavy involvement in sensory perception. However, the amount of output from the agranular insula (anterior insula) is not uniform throughout the whole cortex. Of the five insular subregions analyzed (granular (GI), dysgranular (DI), dorsal agranular (AId), ventral agranular (AIv), and posterior agranular (AIp)), the dorsal agranular insula (AId) produces about 60% of total output, followed by the ventral agranular (AIv) at 6% (Zhang et al.). It is a cyclical process where, when external, setting-based situations change a physiological condition, such as heart rate increasing after seeing a spider, it alters the activity of the insular cortex, indicating changes in the bodily state, and the anterior insula turns that information into subjective feelings, which can then update those physiological conditions in return, like adjustments in heart and respiration rates. Therefore, the insula serves as a bridge between sensation and emotion (Zhang et al.), allowing us to make decisions by utilizing both sensory and emotional information to assess the value of options and further evaluate potential outcomes of each choice. This process can be understood through the somatic marker hypothesis by Antonio Damasio (Bechara, Antoine, and Antonio R. Damasio), a theory proposing that decision-making is also a result of emotional experiences due to bodily states and vice versa associated with past experiences (somatic markers), which also contribute to the cycle of updating our physiological condition. Modern research emphasizes that these somatic signals interact with higher cognitive evaluations rather than dictating one's choices automatically.

Sex Differences in Insula Anatomy and Decision Making

Understanding the impact of neurodegenerative diseases such as Alzheimer's on the insular cortex requires understanding how baseline, healthy male and female brains differ in the structure and organization of the insular cortex before the onset of the disease. A study that looked at the brains of 47 participants, with 24 males and 23 females, utilizing Individual Brain Atlases using Statistical Parametric Mapping (IBASPM) found that males typically have greater total brain volumes (around 1386.91 cm^3), including the insular cortex, than females (around 1244.21 cm^3) (Öz et al.). However, it is important to note that larger overall brain volumes do not signify intelligence. The biological sex of an individual influences cerebral organization due to interactions involving sex chromosomes, exposure to hormones involved in biological sex development/function, and circulating adult gonadal hormones. Similarities between the sexes include the cerebrum's left side being slightly larger (around 575.10 cm^3) than the right side (around 573.17 cm^3). There were slight differences in the insula as well, with the right insula (around 8.25 cm^3) being slightly larger than the left (around 7.82 cm^3) (Öz et al.). This shows that, in general, human brains are naturally lopsided, with the overall brain mass being slightly larger on the left, whereas specific sections like the insula are slightly larger on the right. Crucially, from the source morphometric tables to recalculate insular volumes normalized for total brain volume (TBV) and total cerebral volume (TCV) shows that the larger insular volume in males disappears after accounting for overall brain size (Öz et al.). When adjusted for brain size, females were shown to have similar or slightly higher relative insular volumes than males, while males show greater tissue density (Öz et al.). This suggests that sex differences in the insular cortex are not explained by overall volume alone, but may be due to differences in hemispheric specialization, cellular structure, plasticity, and network connection patterns (Iezzi et al.; Macey et al.; Öz et al.).

Interoception, a key part influenced by the insula, also has sex differences in its accuracy during decision-making. A study conducting a SAG (self-awareness questionnaire) found that females scored higher than males (11 vs 7) for internal gut feelings for visceral sensations and scored higher for somatosensory feelings such as muscle tension or skin sensations (14 vs 12) (Longarzo et al.). Using MRI scans, researchers found the thickness of the brain's outer surface did not differ significantly between sexes. Only the volume, the total tissue space, showed variations. For females, body awareness was linked not just to one region, but to a whole network of interconnected brain regions working together: the left insula, parietal cortex, and cingulate areas. Females with a larger left insula correlated with higher scores and overall body awareness ($R^2 = 0.41$) (Longarzo et al.). For males, on the other hand, bodily awareness was only correlated with one area: the precuneus. This heightened interoception in females closely matched the physical size of their left insula. However, it is important to note that there are variations in studies regarding this. Intriguingly, this self-reported study contrasted with objective testing, where, compared to men, women exhibited worse interoceptive accuracy across blood glucose levels, respiration, blood pressure, nicotine, sexual concordance, and gastric (Prentice et al.). Men frequently scored higher on detecting internal bodily signals such as heartbeats and respiration.

To understand why self-reported bodily awareness differs from objective testing between sexes, psychologists have turned to cognitive and social models. Theories such as Pennebaker

and Roberts' (1992) theory of gender differences in emotion provide one theoretical explanation, hypothesizing that this can happen due to different processing strategies during decision-making: while males tend to rely more directly on autonomic and insular signals to guide emotional reactions, females depend on internal sensations in a wider social and environmental context (Prentice et al.). This shows that decision-making reflects distinct information-processing strategies rather than differences in overall ability. In decision-making tasks such as the Iowa Gambling Task, females showed more sensitivity to loss and active punishment, and prioritized options that minimized repeated negative outcomes.

Males, on the other hand, demonstrated a higher tolerance for frequent loss to optimize final overall gains (Grissom and Reyes). However, when tasks control for losses or remove active punishment, those behavioral differences disappear, which confirms that neither sex possesses an innate advantage in executive decision-making. Instead, they reveal how the brain uses different sensory and physiological factors, like insular interoceptive signals versus contextual cues, to evaluate risks and rewards regarding decision-making. These baseline differences provide a benchmark to distinguish normal sex differences in decision-making from any degeneration caused by diseases such as Alzheimer's.

It is important to consider alternative explanations for these sex differences. Rather than reflecting a biological difference in insular structure, the female advantage in self-reported interoception could partly come from subjective reporting biases or differences in understanding emotional vocabulary, whereas male advantages in factors such as detecting heartbeats can be influenced by physiological factors such as lower resting heart rates rather than naturally advantageous processing. Additionally, differences in decision-making tasks like the Iowa Gambling Task can be influenced by prefrontal cognitive control and differences in risk tolerance rather than because of insular cortex-related interoception alone. Recognizing these confounding variables is essential because it shows that sex differences in decision-making are complex and can not be explained solely by differences in insular anatomy.

Alzheimer's Pathology in the Insular Cortex

During Alzheimer's, many neurons stop functioning properly and lose connections with other neurons. This degeneration involves two proteins: beta-amyloid and tau. Small clusters of beta-amyloid, called oligomers, float between brain cells, blocking cell communication, and over time, these clusters form amyloid plaques. The plaques trigger the immune system, causing swelling and inflammation. The tau protein, on the other hand, goes through hyperphosphorylation, which changes the protein chemically, and it detaches from inside the cell, leading to the proteins clumping together to form neurofibrillary tangles. Beta-amyloid plaques outside the cell trigger chemical changes that cause tau tangles to form inside the cell. Once the tau becomes toxic, it enhances the beta-amyloid inflammation, and the cycle repeats (Armstrong). However, the traditional spread of the amyloid hypothesis, suggesting that amyloid-beta triggers tau pathology and leads to Alzheimer's disease, is debated. Recent biochemical and electrophysiological evidence suggests that soluble forms of amyloid-beta and tau may separately contribute to synaptic dysfunction, which challenged the idea that amyloid-beta is the only cause of early disease stages (Guisano et al.). Together, both proteins

spread through the brain, destroying pre- and postsynaptic sites. More specifically, it alters molecular signaling such as the phosphorylation/activation of cAMP response element-binding protein (CREB) (a transcription factor needed for memory formation). Furthermore, both reduce levels of brain-derived neurotrophic factor (BDNF), a critical factor supporting neuronal survival, synaptic plasticity, and memory (Gulisano et al.).

Bonthius et al. studied post-mortem brain tissue from 17 Alzheimer's patients, investigating how the disease affects the insular cortex. Using specialized staining techniques, it was found that amyloid plaques were spread across different cortical layers, while the neurofibrillary tangles were mainly in the deeper layers (Layers V and VI). Tangle density additionally varied by region, where the agranular insula, in the front, was most affected, while the dysgranular insula showed moderate damage. The granular insula, located towards the back, was least affected (Bonthius et al.). Furthermore, researchers discovered that insular neurofibrillary tangle accumulation correlates directly with the increasingly severe pathology in the entorhinal cortex (EC). Because these regions are directly connected, the findings suggest that tau-related sameg may spread between them. Imaging studies further support this relationship, showing that the anterior insular cortex deteriorates alongside the entorhinal and hippocamal regions during the progression of Alzheimer's disease (Coleman et al., Moon et al., Qu et al.).

However, Alzheimer's does not only affect memory regions such as the hippocampus, entorhinal cortex, and cerebral cortex; it also affects the insular cortex. The disease has exhibited significant gray matter volume reductions due to beta-amyloid and tau accumulation in the insula, along with the hippocampus, temporal lobes, precuneus, cingulate gyrus, and inferior frontal cortex (Wu et al.). The hippocampus is famous for being the main region that shrinks, causing memory loss, but the insula's folding (gyri) also predicts memory consolidation into long-term memory only when hippocampal shrinkage occurs alongside amyloid plaques. So whether an individual with amyloid plaques can consolidate memories may depend on whether their insula is still healthy and folded, or flattened due to atrophy. The insula, specifically the anterior insula, is involved in the brain's salience and default mode network, allowing the brain to assess important cues within internal thoughts and external tasks critical to intuitive processes. As amyloid and tau accumulate, the communication between those two networks starts to fail, which worsens memory when the hippocampus is already atrophied because the insula can not regulate those networks or coordinate with the limbic system, of which the hippocampus is a part (Coleman et al.). The insula has also been associated with neuropsychiatric symptoms (NPS) and changes in cardiovascular and autonomic control (Moon et al.).

Decision-making under normal conditions occurs when your brain actively participates in receiving and sending internal physiological feedback to the body. There is an internal model of what the body should feel like in the brain, so when there is a difference between the expected conditions and what the body feels, such as a rise in heartbeat, that is the prediction error. In Alzheimer's, this system breaks down due to structural damage across the brain networks and insula. Patients with Alzheimer's were tested to show they have significantly lower accuracy in

correctly detecting internal cues in their heart rhythm. The areas responsible for processing those internal sensations- the insular cortex, cingulate cortex, ventromedial prefrontal cortex (vmPFC), and the amygdala- also went through gray matter atrophy (Sun et al.).Overall, interoceptive deficits are tied to atrophy within memory regions like the hippocampus and temporal lobes. Because patients struggle with memory and learning, they can not compare gut feelings with past experiences to determine the risks and rewards of a decision reliably.

Conclusion

Structurally, male brains are larger in volume across the cerebrum and insula, yet both sexes have natural asymmetry in both hemispheres, where there is a larger left cerebrum and a larger right insula (Öz et al.). Importantly, when insular columns are normalized for total brain volume, this sex difference disappears, demonstrating that female and male insular cortices have similar relative insular volumes. Functional differences may instead be rooted in differences in hemisphere specialization, brain network connections, and plasticity (lezzi et al., Öz et al.). Females show heightened interoceptive sensitivity, self-reported, where this awareness is associated with activity in several brain regions centered on the left insular cortex, parietal lobe, and cingulate cortices. Males, however, demonstrated more interoceptive accuracy on objective physiological measures of cardiac and respiratory functions, where they relied on more autonomic feedback when regulating emotional responses. Behaviorally, neither sex holds an innate advantage in decision-making. Instead, they utilize different strategies: females place more importance on immediate negative outcomes, so they use internal signals with environmental cues rather than internal signals, whereas males show higher tolerance for repeated negative losses for longer-term gains by using internal signals more. As Alzheimer's progresses, the accumulation of beta-amyloid oligomers and hyperphosphorylated tau disrupts synapses, which leads to atrophy across regions like the insula. The anterior insula specifically goes through significant gray matter volume loss and a reduction in gyri due to flattening of the surface. This leads to the anterior insula losing its ability to communicate with other brain regions within the salience and default networks, which hinders the ability to evaluate risks and rewards necessary for interoception. This, including the deterioration of memory regions such as the hippocampus, also diminishes the physiological somatic markers from memory-guided intuitive processes (Bechara & Damasio; Damasio).

Combining these findings leads to the theoretical hypothesis, rather than empirical certainty, that sex differences in the baseline insular cortex organization contribute to different patterns of decline in interoceptions and intuitive decision making in Alzheimer’s disease. This theory suggests that cognitive decline is driven not by how differently the insula is atrophied, but by some baseline differences in the structure that the disease dismantles. Because interoceptive awareness in females and decision-making are highly dependent on the brain network consisting of the left insula, cingulate, and parietal regions, early Alzheimer’s pathology is hypothesized to disproportionately disrupt contextual intuition in women, leading to earlier difficulties in utilizing emotional and environmental cues to guide choices.

Conversely, since men rely more heavily on raw bodily autonomic feedback to guide risk-based decisions, the atrophy of the anterior insula due to tau accumulation is hypothesized to deprive them of this primary source of guidance for decision-making, thereby increasing their vulnerability to disadvantageous, risky decisions in uncertain environments. Somatic marker-based decision-making is also deeply affected by Alzheimer's because the process fundamentally requires sensations paired with memories to make calculated decisions (Domasio; Sun et al.). In Alzheimer's, in both sexes, the hippocampal-entorhinal network and the insula experience a severe breakdown in this feedback loop. The difference is that females lose their contextual multisensory strategies, while males lose their direct physiological signals that evaluate their risks. To formally evaluate this hypothesized model, future research must move beyond theory to empirical testing by using longitudinal studies of people with MCI and early stage Azleheimr's disease. These studies may compare biological sexes and APOE-4 allele with using heartbeat tracking tasks to measure interocpetions, the Iowa Gambling Task to assess decision making, and PET scans to track tau buildup in the anterior insula. Addressing this will be essential for developing sex based diagnostic tools that can detect and support intuitive decision-making discrepancies earlier in neurodegenerative diseases such as Alzheimer's.

Work Cited

Armstrong, Richard A. “What Causes Alzheimer’s Disease?” Folia Neuropathologica, vol. 51, no. 3, 2013, pp. 169–88.

Bechara, Antoine, and Antonio R. Damasio. "The Somatic Marker Hypothesis: A Neural Theory of Economic Decision." Games and Economic Behavior, vol. 52, no. 2, 2005, pp. 336-72.

Bonthius, Daniel J., et al. “Pathology of the Insular Cortex in Alzheimer Disease Depends on Cortical Architecture.” Journal of Neuropathology & Experimental Neurology, vol. 64, no. 10, 2005, pp. 910–22.

Castelhano, Joao, et al. “The Role of the Insula in Intuitive Expert Bug Detection in Computer Code: An fMRI Study.” Brain Imaging and Behavior, vol. 13, no. 3, 2019, pp. 623–37.

Coleman, Michelle M., et al. “Interaction of Insula and Hippocampus in Memory Dysfunction in Alzheimer’s Disease.” Journal of Alzheimer’s Disease Reports, vol. 9, 2025, p. 25424823251407323.

Craig, A. D. "How Do You Feel—Now? The Anterior Insula and Human Awareness." Nature Reviews Neuroscience, vol. 10, no. 1, 2009, pp. 59-70.

Damasio, Antonio R. Descartes' Error: Emotion, Reason, and the Human Brain. G.P. Putnam's Sons, 1994.

Damasio, Antonio R. "The Somatic Marker Hypothesis and the Possible Functions of the Prefrontal Cortex." Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, vol. 351, no. 1346, 1996, pp. 1413-20.

Droutman, Vita, et al. “Roles of the Different Sub-Regions of the Insular Cortex in Various Phases of the Decision-Making Process.” Frontiers in Behavioral Neuroscience, vol. 9, 2015, p. 309.

Grissom, Nicola M., and Teresa M. Reyes. “Let’s Call the Whole Thing off: Evaluating Gender and Sex Differences in Executive Function.” Neuropsychopharmacology, vol. 44, no. 1, 2019, pp. 86–96.

Gu, Xiaosi, et al. “Anterior Insular Cortex and Emotional Awareness.” Journal of Comparative Neurology, vol. 521, no. 15, 2013, pp. 3371–88.

Gulisano, Walter, et al. “Role of Amyloid-β and Tau Proteins in Alzheimer’s Disease: Confuting the Amyloid Cascade.” Journal of Alzheimer’s Disease, vol. 64, no. s1, 2018, pp. S611–31.

Iezzi, Daniela, et al. "Sexual Differences in Neuronal and Synaptic Properties Across Subregions of the Mouse Insular Cortex." Biology of Sex Differences, vol. 15, no. 1, 2024, p. 29.

Kortz, Michael W., and Kevin O. Lillehei. “Insular Cortex.” StatPearls [Internet], StatPearls Publishing, 2023.

Laws, Keith R., et al. “Sex Differences in Alzheimer’s Disease.” Current Opinion in Psychiatry, vol. 31, no. 2, 2018, pp. 133–39.

Lighthall, Nichole R., et al. “Gender Differences in Reward-Related Decision Processing under Stress.” Social Cognitive and Affective Neuroscience, vol. 7, no. 4, 2012, pp. 476–84.

Longarzo, Mariachiara, et al. “Gender Brain Structural Differences and Interoception.”

Frontiers in Neuroscience, vol. 14, 2021, p. 586860.
Macey, Paul M., et al. “Sex Differences in Insular Cortex Gyri Responses to the Valsalva Maneuver.” Frontiers in Neurology, vol. 7, 2016, p. 87.
Moon, Yeonsil, et al. “Pathomechanisms of Atrophy in Insular Cortex in Alzheimer’s Disease.” American Journal of Alzheimer’s Disease & Other Dementias®, vol. 30, no. 5, 2015, pp. 497–502.
Namkung, Ho, et al. “The Insula: An Underestimated Brain Area in Clinical Neuroscience, Psychiatry, and Neurology.” Trends in Neurosciences, vol. 40, no. 4, 2017, pp. 200–07.
Öz, Fatma, et al. “The Role of Lateralisation and Sex on Insular Cortex: 3D Volumetric Analysis.” Turkish Journal of Medical Sciences, vol. 51, no. 3, 2021, pp. 1240–48.
Prentice, Freya, et al. “Gender Differences in Interoceptive Accuracy and Emotional Ability: An Explanation for Incompatible Findings.” Neuroscience & Biobehavioral Reviews, vol. 141, 2022, p. 104808.
Prentice et al., “Gender Differences in Interoceptive Accuracy and Emotional Ability: An Explanation for Incompatible Findings.” Neuroscience & Biobehavioral Reviews, vol. 141, 2022, p. 104808.
Sun, Weiyi, et al. “Brain Neural Underpinnings of Interoception and Decision-Making in Alzheimer’s Disease: A Narrative Review.” Frontiers in Neuroscience, vol. 16, 2022, p. 946136.
Uddin, Lucina Q., et al. “Structure and Function of the Human Insula.” Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society, vol. 34, no. 4, 2017, p. 300.
Wang, Xingchao, et al. “Anterior Insular Cortex Plays a Critical Role in Interoceptive Attention.” Elife, vol. 8, 2019, p. e42265.
Wenk, Gary L. “Neuropathologic Changes in Alzheimer’s Disease.” Journal of Clinical Psychiatry, vol. 64, 2003, pp. 7–10.
Wu, Zhanxiong, et al. “Gray Matter Deterioration Pattern during Alzheimer’s Disease Progression: A Regions-of-Interest Based Surface Morphometry Study.” Frontiers in Aging Neuroscience, vol. 13, 2021, p. 593898.
Zhang, Ruohan, et al. “The Insular Cortex: An Interface Between Sensation, Emotion and Cognition: R. Zhang et al.: The Insular Cortex: An Interface Between Sensation, Emotion and Cognition.” Neuroscience Bulletin, vol. 40, no. 11, 2024, pp. 1763–73.

Differential Inflammatory and Cytotoxic Responses to Nanoparticles in HMDM and THP-1 Macrophages By Nikhil Rajagopalan

Abstract

Atherosclerosis, a common cause of death, is partially driven by macrophage-induced inflammation. Nanoparticles can be used in drug delivery, especially for cardiovascular disease treatments, but their effects on immune cells remain unclear. This study performed a meta-analysis of 3 publicly available experimental datasets to examine the cytotoxic and inflammatory effects of five reference nanomaterials on human monocyte-derived macrophages (HMDMs) and THP-1 macrophage-like cells. Data were compared using ratio-to-control normalization to reduce cross-experimental error. Zinc oxide (ZnO) reduced THP-1 cell viability significantly more than that of HMDMs, whereas silicon dioxide (SiO_2) showed the opposite trend. Titanium dioxide (TiO_2) was not cytotoxic in either model. ZnO induced tumor necrosis factor-alpha (TNF-α) release earlier and more severely in HMDMs than THP-1s. A toxicity index revealed that both cell death and inflammation drove ZnO toxicity in THP-1s. Toxicity of SiO_2 was driven by normal macrophage behavior in HMDMs; however, in THP-1s, toxicity increased almost exclusively due to inflammation. Of the three physicochemical properties studied, primary diameter followed expected toxicity trends, but dissolution rate and surface area did not. However, as chemical species confounds the effects of specific physicochemical properties, results cannot be attributed to just one characteristic. Despite this, these findings demonstrate compelling observations that should be further studied and can be used to develop nontoxic nanomaterials capable of life-saving atherosclerosis treatment.

Keywords

Medical and Health Sciences; Material Science; Nanomaterials; Cardiovascular Disease; Macrophages

Introduction

Cardiovascular disease (CVD) is the leading cause of death worldwide, taking the lives of over 17.9 million people per year.[1] CVD includes a broad range of conditions, such as atherosclerosis, stroke, heart failure, and aneurysm. More specifically, atherosclerosis is the hardening of arteries caused by the buildup of plaque, which limits blood flow, leading to heart attacks or stroke.[2] Inflammation driven by macrophage activity has been recognized as a core contributor in the progression of atherosclerosis. This inflammation occurs when monocytes are sent to the site of plaque buildup, which causes inflammation.[3] The monocytes then differentiate into macrophages, which release inflammatory cytokines and transform into foam cells.[3,4] Foam cells absorb lipids such as oxidized LDL, causing plaque buildup and worsening inflammation. This process creates a vicious cycle in which additional macrophages are recruited to regulate inflammation, but only end up causing more.

Nanomaterials have revolutionized the field of medicine, advancing areas such as drug delivery, diagnostic imaging, and medical device manufacturing.[5] Nanotechnology can improve the speed, accuracy, and sensitivity of these applications.[5] Lab-on-nanoparticle technologies have also evolved out of the field, where small devices can perform multiple functions within the body.[5] Nanoparticles can also be designed to target specific tissues and cells.[5] To treat cardiac disease, nanomaterials can be designed to carry life-saving drugs to dissolve plaque. The most commonly used nanomaterials include titanium oxide, zinc oxide, and silicon dioxide. However, these nanomaterials can become toxic to certain cells depending on certain physicochemical properties. Therefore, understanding their interactions with immune cells is extremely important, given their increased use in medical treatments.[5]

Specifically, these nanomaterials can interact with macrophages in harmful ways, causing cell death and higher inflammation, which in turn, can worsen atherosclerosis. Hence, the study of macrophage immunotoxicology is necessary for the creation of safe nanomaterials. This paper is a meta-analysis of several available datasets, reviewing five different nanomaterials and three selected physicochemical characteristics across two types of macrophage cell lines to determine which properties result in the highest toxicity and whether that holds true in both cell lines to establish a trend. Primary particle diameter, surface area, and dissolution rate are three key physicochemical properties of nanomaterials which are known to influence toxicity.[6] Firstly, nanomaterials with high surface-to-volume ratios can absorb relatively large amounts of medication and move easily through the bloodstream, making them ideal carriers.[7] However, this surface-to-volume ratio is directly responsible for reactivity with macrophages, possibly resulting in higher toxicity. Smaller particles by size (diameter) with more surface area generally have greater reactivity, while the relationship is vice versa for larger particles with less surface area. Additionally, dissolution rate has been associated with generic toxicity in literature, especially for ZnO.[8] This involves dissolved ZnO interacting with the immune cells, leading to high amounts of free metal ions, causing cell damage through oxidative stress.[9] This suggests that the generic association of high dissolution rate with toxicity might not hold true, especially when comparing ZnO to SiO_2.

To establish trendlines across multiple cell types, two commonly used macrophage models used in nanotoxicology research were chosen, which are human monocyte-derived macrophages (HMDMs) and the immortalized THP-1 monocyte cell line.[10] Each of these monocyte types differentiates into a macrophage or macrophage-like phenotype.[11] While THP-1 cells are far easier to reproduce than HMDMs, they have a malignant origin, and their development may produce responses that differ from those of HMDMs.[12] This difference in cellular responses produced is crucial to understand in the context of how various nanomaterials could affect different immune cells.

Additionally, biomarkers such as TNF-α and cell viability were chosen to represent the number of cytokines released by the macrophages and the rate at which macrophages were dying or transforming into foam cells. TNF-α, or tumor necrosis factor alpha, is a chemical produced by macrophages to induce inflammation.[13] Therefore, TNF-α is a measure of

macrophage activation, and its prolonged elevation has been linked to endothelial dysfunction and atherogenesis.[14] After TNF-α is released, a decline in cell viability is often observed due to the induced inflammation and immunotoxicity.[15]

However, despite previous findings, comparison of HMDM and THP-1 responses to the same nanomaterials in the same conditions remains limited, as does the correlation between physicochemical properties and toxicity across cell types. This study attempts to address that gap by analyzing cytotoxicity and TNF-α release data for five nanomaterials. The use of both HMDM and THP-1 models is necessary to observe immunotoxicity relevant to CVD. This study also analyzes how various physicochemical characteristics of NM-110 ZnO and NM-200 SiO_2 could relate to toxicity, as those two materials have the most distinct responses and characteristics. This study hypothesizes that differences in cell responses are largely dependent on the characteristics of the surface-area-to-volume ratio and primary diameter size, as these are more indicative of reactivity than dissolution rate.

Methods

Data Source and Study Selection

Experimental data were retrieved from the ENanoMapper database (enanomapper.net), an open-access repository for nanomaterial safety data. Three published datasets were selected based on the use of standardized reference nanomaterials and overlapping experimental endpoints. Additional images were obtained through Creative Commons Licenses.[16, 17] Datasets were drawn from Brzicova et al., Bhattacharya et al., and Farcal et al.[18–20] Brzicova et al. examined molecular responses in THP-1 macrophage-like cells.[18] Farcal et al. provided data on responses from HMDM macrophages.[20] Both studies exposed their respective cell models to the same five nanomaterials: NM-103 (TiO_2), NM-110 and NM-111 (ZnO), and NM-200 and NM-203 (SiO_2). Bhattacharya et al. performed cytotoxicity screening across nineteen nanomaterials and ranked them by inflammogenic potential.[19] No experimental work was done by the author. This is a meta-analysis which solely uses the data in these 3 datasets.

Dataset	Cell model	Nanomaterials	Endpoint	Data used in figures
Brzicova et al. 2019	THP-1	NM-103, -110, -111, -200, -203	Viability & TNF-α	Figs 1-4
Farcal et al. 2015	HMDM	NM-103, -110, -111, -200, -203	Viability & TNF-α	Figs 1-4
Bhattacharya et al. 2017	THP-1	NM-103, -110, -111, -200, -203	Viability	Table 1, Figs 1-4, 6

Nanomaterial Properties

Physicochemical properties of all nanomaterials were sourced from the respective studies, with those reported by Bhattacharya et al. used as the primary reference.[19] All nanomaterials were reference materials provided by the European Commission Joint Research

Centre (EC-JRC) Nanomaterials Repository. Use of these reference materials allows for cross-study comparisons, since their physicochemical properties are the same across laboratories. The three quantifiable properties compared in this study were primary diameter, solubility, and surface area, as these showed the most significant variance across materials. The characterization data detailed in Table 1 are highly internally coherent. The reported surface-area and primary diameter ratios are fully consistent with what particle geometry and the respective material densities require, within a few percent.

Data Synthesis and Normalization

Individual datasets were extracted from the ENanoMapper database and combined into a single dataset for analysis in Microsoft Excel. Each study used different experimental conditions, cell lines, and measurement scales. A normalization procedure was applied to enable cross-dataset comparisons where each data point was divided by its respective baseline (untreated control) value. Results were displayed as fold-changes relative to control. This approach is standard practice in nanotoxicology for multi-study data integration and reduces the effect of inter-laboratory variability in measurements. Additionally, this method of normalization allows for cell viability to exceed 100%, indicating cell growth.

It is important to note that the toxicity index is constructed using data from multiple independent laboratories, combining primary HMDM data from Farcal et al. with THP-1 data from Brzicova et al. and Bhattacharya et al. While ratio-to-control normalization was used to minimize experimental errors between laboratories, synthesizing an index across multiple studies inherently introduces variations in baseline and culture handling. While this approach allows for cross-study comparisons that would otherwise be impossible, the resulting index reflects relative trends rather than cross-laboratory observations that should be taken absolutely.

Statistical Analysis

Data are presented as means ± standard deviation calculated across 3 experimental replicates within each dataset. Student's t-tests were used to determine whether differences between nanoparticle-treated groups and untreated controls were statistically significant, with significance set at $p < 0.05$. The t-tests were conducted between replicate data and untreated controls, not normalized fold-change values. Standard deviations are shown on graphs to convey measurement error. Standard deviations were obtained from normalized data.

Data Visualization

Normalized data was displayed using bar graphs and line graphs, chosen to best communicate observed trends. Bar graphs show endpoint-specific responses across nanomaterial types and concentrations at fixed time points. Line graphs display dose-response relationships. All figures were made to visually compare results across the three datasets and across nanomaterial types. Figures 1-4 were created in Microsoft Excel. Figure 5 was sourced from Creative Commons images.[16, 17] Figure 6 was sourced from Bhattacharya et al.[19]

Results and Discussion

Nanomaterials produced cell-type-specific effects on macrophage viability

To assess the effect of different nanomaterials, data were extracted from eNanoMapper and their effect on cell viability was investigated (Figure 1). Normalized cell viability (%) data for five nanomaterials in HMDM and THP-1 cells at doses of 0-100 µg/mL are shown. At low doses, the viability remained near 100% for both cell types across all materials. Patterns differed by material as the dose increased. NM-110 reduced THP-1 viability to 51.99% at 25 µg/mL and 19.13% at 100 µg/mL, and HMDM viability at the same doses was 65.02% and 52.47% (Figure 1b). For NM-111, THP-1 viability fell to 49.85% and 19.94% at 25 and 50 µg/mL, compared to 77.78% and 79.07% in HMDMs (Figure 1c). It can be observed that ZnO is more toxic to the THP-1 cell line than HMDMs. The relationship was inverse to that of ZnO particles for both SiO_2 materials. NM-200 reduced HMDM viability to 49.09% at 100 µg/mL. THP-1 viability remained at 93.21% at the same dose (Figure 1d). NM-203 reduced HMDM viability to 36.32% at 100 µg/mL, while THP-1 viability was 60.07% at the same dose (Figure 1e). NM-103 showed little to no cytotoxicity in either cell line, and in fact, THP-1 viability for this material exceeded 107% for all doses (Figure 1a). In general, cell viability data became more significant as dosage went up for all nanomaterials. Since cell viability was expressed as a ratio to an untreated control group, due to normalization, it could exceed 100% (Figure 1).

These data show that ZnO materials reduce THP-1 cell viability far more than HMDM cell viability at the same doses, which is consistent with findings in literature.[19, 20] Where ZnO cytotoxicity was likely caused by ion-mediated toxicity, SiO_2 cytotoxicity was likely dependent on phagocytic capacity. Macrophages break down ZnO particles, and Zn^{2+} ions act as catalysts to produce Reactive Oxygen Species, forcing the macrophages to release more cytokines.[21] Zinc-induced toxicity was more severe in THP-1s than in HMDMs, where the THP-1s are limited in their defense against induced toxicity compared to HMDMs.[22] For SiO_2, phagocytic cells like HMDMs have a higher capacity to take in these particles, so at the same dose, HMDMs take in more silica particles than THP-1s, leading to the difference in toxicity observed between cell types.[23] These data demonstrate that cell-specific responses are just as important as physicochemical properties to informing nanomaterial design.

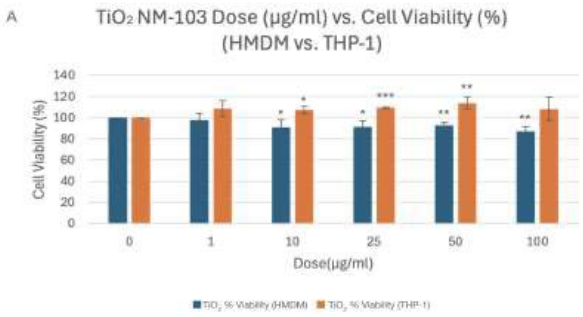

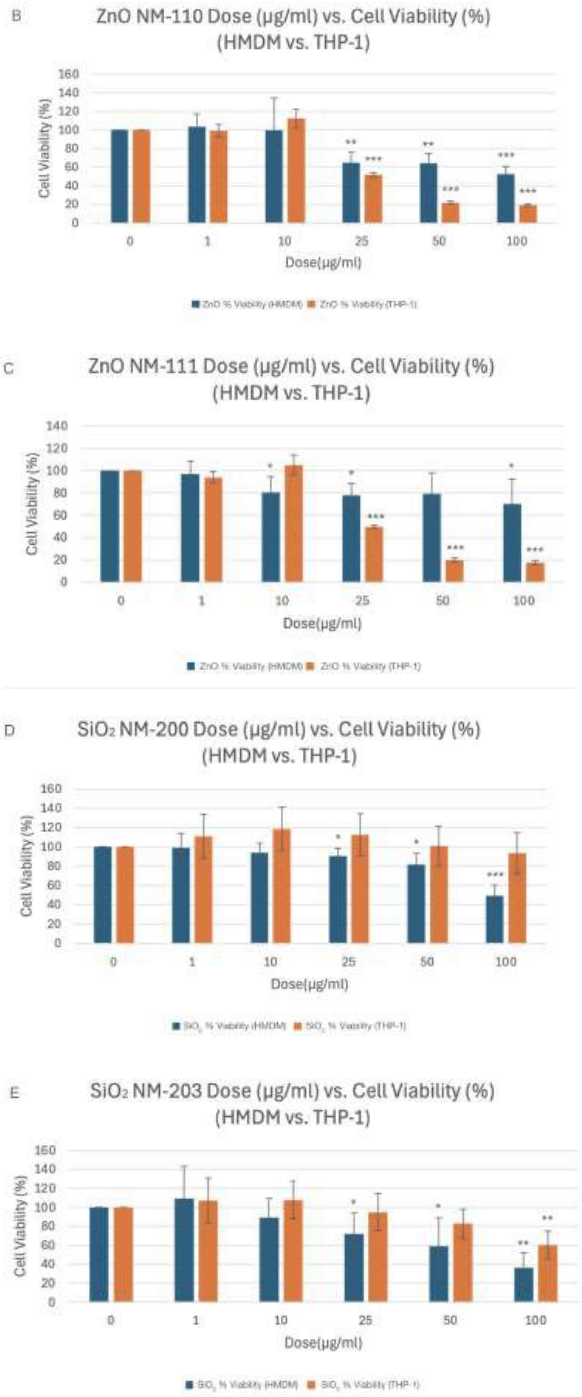

Figure 1. Bar charts showing cell viability (%) of indicated cell lines upon treatment with various nanomaterials, as indicated in each chart. Cell viability is represented as a fold-change over untreated control. Error bars derived from standard deviation. Significance was determined by Student's t-test, *<0.05. HMDM data was obtained from Farcal et al., and THP-1 data from Bhattacharya et al.[19, 20]

ZnO and SiO_2 induced different inflammatory responses in HMDM and THP-1 cells

As zinc oxide (NM-110) and silicon dioxide (NM-200) had the strongest impact on cell viability (Figure 1a), next, their effect on TNF-α was assessed (Figure 2). TiO_2 was proven to be non-toxic, so no further analysis was performed (Figure 1). All nanomaterials, regardless of cell type affected, had the same baseline TNF-α value of 1.00 due to the normalization method of converting data into a ratio of treated data/untreated control data. For NM-110, HMDM TNF-α rose to 1.94 at 1 µg/mL and peaked at 3.73 at 10 µg/mL. It then declined to 2.81 at 25 µg/mL (Figure 2a). In contrast, TNF-α of THP-1s remained near 1.00 for low doses of NM-110 but increased sharply to 8.39 at 25 µg/mL (Figure 2a). This suggests that HMDMs responded earlier to NM-110 exposure, whereas THP-1 cells showed a stronger inflammatory response only at a higher dose. For NM-200 SiO_2, HMDM showed a gradual increase in TNF-α release, increasing from 1.88 at 1 µg/mL to 2.01 and 2.63 at 10 and 25 µg/mL doses for NM-200, (Figure 2b). THP-1 cells also showed a dose-dependent TNF-α increase from 1 to 1.20 at 1 µg/mL, 2.81 at 10 µg/mL, and 3.46 at 25 µg/mL (Figure 2b). Thus, NM-200 induced inflammatory responses in both cell types, with THP-1 cells showing a stronger response at higher doses.

Figure 2 data illustrates ZnO HMDM TNF-α steadily rising, then falling, but ZnO THP-1 and both SiO_2 groups only showed TNF-α growth. This disparity is consistent with the known differences between the two cell types that come from their developmental origins, and the differences between the materials.[24,25] This pattern confirms the normal onset of cytotoxicity, where as more cells die, the macrophage population capable of secreting cytokines declines, in turn leading to more, then less, inflammation. However, a different pattern was observed in the THP-1 group for NM-110 ZnO, leading to the highest toxicity observed in this meta-analysis. When cells die from Zn^{2+}-induced oxidative stress, they release intracellular contents, which signal to neighboring surviving cells to rapidly increase TNF-α production.[25] This is a likely cause for the spike in TNF-α in this group. For SiO_2, TNF-α values for both cell lines steadily rose over time, which are results consistent with increased-dose responses. Overall, increased TNF-α production in THP-1s over HMDMs can be attributed to the absorption capacity of each cell line. HMDMs have been shown to have higher phagocytic capacities compared to THP-1s, implying that they can carry higher amounts of nanomaterials and achieve the same response.[11] Therefore, the same dose of a nanomaterial would elicit a more expressive response in THP-1s.[24] Future work should take into account these fundamental differences in cell biology.

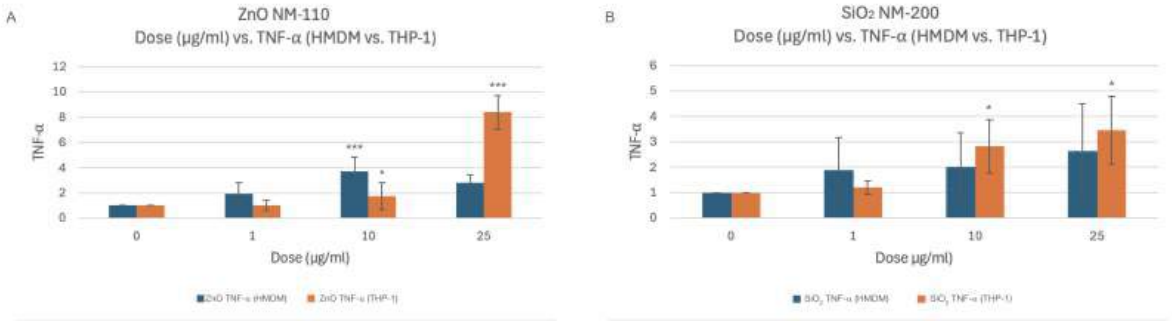

Figure 2. Bar charts displaying normalized TNF-α release, which is the ratio to the non-induced control datapoint, for doses up to 25 μg/mL. HMDM data were extracted from Farcal et al. and THP-1 data from Brzicova et al.[18, 20]

Combined viability and TNF-α analysis revealed distinct toxicity patterns for NM-110 ZnO

Next, the combined toxicity index was calculated for both nanomaterials and in macrophage models (Figure 3). Normalized cell viability (%) plotted and normalized TNF-α release were combined to create a toxicity index by dividing TNF-α by cell viability (%), for ZnO NM-110. This toxicity index was used to compare how inflammation and cell loss contributed to the overall toxic response. Higher toxicity index endpoint values indicate increased cell death (low viability) or increased inflammation (high TNF-α). In HMDMs, viability remained relatively stable from 0 to 10 μg/mL (within 4% of baseline), before decreasing to 65.02% at 25 μg/mL (Figure 3a). TNF-α release peaked at 3.73 at 10 μg/mL and then decreased to 2.81 at the higher dose (Figure 3a). In THP-1 cells (Figure 3c), cell viability increased to 112.63% at 10 μg/ mL before decreasing to 51.98% at 25 μg/mL. TNF-α release showed the opposite trend, remaining close to the baseline at 1.01, before continuously increasing to 8.39. As a result, the toxicity index for THP-1s remained low before sharply increasing all the way to 16.14 at 25 μg/mL (Figure 3d).

Figure 3 reveals that NM-110 toxicity in HMDMs tapers off at higher doses, and the elevated toxicity for NM-110 in THP-1 is driven by the collapse of cell viability. The change in HMDMs was likely due to increasing TNF-α values without a decline in cell viability, meaning that cells survived but suffered increasing amounts of inflammation. At higher doses, the inflammation had a more severe effect on the cells, finally reducing viability. Toxicity for NM-110 in THP-1 was the highest across all cell types and nanomaterials, driven by the collapse of cell viability and/or extreme levels of inflammation. NM-110 caused cell growth faster than it caused inflammation, but inflammation and cell death soon increased, whose combined effects ultimately drove a high toxicity index in the THP-1 cell line. These findings support the need to analyze dose-dependent effects on macrophages, as a dose suitable for one treatment may be catastrophic for another.

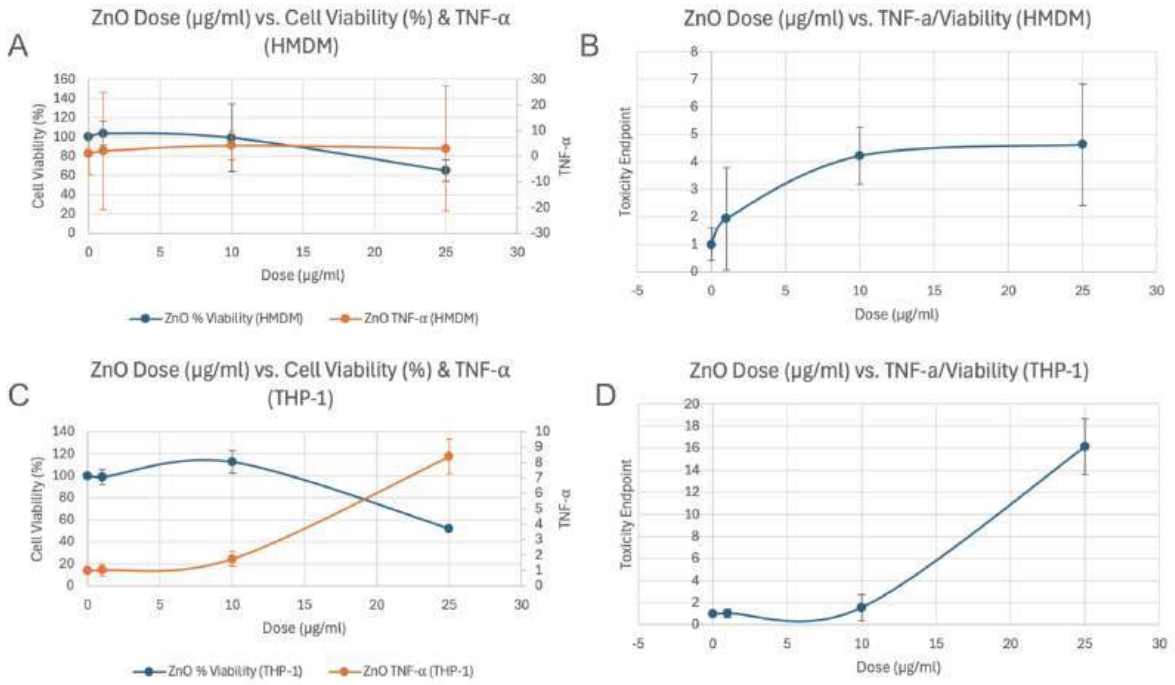

Figure 3. Line charts representing the toxicity index for NM-110 zinc oxide in HMDM and THP-1 cells. Cell viability data were retrieved from Farcal et al. (HMDMs) and Bhattacharya et al. (THP-1s).[19, 20] TNF-α data were retrieved from Farcal et al. (HMDMs) and Brzicova et al. (THP-1s).[18, 20]

NM-200 SiO_2 produced a more gradual toxicity profile than NM-110 ZnO

Next, the effect of silicon dioxide was assessed (Figure 4). In HMDM cells, viability steadily decreased from 100% at baseline to 81.64% at 25 µg/mL in an almost linear fashion (Figure 4a). TNF-α release was more unstable at lower doses, sharply increasing from 1 to 1.88 at just 1 µg/mL, and then continuing to increase to 2.63 at 25 µg/mL (Figure 4a). The toxicity index for HMDMs followed a similar pattern, sharply increasing to 1.90 at 1 µg/mL and then continuing to increase at a slower rate to 3.22 at 25 µg/mL (Figure 4b). In THP-1 cells, data showed an increase in cell viability (100% to 118.49%) and an increase in TNF-α (1 to 2.81) at 10 µg/mL, before viability decreased to 101.12% at 25 µg/mL, and TNF-α increased to 3.46 at 25 µg/mL (Figure 4c). This is consistent with the toxicity index, showing both a mild increase at 1 and 10 µg/mL before increasing at a higher rate to 3.42 at 25 µg/mL, as steady increases in TNF-α release and increases followed by decreases in cell viability drove this change. The toxicity endpoint was obtained with the same method as in Figure 3.

Figure 4 demonstrates a largely different toxicity profile for NM-200 SiO_2 from NM-110 ZnO in THP-1s, but a similar profile for both materials in HMDMs. Early toxicity in HMDMs was driven by inflammation rather than cell viability decreases, which is similar to ZnO in HMDMs (Figure 3a). This reflects normal macrophage behavior when confronted with foreign objects. Behavior in THP-1 cells was far milder, and in some cases, stimulatory. However,

mid-dose inflammation spiked and continued through the high doses. These findings confirm that SiO_2 is far less toxic to THP-1 cells than ZnO, consistent with the cytotoxicity patterns reported in Figure 1 and the hazard rankings by Bhattacharya et al and Farcal et al.[19,20] NM-200 toxicity in THP-1s is mainly driven by a steady increase in TNF-α and an increase and subsequent decrease in cell viability. Overall, SiO_2 induced normal macrophage behavior in HMDMs, and demonstrated benefits for THP-1 cells at low doses. Unfortunately, the practical use of SiO_2 is limited, as THP-1 cells are used almost exclusively in labs. Future work should focus on how to improve toxicity induced by SiO_2 nanoparticles in HMDM cells, as it is less toxic than NM-110 ZnO.

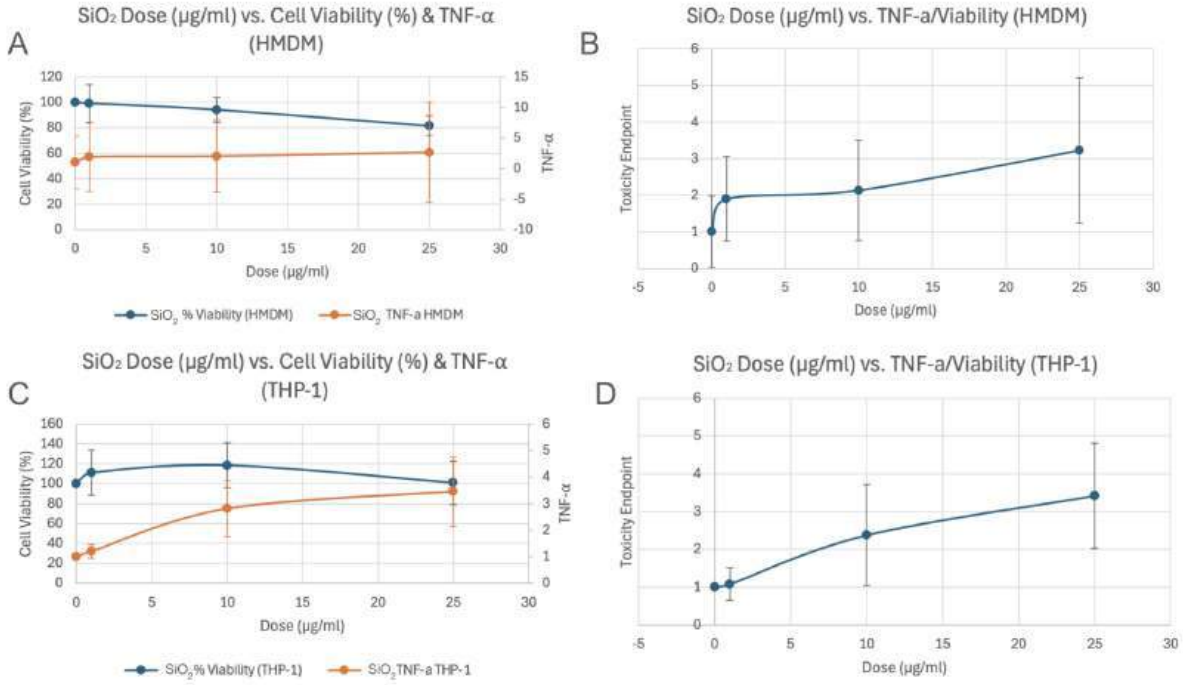

Figure 4. Line charts representing the toxicity index for NM-200 silicon dioxide in HMDM and THP-1 cells. Cell viability data were retrieved from Farcal et al. (HMDMs) and Bhattacharya et al. (THP-1s).[19, 20] TNF-α data were retrieved from Farcal et al. (HMDMs) and Brzicova et al. (THP-1s).[18, 20]

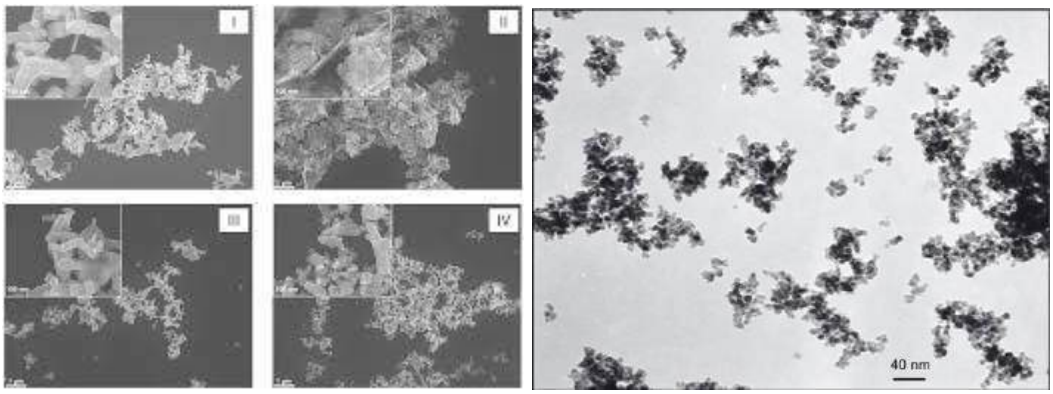

Figure 5. Representative images displaying surface and shape of nanomaterials NM-110 (left) and NM-200 (right).[16, 17] Images are used under Creative Commons license for distribution.

Physicochemical properties differed substantially between NM-110 ZnO and NM-200 SiO_2

Table 1 summarizes the physicochemical characteristics of NM-110 ZnO and NM-200 SiO_2. These characteristics were obtained from Brzicova et al. and Bhattacharya et al.[18,19] Table 1 references a surface area ratio, a primary diameter ratio, and a dissolution ratio. These ratios were obtained by dividing the greater characteristic metric (e.g. surface area for SiO_2) by the smaller characteristic metric (e.g. surface area for ZnO) to obtain a final unitless ratio. In the example provided, SiO_2 has a surface area 17.7 times that of ZnO. Additionally, ZnO has a diameter 8.03 times larger than SiO_2, and SiO_2 dissolves at a rate 44.3 times that of ZnO. These differences indicate that the two materials differ substantially in size, surface area, and solubility, all of which may influence their biological effects. As detailed in Table 1, SiO_2 exhibits a substantially higher absolute dissolution rate (144.0 ppm) compared to ZnO (3.25 ppm), yielding a dissolution ratio of 44.3. However, despite SiO_2 dissolving faster numerically, ZnO induced significantly greater cytotoxicity through intracellular Zn^{2+} ion-mediated oxidative stress (Figure 3, Figure 4). This indicates that the physical rate of dissolution is secondary to the toxicological nature of the released chemical species.

Table 1. Summary of physicochemical properties for NM-110 and NM-200.

Nanomaterial	NM-110 ZnO	NM-200 SiO_2
Primary Particle Diameter Ratio	8.03	1
Primary Particle Diameter	147.0 nm	18.3 nm
Surface Area Ratio	1	17.7
Surface Area (m^2/g)	12.4 m^2/g	220 m^2/g
Solubility Ratio	1	44.3
Dissolution (ppm)	3.25 ppm	144 ppm

Descriptive comparison of physicochemical characteristics and toxicity features of NM-110 ZnO and NM-200 SiO_2

Zinc oxide and silicon dioxide have different properties; therefore, we analyzed how they correlate with the toxicity index (Figure 6). An analysis of the 3 main physicochemical characteristics of NM-110 ZnO and NM-200 SiO_2 in Table 1 was used. Chosen characteristics included: surface area, primary particle diameter, and solubility. For the 3 quantifiable characteristics, the nanomaterial with a higher quantifiable property was divided by the material with the lower quantity for a characteristic ratio. Each ratio was assigned to its respective normalized toxicity value at 25 μg/mL. Bar charts are organized by nanomaterial and then subdivided into cell types. The graphs are designed to reflect the same endpoints in order to most effectively convey which quantifiable properties have the strongest correlation with toxicity outcomes. Since the comparisons are between distinct chemical species, there is no control group and this was a relative analysis.

The first graph, plotting surface area against the toxicity endpoint, shows a significant decrease in toxicity in THP-1s as surface area increases from 1 (ZnO) to 17.7 times that of ZnO (SiO_2). Toxicity decreased from 16.14 to 3.42. In HMDMs, toxicity decreased marginally with surface area (4.63 to 3.22) (Figure 6a).

Primary particle diameter showed a more consistent relationship with toxicity. As particle diameter increased, the toxicity index also increased in both cell types. It was observed that in HMDMs, as diameter increased, toxicity increased (From 3.42 to 4.63), and in THP-1s, the same happened but far more drastically (3.41 to 16.14) (Figure 6b).

Dissolution rate showed an inverse pattern. As the dissolution ratio increased, toxicity decreased mildly in HMDMs (4.63 to 3.22) but sharply decreased in THP-1s (16.14 to 3.42) (Figure 6c).

Table 1 and Figure 6 compare three quantifiable physicochemical characteristics of NM-110 ZnO and NM-200 SiO_2 against their toxicity endpoints in context with their expected behavior. The relationship between surface area did not align with expected trends; higher surface area is thought to be indicative of increased reactivity, which, in turn, should have led to increased toxicity for SiO_2, which was not observed. Primary diameter showed a stronger and more consistent association. Ordinarily, primary diameter, or particle size, affects smaller cells–such as THP-1s–far more than their larger counterparts, due to differences in phagocytic capacity.[23] The smaller the cell, the less foreign material it is able to process. The dissolution rate also had negative correlation, as SiO_2 dissolved far faster than ZnO, yet ZnO is considerably more toxic. This indicates that ion-mediated toxicity from Zn^{2+} ions contributes more to cell damage than dissolution rate alone. Another possible reason for this is that dissolved SiO_2 is easier for the macrophage to process than relatively insoluble ZnO. This would align with higher surface area (which is increased by dissolution) correlating with less toxicity. However, definitive conclusions cannot be drawn, as there was no control for these comparisons, so these are merely compelling observations to be further studied. Future work should involve more controlled experimentation which focuses on physicochemical variations

within chemical species. However, the results discussed present significant findings on dose-dependent toxicity and the effects of nanomaterial type.

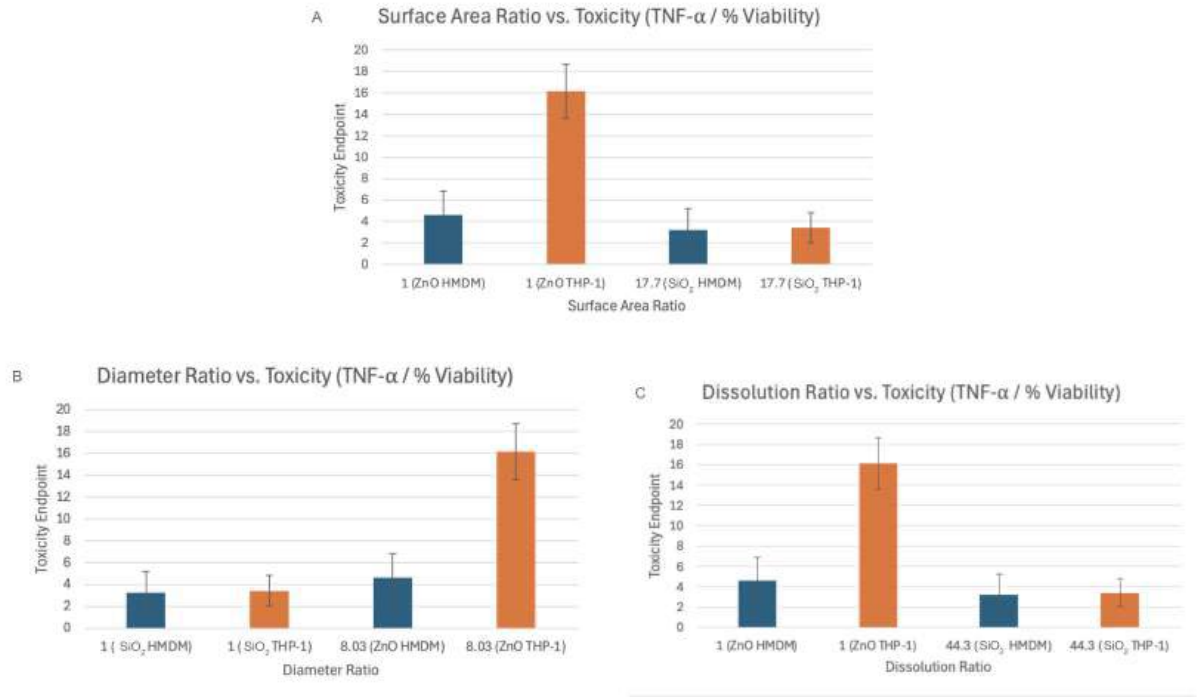

Figure 6. Bar charts showing the association between nanoparticle properties and toxicity index. Data are taken from previously calculated toxicity indices at 25 µg/mL (Figures 3b, 3d, 4b, 4d).

Overall, the results show that nanomaterial toxicity depends on both biological context and material properties. THP-1 cells were more sensitive to ZnO nanomaterials, whereas primary HMDMs were more sensitive to SiO_2 nanomaterials. These differences support the use of both primary and cell-line-derived macrophage models when evaluating nanomaterial safety for cardiovascular applications.

Several limitations apply to this study. Data were drawn from three independent studies with different experimental conditions. Additionally, this study did not include a control, as that data was not available. Normalization reduced variability between datasets, but differences in experimental methods and cell culture conditions may still have influenced the results. The toxicity index combines two endpoints into a single value, which simplifies material comparisons but may obscure how cytotoxicity and inflammation cause toxicity. However, including the original data points that the toxicity endpoints are composed of allows for direct attribution to which data point drove toxicity.

Conclusion

This meta-analysis compared cytotoxic and inflammatory responses in HMDM and THP-1 macrophages exposed to five nanomaterials. It also examined how three distinct

physicochemical properties relate to those responses. It was determined that TiO_2 was nontoxic to both cell types. ZnO nanomaterials were the most toxic, being more toxic to THP-1 cells than HMDMs. SiO_2 was less toxic than ZnO, and no significant difference was observed in the toxicity between HMDMs and THP-1s for SiO_2.

TNF-α data showed that HMDMs exhibited stronger and earlier inflammatory responses than THP-1s to ZnO. SiO_2 toxicity steadily increased with dosage across both cell types. A toxicity index combining cell viability and TNF-α data revealed that ZnO toxicity in THP-1 cells was driven by both cytotoxicity and inflammation, and that ZnO toxicity in HMDMs was driven by increasing TNF-α expression combined with delayed cytotoxicity. HMDM cells exposed to SiO2 materials displayed normal macrophage behavior, including slightly increasing TNF-α and continuously decreasing cell viability as dosage increased. On the other hand, THP-1 cells exposed to SiO2 showed a mild stimulatory effect on cell viability and continuously increasing TNF-α, contributing to increasing toxicity.

One of the three physicochemical properties analyzed showed consistent association with toxicity across both cell types; that being, as particle diameter increased, toxicity increased. Although the larger ZnO particles coincided with a higher toxicity index compared to the smaller SiO_2 particles, this effect is driven by the combined effect of the three properties and ion-mediated toxicity. ZnO caused significantly greater damage than SiO_2, despite lower dissolution rate. This indicates that the physical rate of dissolution is secondary to the toxicological nature of the released components, as seen in how the intracellular release of Zn^{2+} ions induces severe oxidative stress that outpaces macrophage responses. An alternative explanation was that high dissolution rate in SiO2 allowed for easier processing and lesser inflammation for the macrophages, leading to less toxicity. Higher surface area was associated with lower toxicity, contrary to traditional notions that surface area increases reactivity and, by extension, toxicity. Further study is needed to understand the effects of surface area and reactivity on toxicity.

This study highlights the importance of both biological model selection and material composition in nanotoxicity. Future work should compare primary and cell-line derived macrophages under matched experimental conditions and use larger, systematically characterized nanomaterial panels, in which physicochemical variables can be controlled independently. Such studies will be crucial for identifying reliable predictors of nanomaterial toxicity and for improving the safety assessment and design of nanomaterials intended for biomedical applications.

Acknowledgements

I would like to thank Professor Monika Gullerova for mentoring me through the research process. I would also like to thank Bengt Fadeel and Pavel Rossner, Jr. for sending me additional raw data from their experiments. This paper would not be possible without them. Finally, the ENanoMapper database and those who manage it were essential to the production of this paper.

Works Cited

World Health Organization. Cardiovascular Diseases. World Health Organisation. https://www.who.int/health-topics/cardiovascular-diseases#tab=tab_1.

Cleveland Clinic. Atherosclerosis: Types, Causes, & Treatments. Cleveland Clinic. https://my.clevelandclinic.org/health/diseases/16753-atherosclerosis-arterial-disease.

Henein, M. Y.; Vancheri, S.; Longo, G.; Vancheri, F. The Role of Inflammation in Cardiovascular Disease. International Journal of Molecular Sciences 2022, 23 (21), 12906. https://doi.org/10.3390/ijms232112906.

Chistiakov, D. A.; Bobryshev, Y. V.; Orekhov, A. N. Macrophage-Mediated Cholesterol Handling in Atherosclerosis. Journal of Cellular and Molecular Medicine 2015, 20 (1), 17–28. https://doi.org/10.1111/jcmm.12689.

Malik, S.; Muhammad, K.; Waheed, Y. Emerging Applications of Nanotechnology in Healthcare and Medicine. Molecules 2023, 28 (18). https://doi.org/10.3390/molecules28186624.

Gatoo, M. A.; Naseem, S.; Arfat, M. Y.; Dar, A. M.; Qasim, K.; Zubair, S. Physicochemical Properties of Nanomaterials: Implication in Associated Toxic Manifestations. BioMed Research International 2014, 2014, 498420. https://doi.org/10.1155/2014/498420.

Sim, S.; Wong, N. Nanotechnology and Its Use in Imaging and Drug Delivery (Review). Biomedical Reports 2021, 14 (5). https://doi.org/10.3892/br.2021.1418.

Wu, F.; Harper, B. J.; Harper, S. L. Comparative Dissolution, Uptake, and Toxicity of Zinc Oxide Particles in Individual Aquatic Species and Mixed Populations. Environmental Toxicology and Chemistry 2019, 38 (3), 591–602. https://doi.org/10.1002/etc.4349.

Liu, J.; Kang, Y.; Yin, S.; Song, B.; Wei, L.; Chen, L.; Shao, L. Zinc Oxide Nanoparticles Induce Toxic Responses in Human Neuroblastoma SHSY5Y Cells in a Size-Dependent Manner. International Journal of Nanomedicine 2017, Volume 12, 8085–8099. https://doi.org/10.2147/ijn.s149070.

Ding, T.; Du, Y.; Yang, B.; Tian, W.; Li, J.; Xie, J. Comprehensive Review of Macrophage Models: Primary Cells and Immortalized Lines across Species. Frontiers in Immunology 2025, 16. https://doi.org/10.3389/fimmu.2025.1640935.

Daigneault, M.; Preston, J. A.; Marriott, H. M.; Whyte, M. K. B.; Dockrell, D. H. The Identification of Markers of Macrophage Differentiation in PMA-Stimulated THP-1 Cells and Monocyte-Derived Macrophages. PLoS ONE 2010, 5 (1), e8668. https://doi.org/10.1371/journal.pone.0008668.

Tedesco, S.; De Majo, F.; Kim, J.; Trenti, A.; Trevisi, L.; Fadini, G. P.; Bolego, C.; Zandstra, P. W.; Cignarella, A.; Vitiello, L. Convenience versus Biological Significance: Are PMA-Differentiated THP-1 Cells a Reliable Substitute for Blood-Derived Macrophages When Studying in Vitro Polarization? Frontiers in Pharmacology 2018, 9, 71. https://doi.org/10.3389/fphar.2018.00071.

Jang, D.; Lee, A-Hyeon.; Shin, H.-Y.; Song, H.-R.; Park, J.-H.; Kang, T.-B.; Lee, S.-R.; Yang, S.-H. The Role of Tumor Necrosis Factor Alpha (TNF-α) in Autoimmune Disease and

Current TNF-α Inhibitors in Therapeutics. International Journal of Molecular Sciences 2021, 22 (5). https://doi.org/10.3390/ijms22052719.

Zhang, H.; Park, Y.; Wu, J.; Chen, X. ping; Lee, S.; Yang, J.; Dellsperger, Kevin C.; Zhang, C. Role of TNF-α in Vascular Dysfunction. Clinical Science 2009, 116 (3), 219–230. https://doi.org/10.1042/cs20080196.

You, K.; Gu, H.; Yuan, Z.; Xu, X. Tumor Necrosis Factor Alpha Signaling and Organogenesis. Frontiers in Cell and Developmental Biology 2021, 9. https://doi.org/10.3389/fcell.2021.727075.

Wilhelmi, V.; Fischer, U.; Weighardt, H.; Schulze-Osthoff, K.; Nickel, C.; Stahlmecke, B.; Kuhlbusch, T. A. J.; Scherbart, A. M.; Esser, C.; Schins, R. P. F.; Albrecht, C. Zinc Oxide Nanoparticles Induce Necrosis and Apoptosis in Macrophages in a P47phox- and Nrf2-Independent Manner. PLoS ONE 2013, 8 (6), e65704. https://doi.org/10.1371/journal.pone.0065704.

Ziaee, M.; Ganji, Z. Insecticidal Efficacy of Silica Nanoparticles against Rhyzopertha Dominica F. and Tribolium Confusum Jacquelin Du Val. Journal of Plant Protection Research 2016, 56 (No 3). https://doi.org/10.1515/jppr-2016-0037

Brzicova, T.; Javorkova, E.; Vrbova, K.; Zajicova, A.; Holan, V.; Pinkas, D.; Philimonenko, V.; Sikorova, J.; Klema, J.; Topinka, J.; Rossner, P. Molecular Responses in THP-1 Macrophage-like Cells Exposed to Diverse Nanoparticles. Nanomaterials 2019, 9 (5), 687. https://doi.org/10.3390/nano9050687.

Bhattacharya, K.; Gözde Kiliç; Costa, P. M.; Bengt Fadeel. Cytotoxicity Screening and Cytokine Profiling of Nineteen Nanomaterials Enables Hazard Ranking and Grouping Based on Inflammogenic Potential. Nanomaterials 2017, 11 (6), 809–826. https://doi.org/10.1080/17435390.2017.1363309.

Farcal, L.; Torres Andón, F.; Di Cristo, L.; Rotoli, B. M.; Bussolati, O.; Bergamaschi, E.; Mech, A.; Hartmann, N. B.; Rasmussen, K.; Riego-Sintes, J.; Ponti, J.; Kinsner-Ovaskainen, A.; Rossi, F.; Oomen, A.; Bos, P.; Chen, R.; Bai, R.; Chen, C.; Rocks, L.; Fulton, N. Comprehensive in Vitro Toxicity Testing of a Panel of Representative Oxide Nanomaterials: First Steps towards an Intelligent Testing Strategy. PLOS ONE 2015, 10 (5), e0127174. https://doi.org/10.1371/journal.pone.0127174.

Wang, B.; Zhang, Y.; Mao, Z.; Yu, D.; Gao, C. Toxicity of ZnO Nanoparticles to Macrophages due to Cell Uptake and Intracellular Release of Zinc Ions. Journal of Nanoscience and Nanotechnology 2014, 14 (8), 5688–5696. https://doi.org/10.1166/jnn.2014.8876.

Safar, R.; Doumandji, Z.; Saidou, T.; Ferrari, L.; Nahle, S.; Rihn, B. H.; Joubert, O. Cytotoxicity and Global Transcriptional Responses Induced by Zinc Oxide Nanoparticles NM 110 in PMA-Differentiated THP-1 Cells. Toxicology Letters 2019, 308, 65–73. https://doi.org/10.1016/j.toxlet.2018.11.003.

Malugin, A.; Herd, H.; Ghandehari, H. Differential Toxicity of Amorphous Silica Nanoparticles toward Phagocytic and Epithelial Cells. Journal of Nanoparticle Research 2011, 13 (10), 5381–5396. https://doi.org/10.1007/s11051-011-0524-7.

Liu, Y.; Li, H.; Czajkowsky, D. M.; Shao, Z. Monocytic THP-1 Cells Diverge Significantly from Their Primary Counterparts: A Comparative Examination of the Chromosomal Conformations and Transcriptomes. Hereditas 2021, 158 (1). https://doi.org/10.1186/s41065-021-00205-w.

Loeffler, H.; Jonitz-Heincke, A.; Peters, K.; Mueller-Hilke, B.; Fiedler, T.; Bader, R.; Klinder, A. Comparison of Inflammatory Effects in THP-1 Monocytes and Macrophages after Exposure to Metal Ions. Materials 2020, 13 (5), 1150. https://doi.org/10.3390/ma13051150.

Examining the Passage and Legacy of the Civil Rights Act of 1964 By Parvati Kaushik

Introduction

This paper analyzes events from 1950 to the present day, an era encompassing the buildup to, passage of, and effects of the Civil Rights Act of 1964 . During this period, African Americans faced integrated racial segregations.

The Civil Rights Act of 1964 (CRA 1964) is a landmark federal policy prohibiting discrimination in places of public accommodation, integrated schools and other public facilities, and employment discrimination (National Archives and Records Administration). This paper examines the roles of civil rights organizations, activists, policymakers, and segregationist opponents; the legal and protest strategies implemented to achieve policy reform and equality interests; and the act's effects on African American communities, including political representation and access to public resources and institutions. The research also considers the ongoing limitations of the act, particularly how sustained inequalities continue to impact African Americans. To locate and evaluate evidence, this project draws from a combination of archival and primary sources and academic research. Supreme Court cases and Congressional records inform discussions about legal and political implications of the CRA 1964, while scholarly articles assess how the act's outcomes apply today. This approach allows for an examination of how federal legislation and grassroots organizing influenced the accessibility of political and economic resources for African Americans during this time period. The project aims to situate the Civil Rights Act of 1964 within the context of racial politics to report the progress achieved and the remaining limitations that affect African Americans.

Civil Rights Coalition and Strategies

Understanding the Civil Rights Act of 1964 requires exploring the coalescing of actors, interests, and strategies that made its passage possible and assessing how these factors impacted African Americans. The civil rights coalition included political activists and religious leaders along with organizations such as the National Association for the Advancement of Colored People (NAACP), the Student Nonviolent Coordinating Committee (SNCC), and the Congress of Racial Equality (CORE). The NAACP became the leading civil rights organization and combated racial violence and discrimination through legal challenges and public protests. Between the 1910s and the 1950s, the organization built a lobbying membership base led increasingly by African Americans (Shaw et al.). The NAACP, among other actors listed, advocated for the collective interest: to officially eliminate Jim Crow segregation and enact legally enforced political, social, and economic equality. Actors achieved these objectives by dividing protest and litigation strategies into two major cycles. The first cycle garnered national sympathy by exposing brutality and inequality and introducing strategies as methods for achieving the established objectives. The second protest cycle escalated the urgency of the situation by increasing activist participation and spurring the passage of the CRA 1964 (Shaw et al.).

The First Major Protest Cycle

The first major protest cycle was characterized by several legal, political, and social shifts. The NAACP, led by Thurgood Marshall, challenged racial segregation in education. This effort culminated in the U.S. Supreme Court decision in *Brown v. Board of Education of Topeka, Kansas* on May 17, 1954. Justice Earl Warren conveyed the unanimous ruling, asserting that state authorized segregation in public schools directly violated the 14th amendment, and was therefore unconstitutional (National Archives and Records Administration). This landmark case overturned the precedent set forth in 1896 with the *Plessy v. Ferguson* Supreme Court case, which upheld a "separate but equal" doctrine for "white and colored races" (National Archives and Records Administration). Although the decision ordered desegregation "with all deliberate speed," several Southern states resisted the ruling through massive resistance campaigns and the closure of public schools to avoid integration (Shaw et al.). For instance, James Eastland, a Mississippi Senator, declared the Southern states' refusal to adhere to the legislation while Senator Harry Byrd from Virginia denounced the opinion as a "serious blow" to state "authority and welfare" (NAACP Legal Defense Fund). In response to this widespread backlash, civil rights organizers turned to mass mobilization strategies.

The Montgomery Bus Boycott took place from 1955 to 1956 and was sparked by the arrest of Rosa Parks who refused to give up her seat on a Montgomery, Alabama bus to a White passenger (Library of Congress). Dr. Martin Luther King Jr. coordinated this 381 day boycott that successfully ended bus segregation in 1956 through the Supreme Court decision in *Browder v. Gayle* (Shaw et al.). This demonstrated the effectiveness of sustained, nonviolent collective action and pushed toward an all-encompassing legislative transformation. Following the boycott's effectiveness, Martin Luther King Jr. and Black religious leaders founded the Southern Christian Leadership Conference (SCLC) to organize further nonviolent protests across the South. Despite these efforts, discriminatory sentiment from segregationist opponents remained. During the 1957 Little Rock school integration crisis, nine African American students were obstructed by the Arkansas National guard as they attempted to enter their high school. President Dwight Eisenhower sent federal troops to protect the students and enforce the desegregation order established by the courts (Library of Congress). Eisenhower also sent Congress a proposal for civil rights legislation, resulting in the Civil Rights Act of 1957. While this act created a commission to examine complaints of civil rights violations, many criticized the weak enforcement of the act because the commission could not take action to protect Black voters (Lowe). The largely powerless commission was due to Southern Democrat efforts to strip enforcement power and refusal to end literacy tests; however, its inadequate execution paved the way for further protest and eventually the CRA 1964 (Shaw et al.). As the limits of early efforts became clear, a second protest emerged whose strategies ultimately pressured Congress to enact the CRA 1964.

The Second Major Protest Cycle

The second protest cycle transpired in the early 1960s and adopted confrontational methods to accelerate federal action. In 1961, CORE initiated the Freedom Rides campaign to directly test the limits of the federal government's willingness to enforce anti-segregation legislation in transportation. Further, Martin Luther King Jr. directed the SCLC in launching a campaign of sit-ins and marches in Birmingham, Alabama to challenge the city's segregation policies (Library of Congress). Following the student-led sit-ins sparked by the Student Nonviolent Coordinating Committee (SNCC), direct-action tactics diffused in the South (Shaw et al.). After a series of meetings, marches on City Hall, lunch counter sit-ins and a boycott of vendors, King was arrested for violating an anti-protest injunction obtained by the city government (Stanford King Institute). He wrote the "Letter from Birmingham Jail" which challenged White moderates and clergy to engage in "nonviolent direct action," which he contended was key to achieving justice for the oppressed (Stanford King Institute). In 1963, the March on Washington for Jobs and Freedom was the largest non-violent civil rights demonstration in the nation's history to compel policymakers to pass civil rights reforms. The protestors urged Congress to enact "effective civil rights legislation" which included elimination of education, housing, and employment inequities. They conferred with federal representatives and President John F. Kennedy to advocate for structural reforms (Library of Congress). Ultimately, the mass action that took place during the second protest cycle introduced new protest techniques that pushed lawmakers to address civil rights concerns. Notably, the actions of protestors led to the "national unmasking of the viciousness of racial segregation" which created the demand for a comprehensive and legal civil rights safeguard (Rosenberg 1149). Due to this prominent fight for legislative action, President John F. Kennedy pressed Congress to consider a civil rights bill establishing protections for desegregation in public accommodations and nondiscrimination in federally assisted programs.

Passage and Provisions of the Civil Rights Act

The passage of the Civil Rights Act of 1964 was met with obstacles. Southern Democratic segregationists sought to end the bill via filibuster, but enough Republicans supported the bill to allow its ratification. Following Kennedy's death, President Lyndon Johnson signed the bill into law (National Archives and Records Administration). The act contained eleven provisions serving various protections. Collectively, they protected against discrimination in voting practices, commercial businesses, public facilities and schools, federally funded programs, and employment, which ended the Jim-Crow era of racial discrimination (Back).

Economic and Social Effects of the Civil Rights Act

The act improved the economic status of Black Americans following its adoption; for instance, Black men's wages increased both in absolute terms and relative to the wages of White men residing in the same region. The racial wage gap shrunk, with the wage discount for employed Black men declining from 66% in 1960 to 44% in 1970 [see Figure 1] (Bourne 1207). In addition, Black men experienced an increase in managerial positions. Black citizens also

secured more school funding, recreational facilities, and more equitable access to housing loans. Notably, the racial gap in life expectancy significantly decreased; in 1940, the difference between White and Black life expectancy at birth was 11 years, but this number fell to 4.5 years by 2008 [see Figure 2] (Bourne 1222).

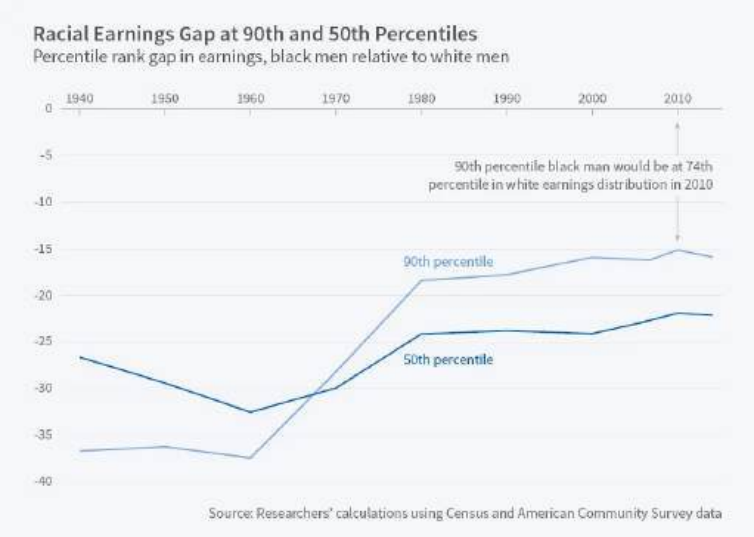

Figure 1: Earnings gap between Black and White men.

Bayer, Patrick, and Kerwin Kofi Charles. "The Earnings Gap between Black and White Men." National Bureau of Economic Research, 2018, https://www.nber.org/digest/jan17/earnings-gap-between-black-and-white-men.

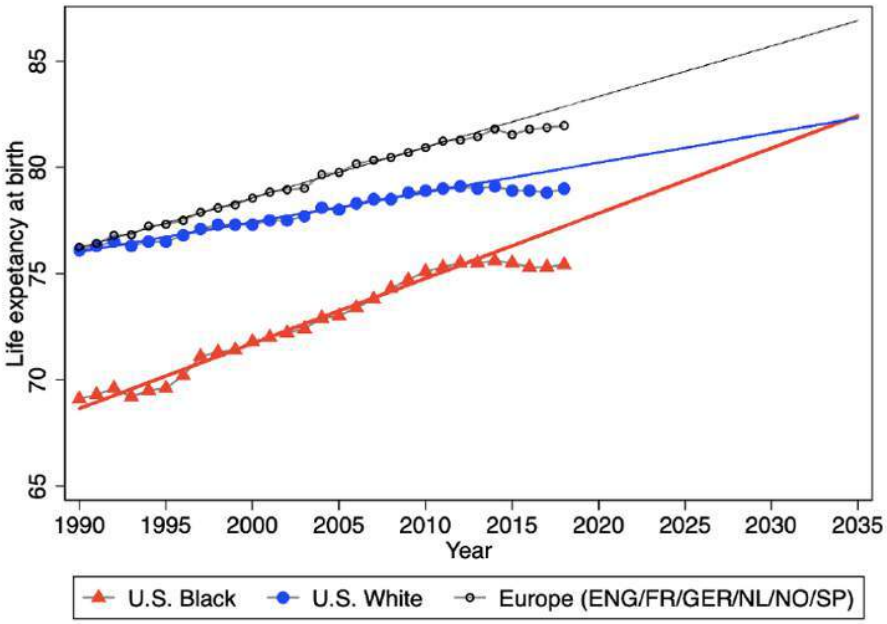

Figure 2: Closing life expectancy gap between Black and White Americans

Currie, Janet. "Life Expectancy Gap between Black and White Americans Closes Nearly 50% in 30 Years." Princeton School of Public and International Affairs, 2021, https://spia.princeton.edu/news/life-expectancy-gap-between-black-and-white-americans-closes-nearly-50-30-years.

Continuing Racial Equalities and Limitations

While the Civil Rights Act of 1964 sought to remove discrimination entirely, racial disparities endure. The U.S. Equal Employment Opportunity Commission was founded by the CRA 1964 and has received over 2.75 million charges of discrimination, brought over 10,000 lawsuits for Americans, and recovered over $11 billion for victims of workplace discrimination (U.S. Equal Employment Opportunity Commission). The disparity between charges received and suits won is substantial, indicating that Black Americans continue to experience workplace inequalities. Moreover, despite the CRA 1964 promises of protection from education discrimination, school inequities persist. Namely, school districts serving predominantly Black and Brown students receive approximately $23 billion less in annual funding than districts serving mostly White students, despite serving the same number of students. This inequity pervades into economic adult life, as White college graduates have "seven times more wealth than Black college graduates" (Ray et al.). Economic inequality also affects businesses; Paycheck Protection Program loans issued during COVID-19 disproportionately excluded Black-owned businesses, and over 40% of Black-owned small businesses closed during the COVID-19 pandemic. While the Civil Rights Act worked to eliminate injustices, these forms of de facto segregation persist and negatively impact African Americans.

Conclusion

The Civil Rights Act of 1964 was made possible by mass political mobilization that pressured policymakers to respond to the issue of racial discrimination (Rosenberg 1150). The act effectively eliminated the Jim Crow-era of segregation, legally prohibiting discrimination in public accommodations, education, employment, and federally funded programs (Back). It improved wages, job opportunities, and increased managerial positions for Black men in addition to enhancing African American communities by providing greater access to public facilities (Bourne 1213). Significantly, the act raised awareness of systemic racial injustice and strengthened political participation among African Americans, particularly with the additional passage of the Voting Rights Act in 1965 (Shaw et al.). In short, African Americans experienced economic gains, increased social mobility, political empowerment, and improved life expectancy due to the CRA 1964. While the act aimed to rectify inequities, racial prejudice continues to impact Black communities. For example, perpetual racial disparities in socioeconomic status cannot be completely solved by improving education access, because a Black man who is equally as educated as a White man is still likely to experience a considerable wealth gap. African Americans continue to endure barriers in education, employment, and business ownership along with ongoing economic inequality (Ray et al.).

The outcomes of the Civil Rights Act of 1964 reflect its implications for racial politics. The political galvanization to achieve its passage demonstrates the importance of coalition building between influential individuals, such as political activists and religious leaders, and organizations with legal and social resources, such as the NAACP. Collectively, the future of systemic reform policies depends on the strength of political mobility (Rosenberg 1148). The act

also reveals the limits of legal reforms without sustained societal change; simply put, addressing de facto segregation is essential in eliminating racial inequality. Furthermore, the act exhibited that federal government intervention is principal in enforcing civil rights. Additionally, by providing federally enforceable protections, the CRA 1964 authorized individuals to bring claims for discrimination under federal law as opposed to relying on state courts (Shaw et al.). By passing the act, Congress exercised its power to institute a standard that prohibited discrimination, overriding state laws that permitted segregation (National Archives and Records Administration). This ultimately established a precedent for future federal interventions in civil rights issues. The Civil Rights Act of 1964 presents how coordinated activism and legal reform can advance racial equality, while also illustrating that inequities persist without broader societal and structural change.

Works Cited

Back, Christine J. “The Civil Rights Act of 1964: Eleven Titles at a Glance.” Congressional Research Service, 2020, https://www.congress.gov/crs-product/IF11705.

Bourne, Jenny. “‘A Stone of Hope’: The Civil Rights Act of 1964 and Its Impact on the Economic Status of Black Americans.” Louisiana Law Review, 2014, https://digitalcommons.law.lsu.edu/cgi/viewcontent.cgi?article=6475&context=lalrev.

Library of Congress. “The Civil Rights Movement.” U.S. History Primary Source Timeline, https://www.loc.gov/classroom-materials/united-states-history-primary-source-timeline/post-war-united-states-1945-1968/civil-rights-movement/.

Lowe, William C. “Civil Rights Act of 1957.” EBSCO Research Starters, 2023, https://www.ebsco.com/research-starters/history/civil-rights-act-1957.

NAACP Legal Defense Fund. “Southern Manifesto and Massive Resistance to Brown v. Board.” NAACP Legal Defense Fund, https://www.naacpldf.org/brown-vs-board/.

National Archives and Records Administration. “Brown v. Board of Education.” National Archives, 2022, https://www.archives.gov/milestone-documents/brown-v-board-of-education.

National Archives and Records Administration. “Civil Rights Act (1964).” National Archives, 2022, https://www.archives.gov/milestone-documents/civil-rights-act.

National Archives and Records Administration. “Plessy v. Ferguson.” National Archives, 2022, https://www.archives.gov/milestone-documents/plessy-v-ferguson.

Ray, Rashawn, Andre M. Perry, David Harshbarger, Samantha Elizondo, and Alexandra Gibbons. “Homeownership, Racial Segregation, and Policies for Racial Wealth Equity.” Brookings Institution, 2021, https://www.brookings.edu/articles/homeownership-racial-segregation-and-policies-for-racial-wealth-equity/.

Rosenberg, Gerald N. “The 1964 Civil Rights Act: The Crucial Role of Social Movements in the Enactment and Implementation of Anti-Discrimination Law.” University of Chicago Law School, 2004, https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?article=2924&context=journal_article.

Shaw, Todd, et al. Uneven Roads: An Introduction to U.S. Racial and Ethnic Politics. 3rd ed., CQ Press, a division of SAGE Publishing, 2023.

Stanford King Institute. “Birmingham Campaign.” Stanford King Institute, https://kinginstitute.stanford.edu/birmingham-campaign.

U.S. Equal Employment Opportunity Commission. “Title VII of the Civil Rights Act of 1964: Requiring Discrimination-Free Workplaces for 60 Years.” U.S. Equal Employment Opportunity Commission, https://www.eeoc.gov/title-vii-civil-rights-act-1964-requiring-discrimination-free-workplaces-60-years.

Providing a Suitable Explanation For The Abandonment of Fatehpur Sikri
By Krishiv Loiya

Abstract

Emperor Akbar had had three offspring before, but all three—the daughter and the twin sons—had passed away during infancy. Desperate, he asked the saints for a prophecy. Sheikh Salim Chisti was the most prominent Sufi Saint of his time, and when Akbar consulted him, he prophesied that Akbar would have sons (Abu'l-Fazl, Akbarnama). After Jahangir (Salim) and two other sons were born, Akbar decided to create the city of Fatehpur Sikri in honor of the Sufi Saints and to show his devotion. Fatehpur Sikri was the capital of the Mughal Empire from 1571 to 1585 under Akbar's rule (Rezavi). It featured a blend of Hindu, Mughal, and other architectural styles prevalent across India. One of its most prominent features was its hydro-engineering prowess. It featured numerous water tanks, dams, 13 major step wells, and even an artificial lake. However, with just fourteen years of serving as the capital, it is the city that held the seat of power for the shortest time. This becomes even clearer when compared to Delhi or Agra, as Agra functioned as the empire's center for over a hundred years and Delhi remained the headquarters for over five hundred years. But perhaps the most surprising part is that, even after nearly four centuries, the reason for the abandonment of Fatehpur Sikri remains debated.

Fatehpur Sikri, at its peak, held over a quarter of a million people, was compared to London by European travelers, and synthesized cultures from across the Indian subcontinent. Yet, it was abandoned in a short time span of just 14 years. There have been several theories for the abandonment, ranging from Oriental Despotism to local folk tales about a dancer getting unjust treatment. However, the two most academically sound theories are that it was abandoned due to a water shortage, which was recently backed by a 2022 study that proved a three-decade mega-drought at the time of Fatehpur Sikri’s abandonment, while the other argues that it was abandoned due to increasing political tensions in the Northwest frontiers. This paper combines both theories: Fatehpur Sikri was abandoned due to increasing political tensions, but the reason Akbar never returned was due to a rapidly increasing water shortage, and uses both paleoclimatic and mathematical data to strengthen this idea.

This paper uses a tripartite methodology: paleoclimatic data, historical primary sources, and a mathematical stress test of the engineering of Fatehpur Sikri to conclude that it was abandoned for political reasons, but that Akbar didn't return after tensions subsided due to hydrological limitations.

Introduction

For decades, historians have advanced several theories for the abandonment of Fatehpur Sikri, with two emerging as prominent: the first is that it was due to a severe water shortage, while the second is that it was due to rising political tensions in Kabul. But a major flaw in both of these theories is the fact that they study their supporting evidence in isolation. The water shortage theory doesn't account for the fact that Fatehpur Sikri had one of the most advanced

hydrological systems to ensure that the city could sustain a peak population of a quarter million people (Rezavi). The increasing political tension in Kabul overlooks the fact that, after Kabul was resolved, Akbar decided to settle in Agra rather than Fatehpur Sikri. This paper argues that the reason for the abandonment of Fatehpur Sikri was increasing political tensions in Kabul; however, the reason why Akbar decided not to return to the city after tensions were resolved was due to a severe water shortage caused by a megadrought that lasted three decades. The discovery of this megadrought was made through recent scientific advancements, where Dr. Gayatri Kathayat and her team measured the ratio of two oxygen isotopes, ^{18}O and ^{16}O, using the growth of cave stalagmites in Mawmluh cave to deduce several climatic changes, including a three-decade catastrophic megadrought in a region including, but not limited to, Fatehpur Sikri (Kathayat et al.). Furthermore, this research paper uses objective, quantitative data to run a mathematical stress test on the city to deduce that it couldn't have survived the drought while containing the population of a Mughal Imperial Court.

The water shortage theory was the older of the two frameworks and has been repeated countless times by older historians and modern-day tour guides. The reason this framework has persisted even when modern historians have provided data against it is that several factors are at play, but the most prominent is that tourist guides favor the more dramatized answer over the more ‘bland’ geopolitical one. The water shortage framework has even led to a famous piece of folklore prominent in the region: that the water of Fatehpur Sikri dried up overnight due to a curse by a court dancer named Zarina, who was falsely accused of theft in the Mughal court. However, these are traditional stories preserved by oral history, not evidence. This framework originated because early European travelers misread the landscape: when travelers like Jean-Baptiste Tavernier encountered the empty sandstone structures that constituted the hydrological system, they assumed a great natural disaster had hit the city (Tavernier).

Colonial Despotism and Leaving ‘On A Whim’

The rediscovery of Fatehpur Sikri occurred in two phases: first through the structural documentation of the sandstone by the Archaeological Survey of India under Edmund W. Smith in the 1890s (E. Smith), and later through physical stabilization efforts initiated during Lord Curzon's viceroyalty in the early twentieth century. Fatehpur Sikri was originally used to demonstrate ‘oriental despotism’ and ‘imperial caprice,’ with the leading theory given by prominent historians of that time, such as Vincent Smith, being that Fatehpur Sikri was abandoned by Akbar on a whim with no particular reason (V. Smith). Oriental despotism refers to a dominant 19th-century British framework designed to make Eastern Kings appear as irrational and driven by personal whim rather than logic. British historian Vincent Smith and the colonial administration were among the first to advance the theory of Akbar leaving on a whim. This narrative codified Fatehpur Sikri's abandonment as ‘imperial caprice’ or a ‘royal whim.’ The British used this narrative as it showed Mughal decisions as chaotic and irrational; this further painted the British Raj as the introduction of logical, data-driven governance to India.

Furthermore, the city in its lifespan has gone from a Mughal capital, an abandoned remnant, a colonial ruin, to finally a UNESCO World Heritage Site in 1986 (UNESCO). Resolving the debate on its abandonment reclaims Mughal agency from colonial myths once again, as travelers and tour guides still retell the story that the site was abandoned due to a water shortage instead of giving a nuanced perspective.

However, as historians like Professor Syed Ali Nadeem Rezavi have pointed out, Fatehpur Sikri was one of the most prominent hydrological wonders of the Mughal Empire; Babur had specifically chosen the land because, in his words, it was 'well watered,' and Akbar had taken that several steps further by creating massive hydrological systems like the above-mentioned thirteen major stepwells and a massive artificial lake (Rezavi; Babur). These conclusions are the result of limited resources. The British colonial administrators had the travel accounts and physical site but didn't conduct archaeological excavations; Rezavi conducted an archeological expedition but was limited because the discovery of the megadrought hadn't occurred yet.

Before we move on to the second theory, we must understand a key term: in Mughal times, the definition of a capital was completely different; in those times, the capital city was wherever the emperor, and by extension the imperial court, was. Relocation of the capital isn't an anomaly that would require a single breaking point for abandonment. This is reinforced by the Jesuit priest Father António Monserrate, who traveled with Akbar and marveled at the sheer scale of the mobile capital (Monserrate 75–79), which was officially known in Mughal administration as Urdu-e-Mualla, or the 'Exalted Camp.' The Mughals went as far as hiring thousands of workers just to set up the Emperor's red tent called the Lal Qamar in a process they called the Farash Khana; this tent housed two hundred and fifty thousand people, over twenty thousand cavalry, and over five thousand elephants.

Richards's Political and Ideological Theory

The second theory is John F. Richards's political and ideological theory for Fatehpur Sikri's abandonment (Richards). He highlighted that before Fatehpur Sikri, the Mughal camp was highly mobile, moving in camps, concluding that Fatehpur Sikri was a transition between camp and imperial city.

Secondly, he argued that the city was built around the spirituality of the Chistis, and that when Akbar, in later times, began developing a new form of religion centered on himself, called Sulh-i-kul and Din-i-Ilahi, abandoning Fatehpur Sikri was a symbolic break from the Chisti culture.

Thirdly, he argued that during the fourteen years of capital rule that Fatehpur Sikri had had, it was actually a shared capital with Agra. He points out that most of the imperial treasury and arsenal were in Agra's Agra Fort, with Fatehpur Sikri only serving as a symbolic and cultural capital; Akbar could abandon it on a moment's notice.

However, the administrative records show that after his expedition to the Northern frontier due to rising tensions in Kabul, he had wanted to come back to Fatehpur Sikri; if he had

abandoned it to free himself from the Chisti spiritual legacy, he wouldn't have wanted to come back (Abu'l-Fazl, Akbarnama).

Furthermore, archeological excavations led by S. A. N. Rezavi revealed that Akbar had created gigantic non-portable infrastructure using a huge amount of state wealth (Rezavi). This included massive stone bazaars, specialized industrial workshops called karkhanas, and the monumental hydraulic system. This directly goes against the claim that Akbar had built Fatehpur Sikri as a transitional phase or a second choice when compared to Agra.

Another thing Richards had overstated was that Akbar had abandoned the Chisti Sufis in favor of his own Din-i-Ilahi, but this can be rebutted by the fact that Akbar's reverence for the Chistis did not end in 1585; even in later capitals like Agra, he regularly made pilgrimages to Sufi sites, like his pilgrimage to Ajmer Sharif, which he visited even after moving to Lahore and even Agra. Furthermore, Akbar went so far as to appoint the Chisti lineage to politically powerful positions through either granting them jagirs like the Ghazipur Jagirs or granting them high military or governing ranks, turning them into powerful aristocrats, or mansabdars. This is also strengthened by the fact that Akbar's doctrines make it clear that Din-i-Ilahi wasn't a mass religion but something limited to about twenty nobles (Abu'l-Fazl, Ain-i-Akbari); this directly undercuts Richards's claim that Akbar left Fatehpur Sikri due to his religious beliefs. Twenty people do not require the relocation of an entire capital.

Another important factor to note here is the historiographical narrative shifting with eras. Early European travelers like Jean-Baptiste Tavernier saw an abandoned city and assumed that Fatehpur Sikri was a ‘ghost town’ that was likely hit by a catastrophic natural disaster; paired with the fact that Fatehpur Sikri was built on a relatively arid area, this led them to believe that Fatehpur Sikri was abandoned due to a massive water failure (Tavernier). As this became more prominent, folklore like the myth of Zarina was created by the locals of the area as a form of entertainment and to explain the unknown. Later British officials saw the city and codified the water failure theory, even going as far as to call the city's abandonment a ‘royal whim’ by Akbar (V. Smith). But by the mid- to late 20th century, historians like Rezavi uncovered the architecture and created systematic counterarguments to the water failure theory, highlighting the hydrological masterclass of the city to show that a water failure was highly improbable (Rezavi). But by the early 2000s, there weren't just counterarguments to the water shortage theory; there was a counter theory. This brings us to the final, and most accepted theory, that Akbar had abandoned Fatehpur Sikri due to rising tensions in Kabul, a theory given by Satish Chandra (Chandra). This argues that Akbar had abandoned Fatehpur Sikri due to rising tensions in the North-Western Frontier, specifically due to the death of his brother, Mirza Hakim, and the rapid expansion of Abdullah Khan Uzbeg, the leader of the Shaybanid Empire in Bukhara. If Uzbeg annexed Kabul, he would have full control over the Khyber Pass, which at that time was the only direct highway to India.

In 1585, Akbar's brother Hakim died in Kabul due to alcohol poisoning, leaving the state completely defenseless (Abu'l-Fazl, Akbarnama 703). He had been ruling Kabul as a semi-independent state and had always been problematic for Akbar, even going so far as to try to

invade India once. His death created a massive power vacuum, as Kabul acted as a buffer state between India and the North, and rulers like the Shah of Persia and, more importantly, the ambitious Abdullah Khan Uzbeg had set their eyes on it. Knowing the risks it would bring to India if Kabul was captured by enemy forces, Akbar decided to take charge and officially annex Kabul into his empire; furthermore, he shifted his capital to Lahore, so he could directly supervise the entire operation. A key detail here is the fact that Fatehpur Sikri remained functional even after Akbar departed for several years, as mentioned by many European travelers. This is supported by several factors, including the fact that the Akbarnama mentions the death of Mirza Hakim and also notes that Akbar left Fatehpur Sikri less than a month after his death; as Abu'l-Fazl states, this was more of a royal march to the frontiers of Punjab than an abandonment of the capital, with some of the royal mint and some record keepers staying behind in Fatehpur Sikri, showing Akbar's intent to return once the conflict in Kabul was resolved (Abu'l-Fazl, Akbarnama). This also solves a major issue that had been faced by the water-shortage theory: the fact that Akbar shifted to Lahore for 13 years, when, if water shortage was the problem, he could have simply moved to Agra, which has an effectively inexhaustible water supply due to the Yamuna River. However, we must also take into consideration recent discoveries like the three-decade megadrought discovered in 2022 (Kathayat et al.). Previous historians dismissed the water shortage theory because they were restricted to the evidence available in their time, but now, with the discovery of the megadrought, the water shortage theory has regained prominence. This also aligns with the travel accounts of European travelers like Ralph Fitch and William Finch (Foster). William Finch goes so far as to explicitly state that Fatehpur Sikri was already in decline by the time he had arrived; however, we must also note that the accounts of Fitch and Finch are accounts of independent European travelers writing decades apart, pointing to a gradual decline rather than an instant collapse.

The Hydrological Prowess of Fatehpur Sikri

On the flip side, we must also truly understand the sheer magnitude of the hydrological prowess Fatehpur had attained. First and foremost was the great lake basin near Hiran Minar with a circumference of nearly twenty kilometers (Rezavi). It collected the monsoon water and was used to replenish local aquifers and regulate the water table of the area. From there, millions of liters of water were lifted forty meters using Persian wheels called rahats. These Persian wheels worked because of blindfolded bullocks turning wooden blocks, which in turn lifted water-filled terracotta pots upwards. The city was divided into two systems: the Northern waterworks at Hathipol and the Southern waterworks near Hakim's Bath. The Hathipol, or elephant gate system, was a raised platform that required the water to be carried upwards over forty meters; this was done using blindfolded bullocks to turn in circles, which rotated a plethora of gears, which rotated a vertical chain of terracotta pots called ghats. However, the forty-meter ridge was very rocky and required the water to be lifted in three different stages from the lake to the overhead distribution aqueducts. On the other hand, Hakim's Bath waterworks was dependent

on a massive four-storied stepwell. Using a series of rahats, the water was lifted from it and used personally for the khwabgah and hammams.

Another thing to note here is the non-hydrological architecture of the city and how it affected Akbar. First and foremost was the gigantic fifty-four-meter-long triumphal gate known as Buland Darwaza (E. Smith). This was added after the successful campaign in Gujarat, representing global ambition.

Secondly was Jama Masjid, where the white marble tomb of Sheikh Salim Chisti was situated, serving as the spiritual heart and primary legitimacy anchor for Akbar's reign in the city.

One of the most important buildings of the entire city was the Ibadat Khana, also known as the House of Worship. Akbar gathered believers of rival religious traditions and told them to discuss their beliefs with each other in the Ibadat Khana. What he wanted to be an idealistic conversation about development often turned into bitter arguments, which pushed Akbar to create Din-i-Ilahi. This also directly refutes John F. Richards, as this architecture clearly shows that Din-i-Ilahi wasn't created before Fatehpur Sikri, but rather, after it.

However, as brilliant as Fatehpur Sikri was, both culturally and hydrologically, there were two fatal flaws in this system; the first was the fact that even though storage was maintained by the artificial lake, which provided water all year round, the active transport was heavily dependent on the rahat wheel system; if the wheel stopped, the gravity-fed fountains, cooling channels, and hammams would dry out. Secondly, although the system was hybrid, combining both rainwater harvesting and baolis, it couldn't depend on groundwater, as the groundwater in Fatehpur Sikri is high in salt content, making it dependent on rainwater to the point where almost every rooftop was used to maximize the collection of rainwater runoff. Compared with other capitals like Delhi and Agra, the differences in hydrological systems become glaringly clear. While Fatehpur Sikri used an intricate yet rain-dependent system to sustain itself, Delhi and Agra used a more traditional approach, sustained by the glacier-fed Yamuna River; a benefit of a glacier-fed river was that if rainfall was below average for a prolonged time, the city could still sustain itself without issue.

Even then, under normal conditions, this system of artificial lakes, open ponds, and channels was designed so brilliantly that it could fill the public and royal baths and sustain a peak population of a quarter million people in a relatively dry climate under India's scorching heat.

However, from 1580 to 1610, it wasn't a normal time.

The Megadrought and a Mathematical Analysis

By examining the oxygen isotope ^{18}O in stalagmites and speleothems from India, scientists were able to reconstruct the history of the Indian summer monsoon (Kathayat et al.). They found that a lighter oxygen isotope ratio meant heavier rain, while a heavier ratio meant drought. Using this data, they found that there was a brutal and severe megadrought in a region including, but not limited to, Fatehpur Sikri from 1580 to 1610; this window of time also matches Akbar's departure.

To find out whether this drought affected the city severely enough that Akbar had to change capitals, we have to look at the details of the hydrological system.

First and foremost, the gigantic artificial lake was created by a 2.4-kilometer-long dam, spanning roughly 3.2 kilometers in length and 1.6 kilometers in breadth, with an estimated depth of 3 meters (Rezavi). We also need to take into account the fact that during these times there was no groundwater replenishment, and in a semi-arid region like Fatehpur Sikri, water evaporated at an estimated rate of roughly two meters annually. This leads us to the following equation:

$$5.12 \times 10^6 \text{ m}^2 \times 2.0 \text{ m} = 10.24 \times 10^6 \text{ m}^3$$

This shows that even in perfect conditions, without monsoon, the lake loses nearly 66 percent of its water to evaporation. Substantial rainfall was almost required for the city to sustain its estimated population.

Furthermore, excavations by Rezavi uncovered that the artificial lakebed and lower stepwells had a thick, undisturbed layer of silt (Rezavi). This is evidence that Akbar's departure also removed the treasury-funded workforce needed to dredge the lake, since ancient lakes require laborers to remove the accumulated mud. Assuming a siltation rate of 1.5% capacity loss per year, over the 13 years between Akbar's departure and the end of conflict in the northern frontiers (1585 to 1598), the lake would have shrunk from a depth of 3 meters to 2.4 meters. A shallower lake means both a smaller water capacity and a faster rate of evaporation. This leads us to the final calculation needed to assess whether Akbar could have returned to the city after his northern campaign.

Maximum volume accounting for silt:

$$5.12 \times 10^6 \text{ m}^2 \times 2.4 \text{ m} = 12.28 \times 10^6 \text{ m}^3$$

Maximum volume accounting for evaporation:

$$12.28 \times 10^6 - 10.24 \times 10^6 = 2.04 \times 10^6 \text{ m}^3$$

An imperial camp with a peak population of a quarter million people, an estimated twenty thousand horses, and an estimated five thousand elephants consumes roughly 2.40×10^6 m^3 of water annually (calculated using standard baseline biological estimates of 150 liters daily per elephant, 40 liters daily per horse, and basic survival rations for humans). This equation shows that the emperor could not possibly have returned to the city after his northern campaign, because the city would not have been able to provide enough water to sustain the entire court.

We must also reinforce the mathematical calculations to show that no matter how generous the assumptions, the city still fails to sustain itself.

An important way to do this is to recount the fact that for 13 years, the siltation in the lake was growing steadily, as there were no treasury-dependent laborers to dredge it. We have already shown that a siltation rate of one and a half percent per year or more is enough to cause the failure of the city; however, the siltation rate could have been as low as one percent per year. In that case, the calculation would look like this:

13 years at a 1.0 percent per year annual loss = 13% capacity loss.

Estimated average depth of the artificial lake = 3 meters.

13% of 3 = a 0.39-meter drop from the initial depth; this results in the lake being 2.61 meters deep.

1.0% siltation results in a depth of 2.61 meters; when factoring in a conservative lower-bound evaporation estimate of 1.8 meters (versus the paper's baseline 2.0 m/year rate), the lake is left with only 0.81 meters of standing water on average.

Furthermore, the usable water is even less than this, because we must also take into account the physical restrictions of the rahat wheel-powered hydrological system. The Persian wheels required a minimum clearance of approximately 0.5 meters above the mud to function (Rezavi).

$$0.81 - 0.50 = 0.31 \text{ meters of usable depth}$$

$$5.12 \times 10^6 \text{ m}^2 \times 0.31 \text{ m} = 1.58 \times 10^6 \text{ m}^3 \text{ of extractable water}$$

When factoring in the fact that the water required by the entire Imperial court was approximately 2.40×10^6 m³, even the 1.0% siltation rate leaves the city massively short of the required amount needed to sustain the Mughal Imperial Court. These calculations show that even under the most conservative assumptions, Fatehpur Sikri still falls short.

Moreover, this matches European travel accounts of the city surviving, yet declining into ruin after Akbar's departure, since the remaining population of the city was about twenty-five thousand local weavers, who would have required less than ten percent of the water the city generated at capacity.

This finally leads us to the reason why Akbar chose Agra over Fatehpur Sikri after resolving the crisis in Kabul. As we have already calculated, Fatehpur Sikri couldn't have survived the three-decade megadrought; this doesn't apply to Agra, as Agra had the massive advantage of the glacier-fed Yamuna River. Even when the summer monsoons completely failed, the river continued to flow, since it drew its water directly from the Himalayas. Moreover, the Yamuna River provided substantial trade and military benefits, as it was one of the major trade routes of the time; pair this with the fact that Akbar didn't even have to construct anything new, since the Agra Fort already existed and, as previously stated, housed most of the armory.

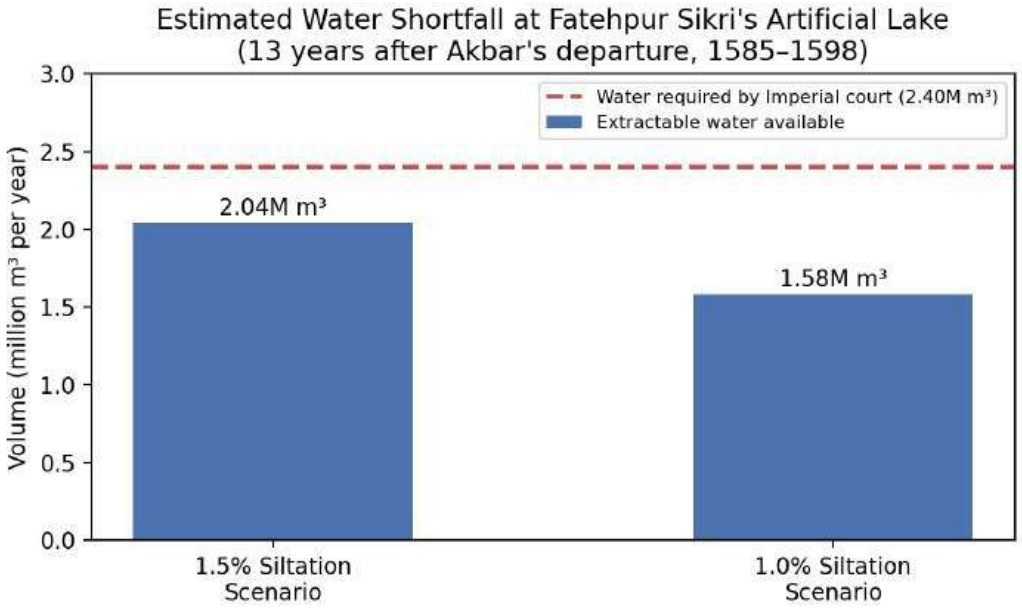

Fig 1: Estimated annual water shortfall at Fatehpur Sikri's artificial lake under two siltation scenarios, compared to the Imperial court's required consumption.

Global Connection

We must also understand that the downfall of Fatehpur Sikri isn't necessarily an isolated Mughal hydrological failure; rather, it fits within a broader global pattern of cities with immense hydrological prowess falling to severe climate-related catastrophes. One of the most prominent examples is Angkor in Cambodia, as shown in a study by Buckley, using tree-ring 'proxy records' created by Fokienia hodginsii cypress trees (Buckley et al.). Buckley used dendroclimatology to show that in the 14th and 15th centuries, Angkor suffered a massive megadrought, which starved its artificial water reservoirs. Just like Fatehpur Sikri, Angkor was a hydro-engineering marvel destroyed by rapid climate change, specifically a megadrought that overwhelmed its rain-dependent water supply. This further shows that Fatehpur Sikri's abandonment reflects a recurring pattern globally, not an isolated Mughal failure.

Another major site abandoned due to a lack of rainfall is the Classic Maya lowlands of Mesoamerica; research by Hodell shows that, using oxygen isotope ratios in stalagmites and speleothems, scientists were able to identify a severe drought coinciding with the decline of major Maya urban centers (Hodell et al.). This connection is important because it shows that speleothem-based paleoclimatology is an established method for identifying climate-driven urban abandonment, not a technique unique to this paper's argument.

Both of these examples show us the same reality as Fatehpur Sikri: a monumental city built with advanced hydrological engineering to manage seasonal water scarcity, a historic megadrought that exceeded the infrastructure's maximum capacity, and a strategic abandonment. We must also address the fact that even after Akbar's formal abandonment, Fatehpur Sikri remained operational as a home for local weavers. From European travel accounts, the estimated

population of Fatehpur Sikri after abandonment was roughly twenty-five thousand, about ten percent of the estimated population at its imperial peak. The surviving baolis and remnants of the hydrological system could comfortably support a small artisan community, but not the entire Mughal imperial court.

Limitations

This leads us to the limitations of this explanation. First and foremost is the limitation of the study showing the three-decade-long megadrought; oxygen isotopes can show a general drought trend over thirty years, but they cannot yield specific year-by-year rainfall estimates. This could be resolved by attaining higher-resolution regional proxy records. Secondly, the evaporation and siltation rates used in the arithmetic calculations are historically grounded estimates, not the exact rates from the sixteenth century, which cannot be recovered directly. This could be resolved through direct sediment-core sampling at the lakebed of Fatehpur Sikri's artificial lake to find the siltation rate; the exact historical evaporation rate will likely remain an estimate. The last limitation of this paper is that, even though there was a regular month-by-month census of the Imperial camp while Akbar resided there, there is no exact month-by-month census of the population of Fatehpur Sikri after his departure. This could be resolved by Persian tax or administrative records of the city, if they survive.

Conclusion

So, in conclusion, Fatehpur Sikri's abandonment was due to rising political tensions in Kabul, caused by the threat of invasion by Abdullah Khan Uzbeg, paired with the death of Akbar's brother, Mirza Hakim. However, the reason Akbar didn't return to Fatehpur Sikri once political tensions had been resolved was that the city lacked enough water to sustain the entire imperial court. The reason for this lack of water was a brutal, three-decade-long megadrought, which scientists recently uncovered by studying the oxygen isotopes present in stalagmites in Mawmluh Cave (Kathayat et al.). To further reinforce this, this paper uses historically grounded estimates to mathematically demonstrate that the city couldn't sustain the water required for an entire Mughal imperial court, and formally connects Gayatri Kathayat et al.'s 2022 isotope paleoclimate discovery with sixteenth-century Mughal administrative historiography and the Kabul geopolitical crisis. Comparing it to the glacially sustained cities of Agra and Delhi, we can note that those cities could have sustained themselves through the drought because of the Yamuna River; furthermore, in Mughal times, the capital city was wherever the emperor and his highly mobile imperial court were. This paper also points out a broader pattern of hydrologically advanced cities being abandoned due to severe climatic shocks, drawing on the Maya lowlands and Angkor specifically to draw similarities with Fatehpur Sikri. It is worth noting, however, that Fatehpur Sikri remains one of the most prominent examples of the sheer hydro-engineering and architectural mastery the Mughals possessed.

Works Cited

Abu'l-Fazl ibn Mubarak. The Akbarnama of Abu-l-Fazl. Translated by Henry Beveridge, vol. 3, Asiatic Society of Bengal, 1897–1939.

---. The Ain-i-Akbari. Translated by H. Blochmann, Asiatic Society of Bengal, 1873.

Babur, Zahiruddin Muhammad. Baburnama. Translated by Wheeler M. Thackston, Modern Library, 2002.

Buckley, Brendan M., et al. "Climate as a Contributing Factor in the Demise of Angkor, Cambodia." Proceedings of the National Academy of Sciences, vol. 107, no. 15, 2010, pp. 6748–52.

Chandra, Satish. Medieval India: From Sultanat to the Mughals: Mughal Empire (1526–1748), Part II. Har-Anand Publications, 2005.

Foster, William, editor. Early Travels in India, 1583–1619. Oxford UP, 1921.

Hodell, David A., et al. "Possible Role of Climate in the Collapse of Classic Maya Civilization." Nature, vol. 375, 1995, pp. 391–94.

Kathayat, Gayatri, et al. "Protracted Indian Monsoon Droughts of the Past Millennium and Their Societal Impacts." Proceedings of the National Academy of Sciences, vol. 119, no. 39, 2022, e2207487119.

Monserrate, António. The Commentary of Father Monserrate. Translated by J. S. Hoyland, Oxford UP, 1922.

Rezavi, Syed Ali Nadeem. Fatehpur Sikri Revisited. Taylor & Francis, 2014.

Richards, John F. The Mughal Empire. Cambridge UP, 1993.

Smith, Edmund W. The Moghul Architecture of Fathpur-Sikri. Archaeological Survey of India, 1894–98.

Smith, Vincent A. Akbar: The Great Mogul, 1542–1605. Clarendon Press, 1917.

Tavernier, Jean-Baptiste. Travels in India. Translated by Valentine Ball, Macmillan, 1889.

UNESCO World Heritage Centre. "Fatehpur Sikri." whc.unesco.org/en/list/255/.

PROK2 and FGFR1 Mutations Responsible for HPG Disruption leading to Kallmann Syndrome in Women By Ashni Bhupathiraju

Abstract

Kallmann Syndrome (KS) has many underlying mechanisms such as genetic mutations that cause neuroendocrine dysfunction. These mutations are associated with symptoms such as anosmia and hypogonadism, hence leading to infertility and loss of smell in men and women. However, KS is 4 to 5 times more prevalent in males (1:8000) than in females (1:40000). While existing literature has identified Fibroblast Growth Factor Receptor 1 (FGFR1) and Prokineticin 2 (PROK2) as genes linked to KS, there is nominal understanding of their distinct roles in the disruption of Gonadotropin-releasing hormone (GnRH) neuron and hypothalamic-pituitary-gonadal (HPG) axis functioning, as mechanisms of female infertility. Existing literature studies FGFR1 and PROK2 in acting as separate causative genes of KS, however both genes work together to affect the function of the GnRH neuronal system through its different stages. The aspect that is minimally developed is how these pathways could ultimately converge to cause GnRH dysfunction and female infertility. Thus, this review synthesizes mutations in PROK2 and FGFR1 by asking the question: How do FGFR1 and PROK2 mutations functionally converge to disrupt GnRH neuronal development and neuroendocrine signaling, contributing to female infertility in Kallmann syndrome?

As GnRH neurons regulate the secretion of reproductive hormones including Follicle Stimulating Hormone (FSH) and Luteinizing Hormone (LH), it is essential for their development and migration to be systematic. For this functioning, FGFR1 primarily plays a role in the development of GnRH neurons, while PROK2 focuses on their migration to the hypothalamus during embryonic development. Furthermore, if these genes are mutated, it causes neuroendocrine dysfunction to the HPG axis. Hence, understanding the mechanisms of these mutations and their effect on neuroendocrine dysfunction is crucial in decoding the link between GnRH development and hormone secretion. This disruption not only causes infertility and anosmia in women, but its understanding also expands our knowledge of reproductive neuroendocrinology and can support the future development of more targeted treatments for infertility.

Introduction

Infertility, defined as the failure to achieve pregnancy after 12 or more months of regular unprotected sexual intercourse, affects nearly 1 in 7 couples worldwide, who experience the inability to conceive naturally (Sang et al.). Kallmann Syndrome (KS) is a rare heterogeneous genetic disorder in which affected individuals experience delayed or absence of puberty and reduced or absence of smell, leading to symptoms including anosmia and hypogonadism (Kumar Yadav et al.). In women, the causation of certain infertility includes endopathies such as hypothalamic amenorrhea and hyperprolactinemia (Concepción-Zavaleta et al.) however, neuroendocrine causes still remain underrepresented. FGFR1 and PROK2 stimulate GnRH

release and trigger the activation of the mammalian reproductive axis. The hindrance of normal episodic GnRH secretion is thought to be the primary cause of KS, as these neurons play a major role in the regulation of the HPG axis (Young et al.). The HPG axis which is responsible for puberty, sexual maturation, and reproductive abilities (Collado-Sole et al.), connects the brain, pituitary gland, and reproductive organs. Hence, a disruption to the HPG axis can lead to several reproductive constraints. These disruptions may transpire through a series of gene mutations, namely FGFR1 and PROK2. With one or more of these mutated genes present within an individual, GnRH dysfunction is likely to take place. As GnRH neurons act as the fundamental regulators of the reproductive axis, their dysfunction is detrimental to the overall reproductive health of women.

GnRH neurons that migrate to the hypothalamus during embryonic development are activated by Kisspeptin neurons. Kisspeptin-producing neurons comprise the major afferents to GnRH neurons. These afferents, which act to regulate the feedback control of GnRH secretion, can lead to a pre-ovulatory surge (Kim). Although existing literature highlights the form and function of GnRH neurons after birth, the importance of their role during embryonic development receives comparatively less attention. As the FGFR1 gene plays an essential role in development of GnRH neurons (Yellapragada et al.), mutations in this gene cause both forms of isolate GnRH deficiency (IGD), accounting for approximately 10% of the GnRH deficiency cases (Miraoui et al.). Mutations in FGFR1 and other genes such as KAL1 account for an estimated 10% of KS cases; therefore, a second gene that is investigated is PROK2 (Bianco and Kaiser). The absence of PROKR2, a primary G-coded receptor for PROK2, in mice has presented abnormal olfactory function and GnRH neuron migration (Martinez-Mayer and Perez-Millan). This indicates PROK2 as a primary gene associated with KS.

Over the past couple of decades, it has been established that mutations in PROK2 and FGFR1 are primary factors leading to KS. While studies examine the role of these two genes as causative genes in KS separately, there is minimal understanding of their synthesis and how they converge to bring upon symptoms of KS through the disruption of the HPG axis, especially in females. This review bridges the connection between both mutations and their collective role in the disruption of the HPG axis, causing KS. While reduced FGFR1 signaling triggers failed neurogenesis (Chung and Tsai), PROK2 affects the signaling pathways and migration of GnRH neurons. When understanding the converging signalling pathways of both FGFR1 and PROK2 genes, it provides insight into how the pathways provide complementary explanations of how GnRH functioning can be disrupted through different stages, leading to KS. Examining the convergence of both genes is instrumental in providing a biological insight into neuroendocrine infertility and advancing reproductive medicine in the future.

The Biological Basis of Fertility

Reproduction is vital for the survival and continuation of species. This life process is regulated and maintained by the HPG axis which consists of the anterior hypothalamus, pituitary gland, and gonads. To ensure the undisturbed function of this axis, GnRH neurons connect the

neuronal pathways of the brain and send projections (Xie et al.) through pulsatile secretion which stimulates the release of FSH and LH (Stamatiades and Kaiser). In standard female fertility, puberty takes place and, as a result, ovulation occurs. Ovulation, defined as the rupture of the dominant follicle of the ovary, releases an egg into the abdominal cavity. This process is governed by two primary reproductive hormones, FSH and LH. Roughly 1000 primordial follicles begin this process of maturation with the granulosa cell layer increasing in size, and production of estrogen commencing through FSH stimulation. FSH plays a role in beginning estrogen synthesis; however, this is later handled by granulosa cells as an autonomous process. As the cycle of follicle development progresses, one follicle called the "dominant follicle" secretes estrogen which wields negative feedback on FSH, hence causing all other follicles to "turn off". Continuing, the role of LH is mainly exhibited on day 14 of the menstrual cycle. As estrogen levels are high, it leads to an increased frequency of GnRH secretion, triggering a surge in LH levels. With LH levels at a peak, intrafollicular proteolytic enzymes increase, weakening the walls of the ovary and allowing for the follicle to be released (Holesh et al.). For couples attempting to conceive, this time period would be optimal as the follicle is exposed. Due to these governing hormones being dependent upon GnRH development and migration, when there is a disruption to the HPG axis due to genetic mutations, it often leads to KS in females through impaired follicular development.

As discussed, successful fertility cannot be attributed to a single hormone alone. Rather, there is a range of hormones that vary between males and females to facilitate conception. For this reason, GnRH dysfunction is detrimental to the HPG axis as it regulates the secretion of multiple of the primary reproductive hormones. GnRH neurons are secreted in an episodic manner through pulses and surges, allowing for the control of the pituitary gland and systematic reproductive function. Migration of these neurons through the neuronal pathways is divided into four distinct stages. These four stages demonstrate what "normal" GnRH functioning consists of. As mutations in FGFR1 and PROK2 are characterized to disrupt this functionality, it is important to comprehend what the systematic GnRH neuronal system would look like.

To begin, after birth in the olfactory placode, the next stage details a specific cue that guides the neurons across the cribriform plate. This plate is a portion of the ethmoid bone which is seen to be located on the base of the skull. The ethmoid bone extends over the nasal septum and contains narrow grooves known as the olfactory fossa. As the olfactory fossa plays a primary role in supporting the olfactory bulbs, when there is a mutation, it can often cause anosmia (Gomez and Pickup). After this crossing, they migrate towards the hypothalamus, finally detaching from their axonal guides and dispersing throughout the hypothalamus (Wierman et al.). Once established in the hypothalamus, GnRH neurons become the primary regulators of reproductive neuroendocrine signaling in the HPG axis. Collectively, GnRH neuron migration to the hypothalamus is the fundamental process that regulates the HPG. Thus, a disruption to any of these stages could impair the ability of GnRH to establish in the hypothalamus, hence causing neuroendocrine dysfunction.

To understand the functioning of the HPG axis, its parts and form should be established. The HPG consists of an intercommunicating set of tissue that function as the complex units of fertility regulation. The gonadotrope cell found in the anterior pituitary gland plays an instrumental role as the mediator between GnRH signal and the germ cell reservoirs. Additionally, it also helps in the steroid hormone productivity of the gonads. However, the gonadatrobe's primary function is to express the GnRH receptor (GnRHR). When GnRH is bound to its receptor, it triggers a plethora of intracellular signals which coordinate the response of these cells to stimulate GnRH and release LH and FSH (Bliss et al.). This demonstrates how the HPG, particularly the gonadotrope cell, is fundamental for female reproduction.

Evidence supports that it's crucial for GnRH to successfully reach the hypothalamus so the function of the HPG can be systematic. However, these stages of GnRH migration are only possible with the signals of guidance molecules including PROK2; therefore, a disruption to this molecule has profound consequences for the reproductive axis as a whole. Moreover, FGFR1 protein is a transmembrane receptor that plays an essential role in the development of GnRH as well as the migration of these neurons from the nasal placode. As a result, a mutation of this gene was first associated with KS in 2003 and is now a common gene involved with KS and Congenital hypogonadotropic hypogonadism (CHH), classified as the complete or partial failure of and infertility due to GnRH deficiency (Fadiga et al.). In addition, when GnRH fails to reach the hypothalamus due to HPG dysfunction it can lead to hypogonadotropic hypogonadism, a group of diseases related to hindrance in secretion and action of GnRH, in the anosmic form (Кокорева et al.). As Kallmann Syndrome is characterized by both CHH and anosmia, these dysfunctions are considered the disorder's underlying causes.

FGFR1 Mutations and their Role in Neuroendocrine Dysfunction

Fibroblast growth receptor 1 (FGFR1), also called KAL2, is a receptor tyrosine kinase (RTK), meaning it's a receptor present on the surface of cells that allows signal transmission from the outside to the inside of the cell. Variants of this protein have been identified in a majority of patients with KS. To detail, FGFR1 regulates many functions and processes including cell proliferation, differentiation, and migration of GnRH neurons. To carry out these processes, FGFR1 protein is encoded by FGFR1 gene which is primarily expressed in the Rathke's pouch and the ventral diencephalon (Soejima et al.). These two structures are parts of the developing embryo that eventually develop into parts of the endocrine system and brain. Although FGFR1 is a fundamental protein that allows for GnRH development and migration, it does not work alone. FGFR1 orchestrates fetal development through neuronal signalling via FGF receptors (Katoh et al.). A specific protein in this family is FGF8 which acts as the primary ligand (activating signal), which enables the activation of FGFR1 after they bind together. Additionally, FGF8 subfamily members also regulate embryonic development processes and are expressed in the heart, limbs, kidney, and central nervous system (Hao et al.). The structure of proteins in the FGF family is what makes them suitable for binding to FGFR1. This structure

consists of all FGFs having an internal core made up of 120-125 amino acids, hence allowing them to easily bind to FGFR.

FGFR1 not only regulates fetal development, it also plays a major role in the development of the nasal placode. The nasal placode, a region that brings about olfactory sensory neurons, infolds to form the primary olfactory epithelium and vomeronasal organ (structures where GnRH neurons begin to migrate from). The olfactory epithelium is a specialized tissue in the nasal cavity that is responsible for detecting odors, and the vomeronasal organ is another small sensory organ located in the nasal septum. After the odor molecules pass through these organs, they reach the olfactory bulb (OB) in the brain that processes the smell. FGFR1 mutations have been identified to disrupt this function. This was tested by studies conducted on mice which show that loss of FGFR1 results in disruption of olfactory bulb development and reduction in GnRH neuron count (Cho et al.). The OB is another crucial component as it acts as the central location for sense of smell and is the region of first synaptic relay involved in processing odor information (Blakemore and Trombley). FGFR1 genes are expressed within these organs, which are essential, because they play a role in odor processing, forming migration pathways, and developing alongside GnRH neurons. Studies conducted in mouse embryos have shown that GnRH mRNA and peptide-expressing neurons are first detected as the medical olfactory placode. Later studies additionally concluded that the neurons migrate along the peripherin-positive olfactovomeronasal fibers through the cribriform plate into the hypothalamus. These studies display how disruption of signalling may cause impairment to the GnRH system even before neuronal migration into the forebrain (Chung et al.).

The location of FGFR1 expression in regions such as the OB provides a basic insight into the gene's association with KS. Since the development of GnRH takes place in the olfactory bulb and surrounding nasal regions before migrating to the hypothalamus, FGFR1 present in these areas may be crucial in initiating the formation of the GnRH neuronal system. This explains the role of FGFR1 in possibly affecting the fundamental process required for GnRH development in the nasal placode. When FGFR1 is found to be mutated, its detrimental effects on GnRH development consequently disrupts the HPG axis, leading to abnormalities in standard ovarian function. Therefore, the expression of FGFR1 in specific regions of the brain is relevant in establishing how GnRH dysfunction due to genetic mutations can contribute to the underlying mechanisms of female infertility in KS, by explaining the effects of GnRH disruption on the secretion of governing reproductive hormones.

To continue, the role of FGFR1 as a fundamental part of maintaining the function of the HPG has been established within many studies. When the FGFR1 gene was examined in seven idiopathic hypogonadotropic hypogonadism (nIHH) subjects, heterozygous FGFR1 mutations were identified in three of the seven. These specific results showcased four rare variants of FGFR1 gene that were found in three different unrelated probands. These variants alter the amino acid sequence of FGFR1. For example, A in exon 6 was predicted to substitute a glycine for serine at the position 237 (G2375) within the extracellular region of FGFR1 (Pitteloud, Acierno, et al.). Furthermore, in 2003, FGFR1 was the first identified gene that was associated

with causing the autosomal dominant form of KS (Fadiga et al.). This was determined by mapping overlapping deletions on chromosome 8p11-p12 in two Kallmann Syndrome patients with contiguous gene syndromes (Dodé, Levilliers, et al.). The observation that two unrelated patients were identified with the disruptions within these regions of the chromosome led researchers to examine the genes located in this area. This is where they identified and determined FGFR1 as a contributor to GnRH deficiency.

To add, it is known that the FGF pathway is important for the development of the entire olfactory system. Thus, when transgenic mice, with targeted gene expression of FGFR1, exhibited delayed puberty and fewer count of GnRH neurons in the hypothalamus; it uncovered the role of FGFR1 in the olfactory system and GnRH ontogeny consistent within the patients of KS with the phenotype (Miraoui et al.). Together, these human genetic studies and mice studies confirm the role of FGFR1 in reducing GnRH neurons numbers, and ultimately contributing to KS. By disrupting migratory pathways and development of GnRH neurons, FGFR1 can be attributed to altering normal reproductive function.

Moreover, the connection between FGFR1 disruption and irregular reproductive function is specifically important to female infertility through KS because of FGFR1's major role in GnRH development. Overall, a disruption to FGFR1 neuronal signalling can meddle with the growth and movement of GnRH neurons, reducing the GnRH count that is established within the hypothalamus. Because of GnRH neurons' role in regulating the production and secretion of FSH and LH, this reduced neuronal population is likely to weaken downstream HPG function. To regulate specific reproductive hormones, GnRH binds to a membrane receptor on the pituitary gonadotrophs which stimulates the secretion of the intended hormone (FSH) (Kaprara and Huhtaniemi). This provides a link between GnRH and hormone secretion, governed by FGFR1. Thus, the phenotype of KS can be understood through the downstream process of HPG disruption.

GnRH migration via PROK2 Signaling

While FGFR1 plays a fundamental role in the successful development of both the OB and GnRH neurons, orderly reproductive function also requires undisturbed migration of GnRH neurons from the nasal placode to the hypothalamus during the embryonic stage. To understand the importance of this migratory pathway, the ligand-receptor relationship between PROK2 and PROKR2 should be understood. Prokineticin receptor 2 (PROKR2) encodes a 384 amino acids long G-protein coupled receptor (GPCR). PROKR2 and its related receptor, PROKR1, share 85% of the same genetic sequence, consequently, binding to the same ligands, PROKR1 and PROKR2. Once these ligands are bound to PROKR2, PROKR2 leads to the activation of several signaling pathways including MAPK, IP3, and cAMP signalling. As PROKR2 shows high expression in the OB, it is considered essential for the development of the olfactory system and migration of GnRH (Martinez-Mayer and Perez-Millan).

One particularly important signalling pathway to look into when beginning to understand the possible functional convergence between FGFR1 and PROK2 is the MAPK signalling

pathway. MAPK or mitogen-activated signalling protein kinase is a group of signalling pathways that regulate a wide range of cellular processes which ultimately activate downstream proteins (Guo et al.). The enzymes in these pathways help transport signals from outside of the cell to its interior; hence, determining how the cell behaves and functions. In relation to PROKR2, MAPK activates multiple signalling pathways after binding with PROK2. Studies investigating mutations within PROK2 have primarily focused on intracellular calcium release, ligand binding capacity, and expression of the receptor on the surface of the cell. These investigations help to determine how mutations can possibly hinder receptor activity and, ultimately, disrupt downstream cellular signaling (Dodé and Rondard).

Additionally, there is also experimental evidence that supports the association between MAPK-signalling and PROK2 activation. One specific experiment demonstrated that wild-type PROKR2 with increasing concentrations of PROK2 ligand produces sigmoidal dose-repose curves for Egr1-luciferase induction which actively serves as an indicator of MAPK-associated activity (Abreu, Noel, et al.). These findings illustrate how PROKR2 activation can elicit several signalling responses within cells, possibly explaining how MAPK signalling may directly influence GnRH migration during development.

GnRH neuronal migration follows a complex pathway commencing from their birth in the nasal placode to the mediobasal hypothalamus, where they complete their differentiation (Bassi et al.). This migration, however, is only possible with the signalling of guidance molecules, a prolific one being PROK2. The PROKR1 gene is located in the human chromosome 2p13.3 and mouse chromosome 6, whereas the PROKR2 gene was reported to be mapped in human chromosome 20p13 and mouse chromosome 2 (Martin et al.). However, more recent studies place human PROKR2 at chromosome 20p12.3 (*PROKR2 Prokineticin Receptor 2 [Homo Sapiens (Human)] - Gene - NCBI*). To study the mechanism behind PROK2 causing neuroendocrine dysfunction and disruption to the HPG axis, studies were conducted on mice. While initial investigations studied PROK2 in smoothening of muscle contraction of the gastrointestinal system, angiogenesis, etc., studies conducted in mice null for PROKR2 presented new information. These mice were shown to have a loss of a majority of GnRH neurons in the forebrain (in adults) and they had also formed complicated webs of olfactory vomeronasal axons. This confirms the taxing neuroendocrine effects that PROK2 and GnRH dysfunction have on mice and humans. To continue, recent studies have also raised the possibility of PROK2 having heparan sulfate binding characteristics, suggesting that this system may cross with anosmin and FGF signaling pathways (Wierman et al.).

Additionally, many studies have confirmed mutations in PROK2 as likely causes of GnRH dysfunction. For example, two brothers with KS and their sister with nIHH were found to have homozygous deletion in the same PROK2 gene. To add, PROK2 knockout mice were studied to display defects in the development of the OB. They also exhibited severely disrupted GnRH migration, leading to a drastic decrease in GnRH population. These studies suggest that homozygous mutations of PROK2 cause both KS and nIHH (Pitteloud, Zhang, et al.). Similar investigations have also been performed in other vertebrates including zebrafish to understand

the effects of PROK2 mutations during early stages of GnRH neuronal migration. These studies indicated that prokr1b (a zebrafish related gene in the prokineticin family) spatial-temporal expression is directly related to GnRH3. Therefore, prokr1b knockout zebrafish hindered development of GnRH3 fibers (Bassi et al.).

To conclude, mutations that alter the ligand-receptor relationship between PROK2 and PROKR2 can disrupt signalling far beyond GnRH migration. With disturbance to the neuronal pathways connecting olfactory regions to the hypothalamus during embryonic development, the entire neuroendocrine signalling in fertility is compromised. Together with KS-like phenotypes presented in PROKR2 knockout mice, it can be strongly argued that PROK2 is a primary protein involved in Kallmann Syndrome (Dodé, Teixeira, et al.).

Link between Genetic Mutations and Female Infertility: Kallmann Syndrome

Although mutations in FGFR1 and PROK2, being discussed within this paper, are two significant factors leading to KS; there are more than 30 other genes that are associated with the syndrome, in which phenotypic correlation and novel mutations have been identified (Chu et al.). For this reason, it is essential to understand that a mutation in one gene is not guaranteed to cause the same phenotypic symptoms in all affected individuals. The importance of understanding the convergence of more than one gene mutation is crucial, as it has been recently proven that KS is not just a monogenic Mendelian disease, rather it is oligogenic in nature (more than one gene is involved). This is proven within a study that investigated 47 Polish males, including 25 with CHH and 22 with KS. After undergoing a customized gene panel screening; 16 patients showed symptoms such as micropenis, 28 cases of delayed/absence of puberty, and 3 cases of infertility were presented (Gach et al.). Within these patients several different variants were shown, determining that KS can be a result of multiple mutations, not simply one. This provides a biological basis as to how KS is contributed to many genetic variants.

After identifying the connection between both genes, the understanding of their neuronal pathways “crossing” is essential. Crossing, in this case, refers to the functional convergence between the neuronal signalling pathways of both genes in which this overlapping causes the influences in the developmental process of the GnRH neurons. Considering phenotypes from studies using PROKR2 or PROK2 deficient mice and the likely effects in human mutations on protein synthesis, these mutations are predicted to be loss-of-function mutations as well. As KS results from insufficient signalling through either FGFR1 or PROKR2, it is expected that both of their signaling pathways interfere at some level. This interference could possibly be the activation of the mitogen-activated protein kinase pathway (Dodé, Teixeira, et al.). This argues that the signalling pathways of FGFR1 and PROK2 converge functionally on the development of the GnRH neuronal system to cause KS by impairing GnRH neuronal development.

To redirect focus back to the bigger picture of HPG axis disruption due to such mutations, the complete biological chain leading to KS should be established. Genetic mutations in FGFR1 cause hindrance to the development of both olfactory and GnRH neurons. As fewer neurons develop before birth, during embryonic development, less neurons reach the hypothalamus. As

proven, GnRH regulates the secretion of reproductive hormones such as FSH and LH in females. Hence, GnRH deficiency presents low basal levels of FSH and LH with the absence of bilateral olfactory bulbs in the brain (Dash and Raj). Continuing, PROK2 primarily affects the migration of neurons, namely GnRH. Mutations in PROK2 and PROKR2 were first identified and associated with KS in 2006 in a study of 192 patients with KS. Mutations were found in 2% and 7% of patients, respectively. Since this finding, roughly nine variants of PROK2 mutations and 24 of PROKR2 have been reported in patients with varying phenotypes and symptoms of KS (Abreu, Kaiser, et al.).

Given the involvement of both FGFR1 and PROK2 in processes relevant to GnRH development and migration, their association with MAPK signaling is a possible functional link between them. Evidence from knockout mice provide insight into how MAPK signalling pathways are productively activated by FGFR1 and PROK2. Therefore, it is plausible that these pathways may overlap in downstream cellular responses. However, due to the fact that available evidence does not directly establish a correlation between FGFR1 and PROK2, further investigation is needed to determine whether the MAPK-mediated signalling influences the shared outcome of disruptions in GnRH causing Kallmann Syndrome.

Together, both genes play different roles but functionally converge to alter the development and migration of olfactory and GnRH neurons, ultimately causing severe disruptions to the HPG axis. Through low GnRH neuron count and an underdeveloped olfactory bulb, the HPG can not function systematically, leading to Kallmann Syndrome (anosmia and CHH).

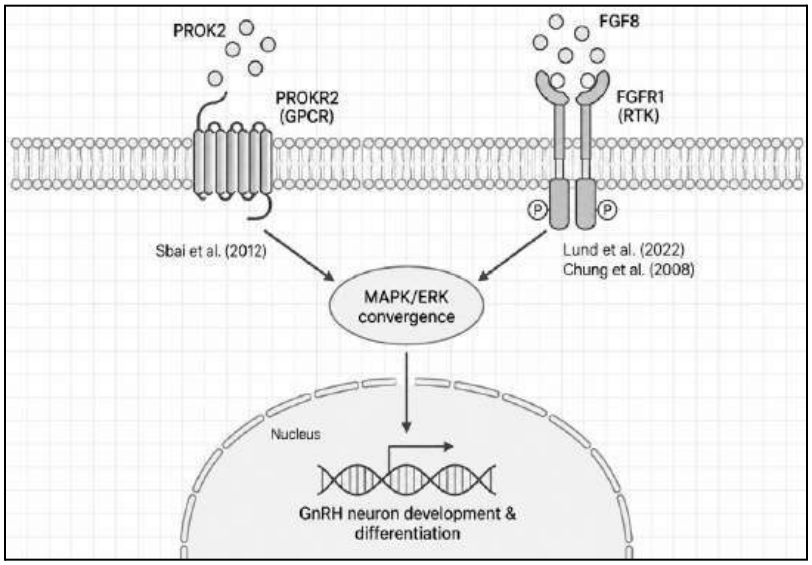

Conclusion

Overall, the evidence synthesized and reviewed in this paper suggests that Kallmann Syndrome is caused by a multitude of gene mutations, more specifically PROK2 and FGFR1. As these genes play their individual roles in disrupting the HPG axis, their convergence becomes

evident through their neuronal pathways and phenotypic symptoms which they cause, including CHH and anosmia. A potential convergence between both PROK2 and FGFR1 provides a possible hypothesis that may help to investigate the mechanisms that can be contributing to reproductive dysfunction, including female infertility. This neuroendocrine dysfunction that leads to insubstantial secretion of female reproductive hormones is the primary cause of Kallmann Syndrome, ultimately leading to infertility in women. As FGFR1 and PROK2 both contribute to different aspects of GnRH and olfactory neuron dysfunction, the disruption of either of their pathways is likely to interfere with neuronal signalling.

Due to the heterogeneous nature of KS, it is appropriate to suggest that variants within these genes, along with others, can cause Kallmann Syndrome in varying forms. Whether it be CHH, nIHH, or KS, these mutated gene variants can have lasting effects on patients. Though KS is oligogenic and doesn't require both PROK2 and FGFR1 mutations to occur, their interconnectedness within GnRH neuronal development increases risk.

Although this review generally focuses on the negative implications that arise with FGFR1 and PROK2/PROKR2 mutations in regard to female reproductive health, a majority of the mechanistic evidence relates to GnRH neuronal development and migration rather than direct female-specific reproductive consequences. For a fact, findings concluded from animal models and human investigations of Kallman Syndrome provide outstanding evidence into the roles that each pathway plays in reproduction function; however, these findings cannot establish that the pathways have distinct effects on females or that their functional convergence causes female infertility. For this reason, further studies involving female patients are likely needed to clarify the said relationship between both pathways.

As studies of Kallmann Syndrome, particularly associated with these two genes, advance, the likelihood of understanding the deeper mechanisms behind these mutations will increase. With evidence from studies consisting of affected female patients increasing, it is possible for a complex insight of these genes to be delivered. However, up to date, there isn't a clear grasp of the junction between FGFR1 and PROK2 in female infertility. To conclude, further research into how combinations of genetic variants affect GnRH neuron development and neuroendocrine signalling could provide break-through evidence to advance the field of neuroendocrine infertility.

Acknowledgements

Thank you to Dr. Jorge Avila and Ms. Bre Calhoun for their guidance in helping me reflect my ideas into a structured paper. Their expertise and feedback ensured that I produce my best work.

The platform **Zotero** was used to organize and cite my bibliography throughout this paper.

Works Cited

Abreu, Ana Paula, Sekoni D. Noel, et al. “Evidence of the Importance of the First Intracellular Loop of Prokineticin Receptor 2 in Receptor Function.” Molecular Endocrinology, vol. 26, no. 8, Aug. 2012, pp. 1417–27. Silverchair, https://doi.org/10.1210/me.2012-1102.

Abreu, Ana Paula, Ursula B. Kaiser, et al. “The Role of Prokineticins in the Pathogenesis of Hypogonadotropic Hypogonadism.” Neuroendocrinology, vol. 91, no. 4, June 2010, pp. 283–90. PubMed Central, https://doi.org/10.1159/000308880.

Bassi, Ivan, et al. “Knocking-down of the Prokineticin Receptor 2 Affect Reveals Its Complex Role in the Regulation of the Hypothalamus-Pituitary-Gonadal Axis in the Zebrafish Model.” Scientific Reports, vol. 10, May 2020, p. 7632. PubMed Central, https://doi.org/10.1038/s41598-020-64077-2.

Bianco, Suzy D. C., and Ursula B. Kaiser. “The Genetic and Molecular Basis of Idiopathic Hypogonadotropic Hypogonadism.” Nature Reviews. Endocrinology, vol. 5, no. 10, Oct. 2009, pp. 569–76. PubMed Central, https://doi.org/10.1038/nrendo.2009.177.

Blakemore, Laura J., and Paul Q. Trombley. “Zinc as a Neuromodulator in the Central Nervous System with a Focus on the Olfactory Bulb.” Frontiers in Cellular Neuroscience, vol. 11, Sept. 2017. Frontiers, https://doi.org/10.3389/fncel.2017.00297.

Bliss, Stuart P., et al. “GnRH Signaling, the Gonadotrope and Endocrine Control of Fertility.” Frontiers in Neuroendocrinology, vol. 31, no. 3, July 2010, pp. 322–40. PubMed Central, https://doi.org/10.1016/j.yfrne.2010.04.002.

Cho, Hyun-Ju, et al. “Nasal Placode Development, GnRH Neuronal Migration and Kallmann Syndrome.” Frontiers in Cell and Developmental Biology, vol. 7, July 2019, p. 121. PubMed Central, https://doi.org/10.3389/fcell.2019.00121.

Chu, Guoming, et al. “Mutation Spectrum of Kallmann Syndrome: Identification of Five Novel Mutations across ANOS1 and FGFR1.” Reproductive Biology and Endocrinology, vol. 21, no. 1, Mar. 2023, p. 23. Springer Link, https://doi.org/10.1186/s12958-023-01074-w.

Chung, Wilson C. J., et al. “Fibroblast Growth Factor 8 Signaling through Fibroblast Growth Factor Receptor 1 Is Required for the Emergence of Gonadotropin-Releasing Hormone Neurons.” Endocrinology, vol. 149, no. 10, Oct. 2008, pp. 4997–5003. PubMed Central, https://doi.org/10.1210/en.2007-1634.

Chung, Wilson CJ, and Pei-San Tsai. “The Initiation and Maintenance of Gonadotropin-Releasing Hormone Neuron Identity in Congenital Hypogonadotropic Hypogonadism.” Frontiers in Endocrinology, vol. 14, Apr. 2023, p. 1166132. PubMed Central, https://doi.org/10.3389/fendo.2023.1166132.

Collado-Sole, Alejandro, et al. “Microglia Rank Signaling Regulates GnRH Neuronal Function and the Hypothalamic-Pituitary-Gonadal Axis.” Science (New York, N.Y.), vol. 392, no. 6800, May 2026, p. eaeb6999. PubMed Central, https://doi.org/10.1126/science.aeb6999.

Concepción-Zavaleta, Marcio José, et al. “Endocrine Factors Associated with Infertility in Women: An Updated Review.” Expert Review of Endocrinology & Metabolism, vol. 18, no. 5, 2023, pp. 399–417. PubMed, https://doi.org/10.1080/17446651.2023.2256405.

Dash, Prafulla Kumar, and Dinesh Harvey Raj. “Biochemical and MRI Findings of Kallmann’s Syndrome.” BMJ Case Reports, vol. 2014, Dec. 2014, p. bcr2014207386. PubMed Central, https://doi.org/10.1136/bcr-2014-207386.

Dodé, Catherine, Luis Teixeira, et al. “Kallmann Syndrome: Mutations in the Genes Encoding Prokineticin-2 and Prokineticin Receptor-2.” PLoS Genetics, vol. 2, no. 10, Oct. 2006, p. e175. PubMed Central, https://doi.org/10.1371/journal.pgen.0020175.

Dodé, Catherine, Jacqueline Levilliers, et al. “Loss-of-Function Mutations in FGFR1 Cause Autosomal Dominant Kallmann Syndrome.” Nature Genetics, vol. 33, no. 4, Apr. 2003, pp. 463–65. PubMed, https://doi.org/10.1038/ng1122.

Dodé, Catherine, and Philippe Rondard. “PROK2/PROKR2 Signaling and Kallmann Syndrome.” Frontiers in Endocrinology, vol. 4, Apr. 2013, p. 19. PubMed Central, https://doi.org/10.3389/fendo.2013.00019.

Fadiga, Lúcia, et al. “A Novel FGFR1 Missense Mutation in a Portuguese Family with Congenital Hypogonadotropic Hypogonadism.” International Journal of Molecular Sciences, vol. 23, no. 8, Apr. 2022, p. 4423. PubMed Central, https://doi.org/10.3390/ijms23084423.

Gach, Agnieszka, et al. “New Findings in Oligogenic Inheritance of Congenital Hypogonadotropic Hypogonadism.” Archives of Medical Science : AMS, vol. 18, no. 2, Sept. 2020, pp. 353–64. PubMed Central, https://doi.org/10.5114/aoms.2020.98909.

Gomez, Jonathan, and Shannon Pickup. “Cribriform Plate Fractures.” StatPearls, StatPearls Publishing, 2026. PubMed, http://www.ncbi.nlm.nih.gov/books/NBK562192/.

Guo, Yan-Jun, et al. “ERK/MAPK Signalling Pathway and Tumorigenesis.” Experimental and Therapeutic Medicine, vol. 19, no. 3, Mar. 2020, pp. 1997–2007. PubMed, https://doi.org/10.3892/etm.2020.8454.

Hao, Yilong, et al. “Roles of FGF8 Subfamily in Embryogenesis and Oral-maxillofacial Diseases (Review).” International Journal of Oncology, vol. 54, no. 3, Mar. 2019, pp. 797–806. www.spandidos-publications.com, https://doi.org/10.3892/ijo.2019.4677.

Holesh, Julie E., et al. “Physiology, Ovulation.” StatPearls, StatPearls Publishing, 2026. PubMed, http://www.ncbi.nlm.nih.gov/books/NBK441996/.

Kaprara, Athina, and Ilpo T. Huhtaniemi. “The Hypothalamus-Pituitary-Gonad Axis: Tales of Mice and Men.” Metabolism: Clinical and Experimental, vol. 86, Sept. 2018, pp. 3–17. PubMed, https://doi.org/10.1016/j.metabol.2017.11.018.

Katoh, Masuko, et al. “FGFR-Targeted Therapeutics: Clinical Activity, Mechanisms of Resistance and New Directions.” Nature Reviews Clinical Oncology, vol. 21, no. 4, Apr. 2024, pp. 312–29. www.nature.com, https://doi.org/10.1038/s41571-024-00869-z.

Kim, Soo-Hyun. “Congenital Hypogonadotropic Hypogonadism and Kallmann Syndrome: Past, Present, and Future.” Endocrinology and Metabolism, vol. 30, no. 4, Dec. 2015, pp. 456–66. PubMed Central, https://doi.org/10.3803/EnM.2015.30.4.456.

Kumar Yadav, Rajiv, et al. “Kallmann Syndrome: Diagnostics and Management.” Clinica Chimica Acta, vol. 565, Jan. 2025, p. 119994. ScienceDirect, https://doi.org/10.1016/j.cca.2024.119994.

Martin, Cecilia, et al. “The Role of the Prokineticin 2 Pathway in Human Reproduction: Evidence from the Study of Human and Murine Gene Mutations.” Endocrine Reviews, vol. 32, no. 2, Apr. 2011, pp. 225–46. PubMed Central, https://doi.org/10.1210/er.2010-0007.

Martinez-Mayer, Julian, and Maria Ines Perez-Millan. “Phenotypic and Genotypic Landscape of PROKR2 in Neuroendocrine Disorders.” Frontiers in Endocrinology, vol. 14, Feb. 2023. Frontiers, https://doi.org/10.3389/fendo.2023.1132787.

Miraoui, Hichem, et al. “Role of Fibroblast Growth Factor (FGF) Signaling in the Neuroendocrine Control of Human Reproduction.” Molecular and Cellular Endocrinology, vol. 346, nos. 1–2, Oct. 2011, pp. 37–43. PubMed Central, https://doi.org/10.1016/j.mce.2011.05.042.

Pitteloud, Nelly, Chengkang Zhang, et al. “Loss-of-Function Mutation in the Prokineticin 2 Gene Causes Kallmann Syndrome and Normosmic Idiopathic Hypogonadotropic Hypogonadism.” Proceedings of the National Academy of Sciences of the United States of America, vol. 104, no. 44, Oct. 2007, pp. 17447–52. PubMed Central, https://doi.org/10.1073/pnas.0707173104.

Pitteloud, Nelly, James S. Acierno, et al. “Mutations in Fibroblast Growth Factor Receptor 1 Cause Both Kallmann Syndrome and Normosmic Idiopathic Hypogonadotropic Hypogonadism.” Proceedings of the National Academy of Sciences of the United States of America, vol. 103, no. 16, Apr. 2006, pp. 6281–86. PubMed Central, https://doi.org/10.1073/pnas.0600962103.

PROKR2 Prokineticin Receptor 2 [Homo Sapiens (Human)] - Gene - NCBI. https://www.ncbi.nlm.nih.gov/gene?cmd=Retrieve&list_uids=128674&utm_source=chatgpt.com. Accessed 17 Oct. 2026.

Sang, Qing, et al. “Understanding the Genetics of Human Infertility.” Science [New York, N.Y.], vol. 380, no. 6641, Apr. 2023, pp. 158–63. PubMed, https://doi.org/10.1126/science.adf7760.

Soejima, Yoshiaki, et al. “Involvement of a Novel Variant of FGFR1 Detected in an Adult Patient with Kallmann Syndrome in Regulation of Gonadal Steroidogenesis.” International Journal of Molecular Sciences, vol. 26, no. 6, Jan. 2025, p. 2713. www.mdpi.com, https://doi.org/10.3390/ijms26062713.

Stamatiades, George A., and Ursula B. Kaiser. “Gonadotropin Regulation by Pulsatile GnRH: Signaling and Gene Expression.” Molecular and Cellular Endocrinology, vol. 463, Mar. 2018, pp. 131–41. PubMed, https://doi.org/10.1016/j.mce.2017.10.015.

Wierman, Margaret E., et al. “Gonadotropin Releasing Hormone (GnRH) Neuron Migration: Initiation, Maintenance and Cessation as Critical Steps to Ensure Normal Reproductive

Function.” Frontiers in Neuroendocrinology, vol. 32, no. 1, Jan. 2011, pp. 43–52. PubMed Central, https://doi.org/10.1016/j.yfrne.2010.07.005.

Xie, Qinying, et al. “The Role of Kisspeptin in the Control of the Hypothalamic-Pituitary-Gonadal Axis and Reproduction.” Frontiers in Endocrinology, vol. 13, June 2022, p. 925206. PubMed Central, https://doi.org/10.3389/fendo.2022.925206.

Yellapragada, Venkatram, et al. “FGF8–FGFR1 Signaling Regulates Human GnRH Neuron Differentiation in a Time- and Dose-Dependent Manner.” Disease Models & Mechanisms, vol. 15, no. 8, Aug. 2022, p. dmm049436. PubMed Central, https://doi.org/10.1242/dmm.049436.

Young, Jacques, et al. “Clinical Management of Congenital Hypogonadotropic Hypogonadism.” Endocrine Reviews, vol. 40, no. 2, Apr. 2019, pp. 669–710. PubMed, https://doi.org/10.1210/er.2018-00116.

Кокорева, К. Д., et al. “Врожденный Изолированный Гипогонадотропный Гипогонадизм: Клинический и Молекулярно-Генетический Полиморфизм.” Problems of Endocrinology, vol. 67, no. 4, Aug. 2021, pp. 46–56. PubMed Central, https://doi.org/10.14341/probl12787.

The Effect of Wing Aspect Ratio on Aerodynamic Efficiency of Regional Jets, Narrow-Body, and Wide-Body Commercial Aircraft By Guanchen Wang

Abstract

Commercial aircraft keep seeking the best balance between fuel efficiency through various aerodynamic designs of their wings. This paper investigates how the wing aspect ratio can affect a commercial aircraft's aerodynamic efficiency. To see how the wing aspect ratio will affect a commercial aircraft's aerodynamic efficiency, 9 different wings with the same surface areas were tested in XFLR5. The aerodynamic efficiency, lift-to-drag ratio, and induced drag were analyzed to answer whether increasing wing aspect ratio affects aerodynamic efficiency in commercial aircraft. The results show that increasing the wing aspect ratio helps with aerodynamic efficiency. However, the simulation did not include structural analysis; this investigation should be considered a suggestion rather than a prediction.

Introduction

Commercial aircraft are designed to carry passengers while meeting aerodynamic and structural requirements. One important design process is to reduce drag. With differences in wing geometry, the wings need to be changed, and their important role is in determining aerodynamic performance.

Commercial aviation relies on aerodynamic efficiency to lower fuel consumption while maintaining a relatively safe and stable flight. One of the biggest factors determining aerodynamic performance is wing design. Within the performance of the wing design is the aspect ratio, defined as the wingspan squared divided by the total wing area (Anderson, Fundamentals of Aerodynamics 451). This specific aspect ratio design is one of the key factors in determining the aircraft efficiency. Even though increasing wing aspect ratio can improve aerodynamic efficiency, having too much wing aspect ratio can be detrimental (Uzun 25; Bradley and Droney). To determine the optimal aerodynamic effects of wing aspect ratio using XFLR5 simulation, this research conducts experiments to assess how aspect ratio affects aerodynamic efficiency across three categories of aircraft. The experiment uses one aircraft from each aircraft category: Embraer E175 from regional jets, Boeing 757 from narrow-body jets, and Airbus A350 from wide-body jets. The three aircraft chosen for this paper are selected because each type is used in different scenarios. For example, the Embraer E175 is primarily used by airlines to serve regional routes. At the same time, the Boeing 757 is chosen as a narrow-body aircraft that typically carries passengers on cross-country flights and even some transatlantic flights. Finally, the Airbus A350 is chosen for the wide-body because it is used for long-haul flights, including flights over 22 hours.

The research question that guides this paper is: How do differences in wing aspect ratio influence aerodynamic efficiency among regional jets, narrow-body airliners, and wide-body airliners operating under comparable subsonic flight conditions? By conducting comparisons in this research, I hope to identify the relationship between aspect ratio and aerodynamic efficiency

in modern aviation.

Literature Review

In this paper, drag plays an important role in the experiment. Drag is treated in this paper as the force acting on the wing during the experiment. In this case, the drag on the wing at subsonic speeds can be divided into two parts: parasite drag, due to skin friction, and induced drag, caused by the lift on the wing's limited surface. When the aircraft is in the air, the wing has higher pressure beneath it than above it. This creates trailing vortices, which cause downward velocity in the direction of the drag (NASA Glenn Research Center; Anderson, Fundamentals of Aerodynamics 436).

In this paper, the primary factor is the wing aspect ratio (AR). Aspect ratio is inversely proportional to induced drag, the force that affects an aircraft's aerodynamic efficiency (NASA Glenn Research Center; Anderson, Fundamentals of Aerodynamics). Induced drag is created by the lift force when in motion. Induced drag is a function of airspeed v, density p, reference area S, drag D_i, and drag coefficient C_{Di}.

$$D_i = \frac{1}{2}\rho V^2 S C_{Di} \quad \text{(Eqn. 1)}$$

$$AR = \frac{b^2}{S} \quad \text{(Eqn. 2) (Uzun 20)}$$

$$K = \frac{1}{\pi \bullet e \bullet AR} \quad \text{(Eqn. 3)}$$

In Eqn. 1, D_i is the induced drag force, and C_{Di} is the induced drag coefficient. C_{Di} depends on the lift coefficient through $C_{Di} = K*C_L{}^2$, where (Eqn. 3) is the induced drag factor, e is the span efficiency factor, and AR is the aspect ratio.

Equation 3 connects Equations 1 and 2 by showing an inverse proportion. Meaning that the higher the aspect ratio goes,s the lower the induced drag is (NASA Glenn Research Center; Anderson, Fundamentals of Aerodynamics 444). In addition, according to Uzun, a higher wing aspect ratio increases the lift-to-drag ratio, thereby improving flight endurance and overall mission capabilities. However, there was a downside to a high aspect ratio: the structural challenges posed by the extra load on the wings (Uzun 25).

This connection to a higher AR (Aspect Ratio), along with the fact that AR is in the denominator of Equation 2, helps explain how a greater AR results in less induced drag than in another wing with similar flight characteristics. Equation 2 predicts that a long, thin wing would produce the same lift but have less induced drag than a shorter, wider wing at the same speed.

In addition, the lift-to-drag ratio is the perfect measure of aerodynamic efficiency because, in flight, the ratio of an aircraft's weight to the required thrust equals L/D. Furthermore, if the aircraft experiences a problem and has no thrust, L/D would indicate how far it can glide before reaching a flight level of 0 (Anderson, Introduction to Flight 487). The equation for lift over drag can be derived as

$$\left(\frac{L}{D}\right)_{max} = \frac{1}{2}\sqrt{\left(\pi e \bullet \frac{AR}{C_{D0}}\right)} \quad \text{(Eqn. 4)}$$

This equation is reached at the angle of attack where induced drag equals profile drag (Raymer 296).

Using existing research and experiments, this paper establishes how the wing aspect ratio, aerodynamic efficiency, induced drag, and structure are considered in this experiment.

Furthermore, the experiment in this paper can be compared with that conducted by Kretov and Tiniakov on the Boeing 777 (Kretov and Tiniakov 4). This paper's experiment tested 3 different wingspans across 3 aircraft: the Embraer E175, Boeing 757-200, and Airbus A350-900. Despite having three different aircraft of varying sizes and purposes, the Airbus A350-900 can be closely compared with the Boeing 777 due to their similar wingspans.

Methodology

To find the most aerodynamically efficient wings for each of the three jets, this experiment analyses nine different wing designs across three aircraft, tested in a simulation using XFLR 5, software for evaluating aircraft performance. The three aircraft chosen for this experiment are from each category of commercial jets: Embraer E175, regional jet; Boeing 757-200, narrow-body; and Airbus A350-900, widebody. The first step in the experiment was to design nine wings with different aspect ratios and roughly the same wing area for each of the three aircraft I chose. Furthermore, I used the same NACA 2412 airfoil for all 9 different wing designs.

Next, in XFLR5, each wing of all three aircraft was analyzed at different attack angles and compared in terms of lift and drag. After collecting all the data from the experiment, it is sufficient to determine how different aspect ratios affect the aerodynamic efficiency of the Boeing 757, Airbus A350, and Embraer E175. Furthermore, in this experiment, most variables are controlled, so even though the wings had different aspect ratios, the controlled environment allows for comparison and analysis of the different wing types across the three aircraft. In addition, this allows the research to change only the Aspect ratio so that the results could be compared with minimal change. The three-dimensional analysis used the VLM1 method in XFlR5 with the viscous option enabled, and the profile drag was taken from XFoil polars of the NACA 2412 airfoil (Ncrit = 9). In addition, a fixed-speed polar was run at a freestream velocity of 70 m/s using standard sea-level air properties, with the angle of attack ranging from -4° to 12° in 0.5° increments. Furthermore, the Reynolds numbers are based on the mean chord for each aircraft: $1.2\text{-}1.4\times10^7$ for the E175 wings, $1.9\text{-}2.1\times10^7$ for the Boeing 757-200, and $3.1\text{-}3.4\times10^7$ for the Airbus A350-900. The output for each run consisted of the lift coefficient C_L; the induced-drag coefficient $C_{D,i}$; the viscous-drag coefficient $C_{D,v}$; the total drag coefficient $C_D = C_{D,i} + C_{D,v}$; and the pitching-moment coefficient, all referenced to the wing area.

Result

Aircraft	Aspect Ratio	Wing Span (m)	Wing Area (m²)	Air foil	Run
E175	9	25.972	74.953	NACA 2412	1
E175	10	27.378	74.955	NACA 2412	2
E175	11	28.714	74.955	NACA 2412	3
B757-200	9	40.89	185.777	NACA 2412	4
B757-200	10	43.102	185.769	NACA 2412	5
B757-200	11	45.206	185.775	NACA 2412	6
A350-900	9	64.16	457.445	NACA 2412	7
A350-900	10	67.639	457.443	NACA 2412	8
A350-900	11	70.936	457.433	NACA 2412	9

Table 1. Description of nine wing configurations in this experiment.

Table 1 describes the nine wing configurations tested during the experiment. All three aircraft types' aspect ratios range from 9 to 11, using the same airfoil.

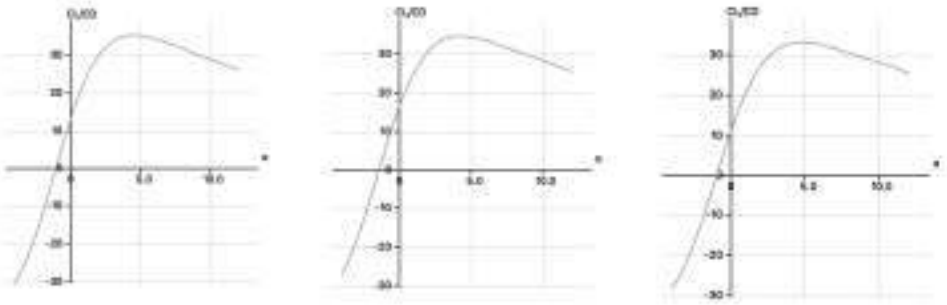

Figure 1. Aspect ratio = 11. (a) Analysis of the Airbus A350-900: lift, drag, and angle of attack using XFLR. (b) Analysis of the Boeing 757-200: lift, drag, and angle of attack using XFLR.(c) Analysis of the Embraer E175: lift, drag, and angle of attack using XFLR.

For the group with the highest aspect ratio, the Airbus A350, Boeing 757, and Embraer 175, the Airbus A350-900 produced the highest lift-to-drag ratio of 35.24. In contrast, the Embraer 175 produced the lowest lift-to-drag ratio of 33.19, with the Boeing 757-200's lift-to-drag ratio between those two aircraft at 34.63.

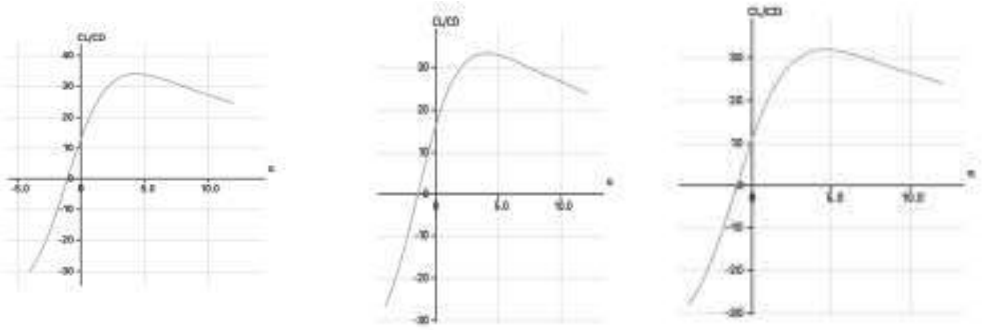

Figure 2. Aspect ratio = 10. (a) Analysis of the Airbus A350-900: lift, drag, and angle of attack using XFLR. (b) Analysis of the Boeing 757-200: lift, drag, and angle of attack using XFLR. (c) Analysis of the Embraer E175: lift, drag, and angle of attack using XFLR.

For the group with the highest aspect ratio, the Airbus A350-900 produced the highest lift-to-drag ratio of 33.87. In contrast, the Embraer 175 produced the lowest lift-to-drag ratio, 31.91, while the Boeing 757-200's lift-to-drag ratio fell between those two aircraft at 33.3.

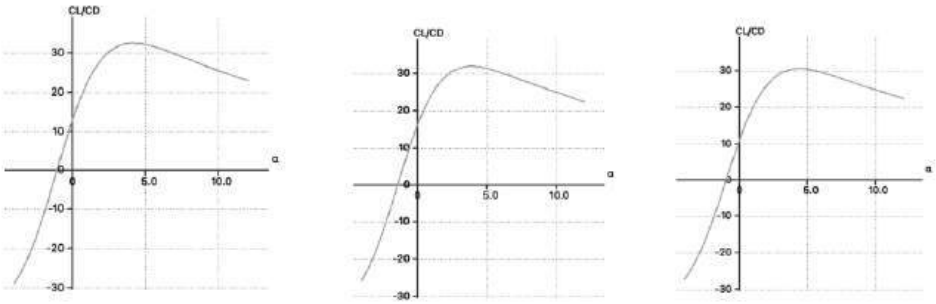

Figure 3. Aspect ratio = 9. (a) Analysis for Airbus A350-900; variables of lift, drag, and angle of attack using XFLR. (b) Analysis of the Boeing 757-200: lift, drag, mass, and angle of attack using XFLR. (c) Analysis of the Embraer E175: lift, drag, and angle of attack using XFLR.

For the group with the highest aspect ratio, the Airbus A350-900 produced the highest lift-to-drag ratio of 32.42. In contrast, the Embraer 175 produced the lowest lift-to-drag ratio of 30.52, with the Boeing 757-200's lift-to-drag ratio falling between those two aircraft, at 31.83.

Aircraft	AR	Maximum L/D	AoA at Max (°)	CL	CDi	CDV	CD	e	Gain vs. AR = 9
A350-900	9	32.42	4.0	0.405	0.00626	0.00624	0.01249	0.926	0
A350-900	10	33.87	4.5	0.453	0.00704	0.00633	0.01336	0.926	4.47%
A350-900	11	35.24	4.5	0.459	0.00666	0.00638	0.01303	0.917	8.70%
757-200	9	31.83	4	0.438	0.00706	0.00670	0.01377	0.962	0
757-200	10	33.30	4	0.446	0.00662	0.00677	0.01339	0.956	4.61%
757-200	11	34.63	4.5	0.494	0.00738	0.00688	0.01425	0.956	8.79%
E-175	9	30.52	4.5	0.456	0.00770	0.00723	0.01493	0.954	0
E-175	10	31.91	4.5	0.465	0.00726	0.00731	0.01457	0.948	4.55%
E-175	11	33.19	5	0.515	0.00810	0.00743	0.01553	0.950	8.75%

Table 2. Maximum lift-to-drag ratio, the AoA at which it occurs, and the lift and drag coefficients at that point. e is the span-efficiency factor from equation 3. The last column shows the percent increase in lift-to-drag ratio compared to the lowest-aspect-ratio wing of the same aircraft.

Discussion

The results from the experiments show that wings with increasing aspect ratios consistently increase each aircraft's aerodynamic efficiency. Across the Airbus A350, Boeing 757, and Embraer E175 models, different aspect ratios affect the lift-to-drag ratio. That said, the hypothesis that increasing the aspect ratio while keeping the wing area constant will raise the maximum lift-to-drag ratio is correct.

The results for the Airbus A350 showed an increase in the lift-to-drag ratio when the aspect ratio increased. The aircraft's highest lift-to-drag ratio was 35.24, compared to a lowest of 32.42. This was an 8.7% increase from the lowest lift-to-drag ratio to the highest lift-to-drag Ratio.

The results for the Boeing 757 showed an increase in the lift-to-drag ratio when the aspect ratio increased. The aircraft's highest lift-to-drag ratio was 34.63, compared to a lowest of 31.83. This was an 8.79% increase from the lowest lift-to-drag ratio to the highest lift-to-drag Ratio.

The results for the Embraer E175 showed an increase in the lift-to-drag ratio when the aspect ratio increased. The aircraft's highest lift-to-drag ratio was 33.19, compared to the lowest of 30.52. This was an 8.75% increase from the lowest lift-to-drag ratio to the highest lift-to-drag Ratio.

At every aspect ratio, the three aircraft rank in the same order: A350-900 highest, 757-200 second, and E175 lowest, with the A350 ahead of the E175 by 1.9-2.1 points of L/D (6–7%). Because the three wings in each group now share the same aspect ratio, airfoil, and freestream speed, this ranking cannot be explained by Equation 3, which depends only on aspect ratio and span efficiency. In fact, the induced-drag coefficients at (L/D)max are nearly the same for all three aircraft at a given aspect ratio (Table 2), and the A350 has the lowest span-efficiency factor of the three (0.917–0.926 versus 0.948–0.962), which by Equation 3 alone would make it the least efficient. The viscous-drag column instead sets the ranking: at AR = 9, C_D,v is 0.00624 for the A350, 0.00670 for the 757, and 0.00723 for the E175, a spread of 16%; the same order holds at AR = 10 and AR = 11. Wing size, therefore, enters the comparison not through Equation 3 but through C_D,0 in Equation 4, and the Table 2 viscous-drag column is direct evidence for this mechanism.

Kretov and Tiniakov examined the Boeing 777-300ER wing, with an aspect ratio of 9.6, and compared it with the 777-9 wing, which has an aspect ratio of 11. They found that raising the aspect ratio by about 1.5 cuts induced drag by roughly 22.5% and total wing drag by about 13% (Kretov and Tiniakov 10). At first glance, this seems to contradict Equation 3. Initially, this suggests that increasing the aspect ratio from 9.6 to 11 should reduce the induced-drag coefficient by about 12.7% due to 9.6/11 = 0.873. The difference is that the 777-9 wing is not a constant-area stretch: its area increased from about 437 m2 to 467 m2 along with the span. Since the airplane weighs the same, the larger wing needs a smaller lift coefficient in cruise, and induced drag scales with the square of the lift coefficient. So two effects are working at once, and they multiply: (9.6/11) × (437/467)2 ≈ 0.76, which is a drop of about 24%. That is very close to the 22.5% reported by Kretov and Tiniakov. In this paper, the wing area remains the same for each aircraft, and each wing is compared at its own best L/D point, so only the 1/AR term in Equation 3 matters. Equation 4 then says that (L/D)max should grow with the square root of the aspect ratio, which, for a change from AR = 9 to AR = 11, represents a gain of 10.6%. The gains in this study were 8.7 -- 8.8%, or about 82% of that, and Table 2 shows where the rest went. The viscous-drag coefficient increases by 2.3–2.8% from AR = 9 to AR = 11 for every aircraft because the longer wings have narrower chords and therefore operate at lower Reynolds numbers, while the span-efficiency factor drops by about 1%. Neither effect is large, but together they reduce the benefit. Once the change in wing area is accounted for, the two studies agree. The fact that the gains were almost the same for all three aircraft, 8.70%, 8.79%, and 8.75%, even though their wing areas differ by a factor of six, shows that the aspect-ratio effect depends on the planform and not on how big the wing is. Uzun observed the same pattern in a small fixed-wing UAV, where a higher aspect ratio increased L/D but also increased the structural load at the wing root (Uzun 20). A UAV wing is much smaller and flies at a much lower Reynolds number than the wings in this study, so the effect clearly holds across a wide range of wing sizes. Finally, the maximum L/D values of 31–36 found here are higher than the typical value of about 17 for an airliner in cruise (Raymer 26). The reasons for this are discussed below. However, the limitations of this research include that the study examined only wings,

without engines, fuselage, winglets, or horizontal stabilizers. Furthermore, in XFLR5, everything is tested with low-speed flow, meaning it cannot capture the compressibility and wave-drag effects that occur at normal airline cruising speeds of Mach 0.78-0.85 (Deperrois; Raymer).

Conclusion

This research paper investigated how wing aspect ratio affects the aerodynamic efficiency of three commercial aircraft. Each aircraft, Boeing 757, Airbus A350, and Embraer E175, was modeled with different wing aspect ratios and wingspans while maintaining the same wing area. Furthermore, using XFLR5, this study examined how different wing shapes affect aerodynamic performance.

The findings support the conclusion that wing aspect ratio is directly related to aerodynamic efficiency, meaning that a higher aspect ratio yields greater aerodynamic efficiency. Furthermore, these findings support the industry's continued movement toward higher-aspect-ratio wings and explain why it is not doing so immediately (Bradley and Droney). The findings in this paper show that the wing aspect ratio is important for commercial aviation, helping airlines boost their aircraft efficiency. Furthermore, future experiments should include designing different wings from cardboard or metal and testing them in a wind tunnel. This would address the limitations of the experiments in this paper and allow us to fully understand the extent to which changes in the wing aspect ratio affect aerodynamic efficiency.

Works Cited

Anderson, John D. Fundamentals of Aerodynamics. 6th ed., McGraw-Hill Education, 2017.

Anderson, John D. Introduction to Flight. 8th ed., McGraw-Hill Education, 2016.

Bradley, Marty K., and Christopher K. Droney. Subsonic Ultra Green Aircraft Research: Phase I Final Report. NASA/CR-2011-216847, National Aeronautics and Space Administration, 2011,ntrs.nasa.gov/citations/20110011321.

Deperrois, André. Guidelines for QFLR5 v0.03. Oct. 2009, www.xflr5.tech/docs/Guidelines_QFLR5_v0.03.pdf.

Kretov, Anatolii, and Dmytro Tiniakov. "Evaluation of the Mass and Aerodynamic Efficiency of a High Aspect Ratio Wing for Prospective Passenger Aircraft." Aerospace, vol. 9, no. 9, 2022, p.497, https://doi.org/10.3390/aerospace9090497.

NASA Glenn Research Center. "Induced Drag Coefficient." Beginner's Guide to Aeronautics, National Aeronautics and Space Administration, www.grc.nasa.gov/www/k-12/VirtualAero/BottleRocket/airplane/induced.html.

Raymer, Daniel P. Aircraft Design: A Conceptual Approach. 6th ed., American Institute of Aeronautics and Astronautics, 2018.

Uzun, Metin. "Impact of Aspect Ratio on Structural Integrity and Aerodynamic Performance in Fixed-Wing UAV." Journal of Aviation, vol. 9, no. 1, 2025, pp. 19–27, https://doi.org/10.30518/jav.1583881.

The Ethics of Artificial Intelligence in Cybersecurity: Toward a Framework for Responsible Integration By Tahir Ibayev, Mentor By Emily Sheetz

Abstract

This article examines a set of ethical issues that arise when integrating artificial intelligence (AI) technologies into the field of information security. The author analyzes the contradictions between the need to improve the effectiveness of cyber threat detection and the requirements of transparency, accountability and respect for users' privacy rights. Special attention is paid to the problem of the "black box" in the context of security information and event management systems (SIEM) and autonomous security systems. The paper compares existing approaches to ensuring the ethics of AI, including the concept of explicable AI (XAI) and regulatory regulation. Based on the literature analysis, a model for integrating ethical standards into the lifecycle of security system development is proposed.

Keywords artificial intelligence, cybersecurity, AI ethics, algorithmic transparency, explanation of AI, information security, algorithmic bias, data privacy.

1.Introduction

Artificial intelligence has become a key tool in the field of cybersecurity, facilitating threat detection, real-time monitoring, and automated response mechanisms. The adoption of AI-based security measures has shown steady growth driven by the increasing complexity and scale of cyber threats. AI's ability to analyze data, recognize patterns, and predict potential security breaches has made it the cornerstone of modern cybersecurity. However, the implementation of AI in this field is not without ethical challenges. We may encounter bias and privacy violations when using machine learning and deep learning algorithms.

It is precisely these issues of bias in artificial intelligence systems that contribute to a serious ethical problem that arises when algorithms reflect and reinforce existing biases embedded in the training data. Thus, biased AI in the field of cybersecurity can lead to unfair risk assessments, thereby disregarding certain groups. Another significant issue is privacy, because AI-based security measures often require extensive access to personal and organizational data, leading to concerns about data protection and misuse.

Additionally, the opacity of AI in cybersecurity raises concerns because the complexity of the models makes it difficult to understand the reasons behind their decisions and to assess their reliability.

1.1 Problem Statement

The development of artificial intelligence (AI) and machine learning (ML) technologies is significantly changing the field of cybersecurity. Companies are actively using AI to automate intrusion detection, malware analysis, and vulnerability management processes. However, providing autonomous systems with control over critical infrastructures raises serious ethical

concerns. The difficulty lies not only in the technical reliability of algorithms, but also in their ability to make decisions that comply with moral and legal norms.

The relevance of this study is related to the fact that in cybersecurity, AI errors or its bias can have serious consequences. This may result in a violation of citizens' right to privacy, restriction of access to resources, or unfounded accusations by users. The research question is to find a compromise between the high operational efficiency of AI-based security systems and the need to ensure their transparency and accountability.

This study examines the ethical aspects of using AI in monitoring, incident management, and data protection systems. The paper analyzes modern approaches to AI regulation, both Russian and international, and examines ethics as the foundation of trust between technology, government, and users.

2. Causes and the Current State of the Issue

One of the most important moments of the ethical crisis in the field of AI is the phenomenon of the "black box". Current neural network models are capable of detecting anomalies with high accuracy, however, even for developers, the logic of their decision-making becomes a difficult task . In the context of cybersecurity, this creates a situation where the system can block a critical business process or user access, thereby leaving no justification for such an action. As a result, the legitimacy of the use of fully autonomous systems in the legal field becomes highly questionable.

The second important reason is the problem of algorithmic bias. Thus, the AI's work directly depends on the data on which it was trained. We want to give an example on a training sample. For example, if a training sample contains hidden patterns that reflect human biases or mistakes of the past, AI can cause the emergence of certain user groups or traffic types . In the field of information security, this is manifested in false positives, which can be selective, undermining the principle of fairness.

At the moment, the problem is characterized by a lag between the regulatory framework and technological progress. Researchers note that the lack of clear mechanisms for allocating responsibility for mistakes made by artificial intelligence (AI) systems in the field of cybersecurity complicates law enforcement [4]. In cases where an autonomous data protection system leaks information or, conversely, unreasonably restricts access, it remains unclear who should be responsible. This question concerns the algorithm developer, the system owner, or the operator who relied on the AI's conclusions.

The security of the data used for AI training is crucial. The use of confidential information to train cyber defense models carries risks. There is a possibility of re-identification of data subjects or intellectual property leakage as a result of model inversion attacks . Thus, ethical privacy issues conflict with the need for large amounts of data to improve the accuracy of threat detection.

2.1 The problem of the "black box" and lack of transparency in AI systems

The concept of a "black box" has existed since the middle of the twentieth century, but it was in the 2010s that it became particularly relevant. This is due to three key factors. First, there has been an explosive growth in the amount of data used to train models. Secondly, computing power has increased significantly, in particular, due to the development of GPU and TPU. Thirdly, artificial intelligence has begun to be actively implemented in critical areas where its decisions directly affect people's lives and safety.

Thus, the problem of the "black box" has been transformed from the field of academic reasoning into an urgent issue affecting technological ethics, legal norms and philosophical aspects.

The misunderstanding of the work of modern artificial intelligence systems is due to several factors.Firstly, because of their architecture, it is very complex. In deep learning models such as transformers described in "Attention Is All You Need", each parameter represents the weight of the connection between virtual neurons. Their complex interaction makes it impossible to trace the logic of their work .

Secondly, AI often exhibits emergent properties. This means that systems have this ability as a result of the interaction of parameters. Large language models trained on texts can, for example, translate or write program code, although they have not been specifically taught this. But even their mechanism is transparent to developers.

Thirdly, the problem lies in the data on which AI systems are trained and the hidden biases they contain. These biases can be cultural, social, or economic, and can lead to unexpected decisions by the model. An MIT study conducted in 2018 showed that facial recognition algorithms work more accurately for light-skinned men than for dark-skinned women. This is all due to the uneven and incomplete training data.

Finally, modern neural networks lack the direct causal chain typical of classical algorithms. This means that under the condition of X and Y, neural networks make decisions as a result of complex statistical correlation, and not by logic.This means that he gives different answers depending on the context.

3. Existing approaches to the solution

In computer science and in practice, several basic ways of solving these ethical problems have emerged.

The first method focuses on the transparency of algorithms. It is believed that transparency is necessary in order to be accountable for the operation of systems and for people to trust technology. As part of this approach, it is suggested to check the algorithms and disclose how they work. However, this is difficult to do because companies do not want to disclose their trade secrets, and there is also a risk that this will help attackers bypass security systems.We see a special example of this problem from Proofpoint, which deals with Cybersecurity. Proofpoint directly shows users that they do not disclose the artificial intelligence models that they use in their cybersecurity and protection services. They are justified by the fact that information about the AI models used is confidential information, and maintaining their confidentiality status is

important for the company's ability to continue to detect threats. Confidentiality also applies to training and production data related to models, requests to models and the results of their work.We can also see a similar approach from CrowdStrike. This company clearly describes Threat Graph as its own technology that uses AI algorithms, behavioral pattern analysis, and cyber threat data.These algorithms, analysis methods, and related data represent significant technological value for the company.The last example will be from FireEye. The company describes several threat detection technologies as proprietary developments. Among them, we can see examples such as anomaly detection, virtual execution, monitoring of operational stages, data correlation and behavior analysis. This means that these technologies are a significant part of its cyber threat analysis.Thus, these indicators from Proofpoint, CrowdStrike and FireEye show that cybersecurity companies protect AI models and threat detection methods. Therefore, there is a compromise between the need for transparency of AI systems and the protection of technologies that may be important for the company's security.

The second way is to develop "explicable artificial intelligence" (XAI). XAI methods explain why the AI system made a certain decision, for example, by showing which signs most strongly influenced the classification of an event as a threat. Such explanations are usually presented in the form of the significance of features, their ranking, or simple visualizations that allow the operator to understand the logic of solving the model. For example, a decision tree can represent the classification process as a sequence of simple yes/no questions, which makes the decision-making process more understandable to humans.It is important that the operator understands why the security system has decided that an event is a threat. If we compare these approaches, it turns out that there is a compromise between the accuracy of the model and its explainability. The most accurate models, such as deep neural networks, are more difficult for humans to understand. And simpler models, like decision trees, may be worse at detecting complex attacks.

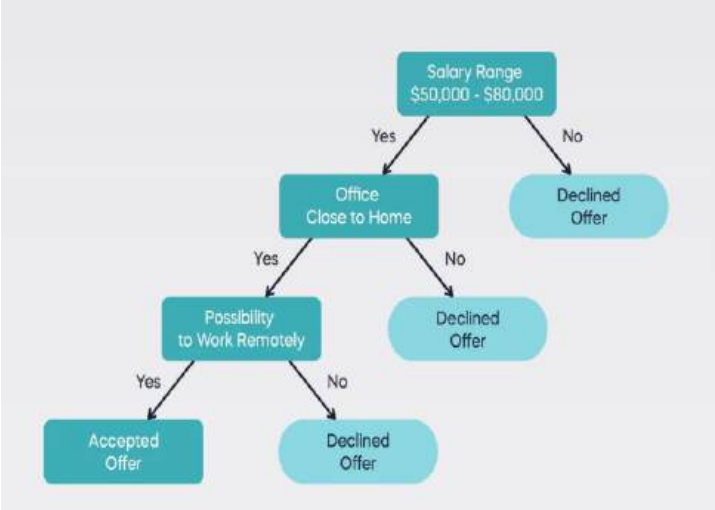

Figure 1. Practical example of a decision tree. Source: 365 Data Science, "Decision Trees."

The third way is to create ethical codes and legislative norms. They are actively discussing standards that should ensure that AI will work for the benefit of people and not violate their rights. But the problem is that these standards are often only advisory in nature, they are difficult to enforce, and they do not keep up with rapidly changing technologies.For example, a specific example of ethical standards in Europe is : "Ethics Guidelines for Trustworthy AI", published by the European Commission in 2019. A group of experts developed these recommendations and contained seven key requirements for reliable AI. These Guidelines were recommendations for AI developers and users.A more mandatory example is the European Union Artificial Intelligence Act (EU AI Act)." Unlike Ethics Guidelines, this regulatory act was officially adopted and entered into force on August 1, 2024, and the basic rules for their enforcement began to apply on August 2, 2026. The law also provides requirements related to the security and cybersecurity of certain AI systems.Finally, another example is: "Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law". The Convention was adopted by the Committee of Ministers of the Council of Europe on 17 May 2024 and opened for signature on 5 September 2024. This is an international legal treaty.

The fourth method is "design ethics". It means that ethical checks should be built into the software development process from the very beginning. For example, you need to check the data for bias before training the model and install control modules while the program is running. This approach seems promising, but it requires a lot of computing resources and specialists of different profiles: not only engineers, but also experts in ethics and law.

4. The result of the study or the proposed solution

The author proposes a comprehensive "ethical contour" model for AI systems in cybersecurity to solve existing problems. This model combines technological solutions such as explicable AI (XAI) with organizational control measures.

The main idea is to implement a two—tier system for verifying AI solutions. At the first level, a fast algorithm analyzes threats in real time to ensure a fast response. At the second level, a special module explains how the AI made the decision, making it understandable for humans. This allows the security specialist to check whether the AI's action was ethical, for example, whether the account was legitimately blocked.

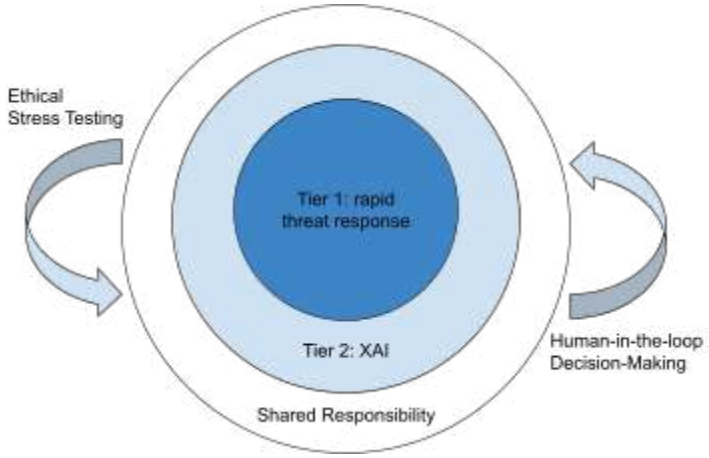

Figure 2.

The author believes that instead of searching for the "guilty algorithm", we need to move on to shared responsibility. Responsibility should be distributed depending on how autonomously the AI works. If AI makes decisions on its own, then stricter requirements should be imposed on its development and the quality of input data from the supplier.

To reduce algorithm bias, it is proposed to conduct "regular ethical stress testing." This means testing the security system for different scenarios that mimic the behavior of different user groups. The goal is to identify hidden discriminatory patterns. This approach allows us not only to talk about fairness, but also to measure it using error equality metrics.

It is also proposed to create a registry of "critical AI solutions" that must undergo mandatory human verification. Such decisions include actions that have irreversible consequences or severely restrict user access. This ensures that the person remains in the decision-making process (human-in-the-loop) and can cancel the automation action.

5. Discussion and limitations

The model that tries to integrate ethics into cybersecurity has its own problems. Firstly, when we use explainability modules (XAI), the system may start to run slower. If a cyberattack occurs, even a few milliseconds of delay in explanation can be critical and lead to hacking. It turns out that transparency, which is important for ethics, conflicts with the need to react quickly.

Secondly, the complete openness of security algorithms can be dangerous. Hackers can find out how our AI systems work and use this information to circumvent security. Therefore, the researchers are faced with the task: how to make AI understandable for our own people, but not for those who want to attack us.

Thirdly, the ethical standards themselves are rather relative. What is considered a "fair" or "acceptable risk" may vary from country to country. This makes it more difficult to create uniform international rules for AI in cybersecurity, even though the Internet and cyber threats are everywhere now.

Finally, there is a risk that we will trust AI too much. If the system almost always makes the right and ethical decisions, a person can relax and stop checking its conclusions. Then the benefits of human control disappear.

6. Conclusion

The integration of artificial intelligence into cybersecurity systems is a necessary step to counter modern threats, but it should not be carried out at the expense of degrading ethical standards and violating human rights. The analysis showed that the main challenges are related to the opacity of algorithms ("black box"), the risk of algorithmic bias and uncertainty about responsibility.

The study found that the solution to the problem lies in the creation of explicable systems (XAI) and the implementation of ethical control at all stages of the AI lifecycle. The proposed model of the "ethical contour" partially smooths out the contradictions between efficiency and transparency, but requires further technical refinement to minimize delays in the operation of protection systems.

Compliance with ethical principles in the field of AI for cybersecurity is not just a matter of moral choice, but a necessary condition for the sustainability of a digital society. It is only through ensuring transparency, accountability, and fairness that a level of trust in technology can be built that is sufficient for its safe and effective scaling in critical areas of life.

Acknowledgment

I would like to express my sincere gratitude to Emily Sheetz, PhD at the University of Michigan, for her mentorship, guidance, and valuable feedback throughout this research. I am also grateful to her for sharing relevant academic works with me.

Works Cited

Bruschi, D., & Diomede, N. (2023). A framework for assessing AI ethics with applications to cybersecurity. AI and Ethics, 3, 65–72. https://doi.org/10.1007/s43681-022-00162-8

Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 77–91.

European Commission, High-Level Expert Group on Artificial Intelligence. (2019). Ethics Guidelines for Trustworthy AI. European Commission. https://doi.org/10.2759/346720

Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (pp. 1322–1333). ACM. https://doi.org/10.1145/2810103.2813677

Kuipers, B. (2020). Perspectives on ethics of AI: Computer science. In M. D. Dubber, F. Pasquale, & S. Das (Eds.), The Oxford Handbook of Ethics of AI (pp. 420–441). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190067397.013.27

Kuipers, B. (2022). Trust and cooperation. Frontiers in Robotics and AI, 9, 676767. https://doi.org/10.3389/frobt.2022.676767

Kuipers, B. (2023). AI and society: Ethics, trust, and cooperation. Communications of the ACM, 66(8), 39–42.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). ACM. https://doi.org/10.1145/2939672.2939778

Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

Vemuri, N., Thaneeru, N., & Tatikonda, V. M. (2023). Securing trust: Ethical considerations in AI for cybersecurity. Journal of Knowledge Learning and Science Technology, 2(2), 167–175. https://doi.org/10.60087/jklst.vol2.n2.p175

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2).

European Commission, High-Level Expert Group on Artificial Intelligence. (2019). Ethics Guidelines for Trustworthy AI. European Commission.

European Commission. (2024). Artificial Intelligence Act. European Commission.

Council of Europe. (2024). Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. Council of Europe.

Proofpoint. (n.d.). Commitment to responsible development and use of AI and ML. Proofpoint.

CrowdStrike Holdings, Inc. (2025). Annual report (Form 10-K) for the fiscal year ended January 31, 2025. U.S. Securities and Exchange Commission.

FireEye, Inc. (2016). Annual report (Form 10-K) for the fiscal year ended December 31, 2015. U.S. Securities and Exchange Commission.

Crafting Catalonia: Tradition and Reform in the Architecture of Casa Batlló

By Kasin Stone Obarski

Abstract

This paper examines Casa Batlló as an example of Catalan Modernism and explores how it helped to shape modern architecture. It analyzes how organic forms, fantasy, craftsmanship, and regional symbolism reflect the evolving region of Catalonia during the nineteenth and early twentieth centuries. Through an analysis of both the region's history and specific design elements found in Casa Batlló, the paper connects architect Antoni Gaudi's work to the larger contexts of Art Nouveau and universal forms of international modernism. The historical analysis of these movements provides a deeper understanding of how different expressions of modernity emerged in response to changing cultural, political, and technological traditions. Ultimately, this paper argues that Gaudi's work should not be understood solely as an example of international modernism or through the lens of critical regionalism but as an essential blend of both.

Introduction

Along one of Barcelona's famous streets stands a building that appears to be a pool of emotions and colors. Its roof arches like the spine of a dragon, covered in iridescent scales that shimmer in the sun. These scales are luminescent rays of color that sit asymmetrically to the house, making the mosaics on the front of the house appear colorful and symmetrical. Near the bottom of the building, marble columns near the windows stretch out like bones across the front facade, spiraling and curving into shapes that resemble knotty tree limbs. On the floors above are balconies that resemble hollowed skulls with convex railings that sit central to the top of the building. Any local or visitor who walks by this piece of art is invited to experience an enchanting story that unfolds through the home's unique materials and textures.

Many interpretations connect the inspiration for Casa Batlló to the legend of Sant Jordi, the patron Saint of Catalonia. According to this medieval tale, a dragon terrorized the kingdom until Saint George defeated the beast and managed to restore peace to the land (Howard 52-56). Whether literal or symbolic, the legend reveals something essential about the building: Casa Batlló was intended to be more than a functional dwelling. It was designed as an experience for visitors and for the Batlló family, and as a cultural statement that embodies Catalan traditions.

Gaudi took over the design of Casa Batlló in 1877 from his professor, Emili Sala i Cortés, who had originally envisioned it as a more classical and sleek apartment building. Remodeled and expanded between 1904 and 1906 Casa Batlló emerged at a moment when Barcelona was redefining itself politically and culturally. Across Europe, architects interested in Art Nouveau rejected historical imitation of previous classical styles that had been popular during the past few centuries by embracing new, flowing forms inspired by nature. In Barcelona, however, this movement evolved differently. Catalan Modernism infused this new artistic experimentation with regional identity. Gaudi transformed architecture into a living display of Catalan culture, merging technological innovation with organic symbolism in ways that

challenged conventional understandings of modernity ("The Facade").

This architectural shift occurred during a period of rapid industrial development that widened class divisions in Spain. Expanding industrialization and an increased urban population created unprecedented economic opportunities while increasing tensions between workers and the upper classes. Around the same time, Spain faced a period of uncertainty following the loss of its remaining overseas colonies, such as Puerto Rico and Cuba, after the Spanish American War of 1898. This prompted debates about national identity and the country's future. In Catalonia, one of Spain's most industrialized and culturally distinct regions, broader European technological developments clashed with a growing movement to revive Catalan language, history, and culture (Mackay 36-38). The building reflects broader questions about what it means to be modern at the beginning of the twentieth century and demonstrates how architecture can serve as both a cultural symbol and a response to social transformation. Unlike the minimalism of the international modernism movement that would dominate much of the rest of the twentieth century, Gaudi's architecture embraced mythology and nature, creating an immersive space that remains influential in contemporary architecture.

This atmosphere gave rise to Catalan Modernism, a regional perspective of modern architecture that embraced many of the artistic ideals of Art Nouveau while expressing a distinctly Catalan identity. One of the movement's most celebrated monuments, Casa Batlló, emerged from that unique cultural moment. Understanding this broader context helps explain why Casa Batlló became not just an architectural landmark but a powerful symbol of Catalonia's cultural ambitions at the turn of the twentieth century.

Literature Review

Numerous scholars have examined Casa Batlló and Antoni Gaudi through the lenses of Art Nouveau and symbolism. Scholars such as Jeremy Howard & Ben Orman have explored Gaudi's use of natural motifs and handcrafted decorations. These scholars fit Casa Batlló within the broader European Art Nouveau movement, highlighting the building's rejection of rigid architecture. Their work helps explain how Gaudi drew inspiration from nature, which translates into drawings that develop into the space. Other scholars focus specifically on Catalan Modernism and its relationship to the political and cultural climate of late nineteenth-century Barcelona. Scholars such as William H. Robinson, Jordi Falgàs, Carmen Belen Lord, and Josefina Alix argue that Catalan Modernism was more than an artistic style. Rather, it emerged from the period of industrialization and cultural revival in Catalonia. These scholars emphasize the role of architecture in expressing Catalan identity and promoting a distinct regional culture during this period of growing nationalist sentiment. Within the realm of scholarship, Casa Batlló is seen as a symbolic expression of Catalan mythology and cultural pride.

This paper builds upon these scholarly conversations but steers away from viewing Casa Batlló as simply a work of art. Instead, this paper analyzes how the building reflects both change and continuity in Catalan society during an era of rapid regional transformation. The building

represents change through innovative structural systems, advanced ventilation and lighting strategies, and organic forms. At the same time, it reflects continuity through its use of traditional craftsmanship and regional symbolism. By examining how Gaudi combined innovation with historical and cultural traditions, this paper argues that Casa Batlló reveals a distinctly Catalan interpretation of modernity, one that challenges and changes the assumption that progress required a complete break from the past.

Map of paper

To understand how Casa Batlló became one of the defining monuments of the Catalan Modernism, this paper first will examine the broader political, social, and cultural contexts that shaped Europe at the turn of the twentieth century. It begins by exploring the effects of industrialization, urban growth, widening class divisions and political tensions that encouraged artists and architects to search for a new artistic language capable of expressing modern life.

From this global context, the discussion shifts to Barcelona and its participation in the broader Art Nouveau movement, comparing Catalonia's architectural developments with contemporary movements in France, Germany, and Austria. However, demonstrating how Barcelona's version of Art Nouveau became tied to regional identity. The paper then distinguishes Catalan Modernism from Art Nouveau more broadly, arguing that it represented not only an artistic movement but also a cultural and political expression of Catalan aspirations.

After establishing this historical foundation, the study turns to Antoni Gaudi himself, examining the origins of his architectural philosophy through his fascination with nature, biomorphic forms, local traditions, and historical influences, all of which he integrated into Casa Batlló to create an architecture that feels simultaneously fantastical and alive.

Finally, the paper compares Casa Batlló with the principles of international modernism, demonstrating how Gaudi's work offered an alternative vision of modernity. By exploring both innovations and continuities present in Casa Batlló, this paper demonstrates that Gaudi's architecture was not merely a product of its time moreover a forward-looking interpretation of modern design whose influence continues to resonate in contemporary architecture.

Barcelona in Transition

At the turn of the twentieth century, Barcelona was undergoing a period of heavy transformation. Rapid industrialization had reshaped cities and social structures, creating unprecedented wealth for industrial elites such as the Batlló family, who made their fortune from textile production in the late nineteenth century (Robinson 34-40). However, this simultaneously widened the gap between people like the Batllós and the working class. Across the continent, growing urban populations, labor unrest, and political tensions challenged traditional order. In Spain, the Cerdà Plan of 1859 , developed by civil engineer Ildefons Cerdà, aimed to expand the

city of Barcelona in the form of an egalitarian urban grid. This plan established the Eixample District in Barcelona, arranging buildings in a grid-like fashion (Mackay 1-3). Cerda provided a new health driven skeletal framework for the city that prioritizes movement of people. New technologies and opportunities fueled optimism about progress, yet many artists and intellectuals feared that industrial society was becoming increasingly detached from craftsmanship and regional identity. As movements such as Art Nouveau emerged throughout Europe, Spanish artists looked to them for inspiration.

In Catalonia, one of Spain's most industrialized regions, a growing movement known as the Renaixenca aimed to revive Catalan culture between the 1830s and 1890s (Robinson 22-24). This movement spread due to wider literacy and broader circulation of published materials such as like the journal *La Renaixença*, Lo Gay Saber, La Il·lustració Catalana. All of these publications attempt to revive artists and writers using photographs and writing pieces to do so. While the Renaixenca was in some ways a push for more political autonomy for the region, it was fueled by the industrial bourgeoisie, who wanted to demonstrate their growing economic power. The Renaixenca most notably changed Catalan from a marginalized language to a symbol of regional pride. One of the most remarkable people that arose from this movement is Lluís Domènech i Montaner, a Spanish architect with a background in physics and the natural sciences who became a vocal supporter of Catalan politics. In his 1878 manifesto "En busca d'una arquitectura Nacional" (In Search of a National Architecture), Domènech i Montaner argued that Catalan architects should honor academic tradition and progressive ideas by developing a style that reflected the cultural character of Catalonia. His buildings place applied arts such as stained glass, mosaics, and ceramics on the same level as architecture itself. As Barcelona expanded into a thriving industrial and cultural center, architecture became an important tool for expressing both modern ambitions and regional pride.

Art Nouveau

The Second Industrial Revolution (c.1870-1914) introduced electricity, steel-frame construction, mechanized manufacturing, and railroads that reshaped everyday life. Cities expanded rapidly as factories attracted workers from rural areas, creating new wealth for industrialists while widening economic inequality and fueling labor unrest. Architecture heavily reflected these changes as well. New materials such as iron, steel, concrete and glass made previously impossible forms achievable. Decorative elements could now be forms that curved and installations could be easier than ever. Yet many artists feared that industrial mass production was stripping cities of their local identities. In *Geographies of Modernism*, modernism emerged from this period of industrial expansion and urban transformation, It developed through cultural exchange (Huyssen). Architects across Europe therefore searched for a new artistic language capable of expressing the modern age while responding to the social consequences of industrialization.

This search produced Art Nouveau, an international movement that rejected rigid academic classicism and embraced natural forms and the collision of art and architecture.

Although unified by these ideals, Art Nouveau developed differently across Europe. In France, architects such as Hector Guimard emphasized elegant decorative forms that transformed ordinary urban infrastructure, most famously through the entrances to the Paris Metro. In Belgium, Victor Horta pioneered the use of iron and glass to create fluid, open interiors that celebrated both craftsmanship and modern engineering. Austria's Vienna Secession movement, led by artists such as Otto Wagner and Josef Hoffman, pursued a more geometric and restrained interpretation of Art Nouveau gradually moving toward abstraction and the functional principles that would influence later international modernism. And Germany's Jugendstil balanced organic ornament with increasing structural rationalism (Grady 187-92). Together, these movements reflected a shared desire to create a modern artistic identity while responding to the challenges of industrial society.

Barcelona participated fully in this international exchange while developing a distinctly regional interpretation of modernity. The city's transformation accelerated after the expansion of the Eixample district and the 1888 Barcelona Universal exposition, which introduced Barcelona as an emerging industrial and cultural capital (1888 Universal Exposition). The wealth generated from Catalonia's textile industry financed ambitious architectural commissions, allowing the bourgeoisie to express both economic success and cultural ambitions through architecture. The exposition introduced Barcelona into an international audience, demonstrating to other countries that Catalonia possessed a culture equivalent to Paris, Vienna, or Brussels (Mackay 4-8). Unlike other European Art Nouveau centers, however, Catalan Modernism became deeply connected to Catalan identity. In this sense, Catalan Modernism was not simply Barcelona's version of Art Nouveau but a unique expression of regional identity (McCarthy 379-95). This distinction provides the historical foundation for understanding why Catalan Modernism, and ultimately Casa Batlló, cannot be interpreted solely as decorative architecture but as a response to the social, political, and cultural tensions of the early twentieth century.

Catalan Modernism, or "Modernisme"

Casa Batlló embodied the aspirations of a society searching for a distinctive Catalan modernity. Architects throughout Europe rejected rigid ideas from the nineteenth century academic architecture while also resisting the impersonality of industrial mass production. Artists and architects turned towards organic forms and the integration of architecture with decorative arts, ironwork or mosaics (Mackay vii). Art Nouveau therefore represented an attempt to create a distinctly modern artistic language capable of expressing contemporary life without sacrificing beauty of craftsmanship.

Barcelona contributes heavily to this international movement, yet its architectural response developed under markedly different historical circumstances. Architecture became one of the principal means through which this ambition could be expressed. For Catalonia scholars, modernization did not need to abandon ideas from the past. Instead, it demanded reinterpreting local history within a modern context. According to some scholars, Catalan Modernism, or "Modernisme," emerged alongside broader efforts to define Barcelona as an international power

and rising cultural capital (Robinson et al. 34-40). The relationship between Catalan Modernism from other expressions of Art Nouveau can be found as distinctly different. In fact, Catalan Modernism was built off the existing trends of Art Nouveau that explores materialistic innovation with iron and glass. However, Barcelona's modernism places nationalist and regional history in the broader context of architecture in the form of materials and elements such as medieval references, gothic traditions, Christian symbolism, and Catalan folklore. M.J. McCarthy argues that Catalan Modernism cannot be separated from contemporary nationalist thought because it developed alongside efforts to preserve Catalan language and distinguish Catalonia culturally from the rest of Spain (397). In this sense, architecture became an instrument of cultural continuity as much as artistic innovation.

Understanding Catalan Modernism in this way reveals why the movement cannot be classified simply as Spanish Art Nouveau. Although it shares the international movement's admiration for organic forms, motifs, and elements of handcrafted design, this movement attaches itself to the Catalan narrative. Nature itself acquired symbolic meaning, functioning as an expression of local craftsmen. As Ben Orman suggests in *Art Nouveau & Gaudi*, nature served as a structural model instead of merely an ornamental source, allowing architecture to imitate the movement and efficiency found in living organisms.

From this perspective, Catalan Modernism represented both change and continuity. It embraced new construction, technologies and engineering methods made possible by the process of industrialization while preserving regional traditions. Rather than rejecting history in pursuit of a universal modern architecture, the movement argued that genuine modernity could arise through the reimagination of local traditions, giving new meanings. This distinction becomes especially significant when contrasted with the later principles of International Modernism, which increasingly favored universal design over regional expression.

Case Study: Casa Batlló

While growing up in Baix Camp, a coastal region in Catalonia, Gaudi was fascinated by nature, natural geometry, medieval Spanish architecture, and the mix of Islamic and Mudejar architecture that dominated the region. Unlike most architects at the time, Gaudi's professor, Lluís Domènech i Montaner, and his mentor, Joan Martorell i Montells, both explored ornate details and innovative ideas. Gaudi was inspired by their example, and his early designs reflected a deep connection to the world around him. Gaudi's first few projects, including homes like Casa Vicens and Palau Guell and the gardens at Park Guell, allowed him to experiment with ceramics, ironwork, and structural elements on a smaller scale.

By the time Gaudi created Casa Batlló, he had already refined techniques like overlap structure, ornament, engineering, and light (Mackay 21-30). This house shifts throughout the day as the sunlight moves across reflective surfaces and begins to play with the carefully engineered exterior and interior. Instead of simply decorating a conventional structure, Gaudi designed every component from the facade to the attic as a part of the larger interconnected organism whose appearance dramatically changes through the movement of light. Historical photographs and

studies done on illumination inside the complex reveal how the reflective surfaces were intended to create the impression of constant motion (Antoni Gaudi Archives). The roof's undulating ceramic tiles catch sunlight at different angles, creating the roofline's dragon-like appearance. This continues inside the house thanks to Gaudi's sophisticated use of natural light. Instead of allowing daylight to simply enter through windows, Gaudi manipulated its intensity and distribution through proportional architecture and strategic color gradients. Because light loses intensity as it descends, Gaudi used natural lighting to illuminate the central staircase creating lighting without mechanical intervention. These subtle design elements demonstrate an advanced understanding of optics and environmental engineering decades before passive environmental design became a central pillar of modern architecture.

Gaudi's structural innovations reinforce this impression while revealing his distinctive understanding of nature as a model for engineering. For example, the attic was intended as both a thermal buffer and a service space that consisted of more than sixty catenary brick arches extending across the floors. It was built entirely from lightweight brick and plaster to create a continuous white surface that feels hollow and expansive. These "empty" bricks efficiently distribute structural loads from the heavy roof mosaics while minimizing interior materials. While most architects concealed engineering behind the ornament of a building, Gaudi chose to use structural logic to become an architectural expression. This integration of function and form reflected his belief that the most efficient structures already existed in nature and that architecture should mimic some of the same principles.

Gaudi Beyond Decorations: Radical Innovations

Gaudi's work reflected his natural curiosity towards the environment around him. In addition, he combined structural innovations to completely change the atmosphere of his buildings: ventilation, light and general aesthetic. Casa Batlló was carefully engineered to allow Gaudi to push the boundaries of the materials he chose. Not only did he challenge what decorative elements were incorporated into the design but also how each element was used. Each of the developments and choices Gaudi made for the building contributes to an understanding of how he uses architecture to reflect innovation.

1. Advanced ventilation systems

   Instead of relying on machines, Casa Batlló naturally regulates temperature and airflow. Fresh air enters through adjustable wooden windows and vents while warm air naturally rises through the central lightwell and exits through roof vents. Operable vents above doors allow air to circulate comfortably while reducing humidity (Huerta). While most buildings designed in the early 1900s contained passive ventilation systems that relied primarily on standard windows and external vents, Gaudi designed Casa Batlló around an array of internal and external ventilation components that operated in harmony to maximize airflow throughout the entire structure.

2. Strategic natural lighting

Casa Batlló is carefully engineered to allow daylight to travel throughout the building. The central courtyard creates a giant light shaft that varies in color and gradient (Hinojosa 136-42). For example, the darker blue tiles near the roof compensate for the stronger sunlight while the lighter tiles down below brighten the darker levels. Window openings are also progressively smaller toward the roof because upper floors naturally receive more light.

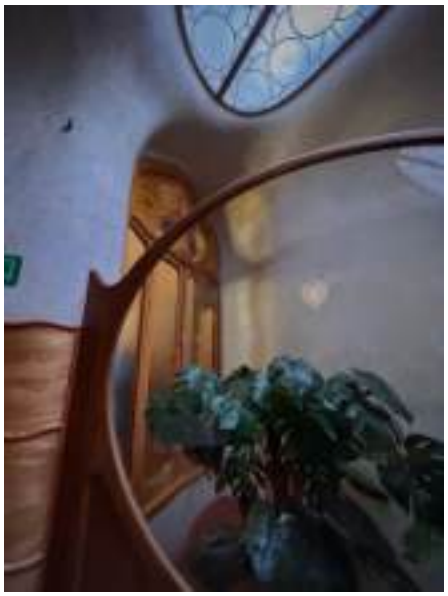

Fig. 1: Mosaics at Casa Batlló, Photograph by Kasin Obarski, 2026

3. Structural curves (Catenary engineering)

Instead of using traditional straight beams and right angles, Gaudi designed curved walls, arches, and vaulted ceilings using catenary engineering techniques. Curves distribute compression more efficiently than straight beams, this reduces material while increasing strength (Huerta). The catenary curve is based on the natural shape of a hanging chain, which evenly distributes the weight of the chain along its entire length. When inverted, it becomes one of the longest self-supporting shapes. So using precise measurements Gaudi found himself once again relying on proven mathematical formulas and theories that had also been used in classical architecture.

Fig. 2: Casa Batlló Furniture, Photograph by David Cardelús, [n.d.]

4. Trencadis Mosaic Technique

Trencadis is a mosaic technique made from broken ceramic pieces. Instead of cutting tiles into regular shapes, artisans reused discarded ceramics, glass, porcelain, and dishes. These fragments can cover complex curved surfaces that ordinary tiles could not. The technique is visually dynamic. This gives Casa Batlló's facade and roof their glowing quality. And while it's an old technique rooted in Catalan tradition, today many scholars view Trencadis as an innovative form of architectural recycling (Hinojosa 139-43).

Fig. 3: Mosaics at Casa Batlló, Photograph by Kasin Obarski, 2026

5. Sculptural Plaster Modeling

Instead of drawing and redrawing many iterations of the same mold for the facade, Gaudi made full scale plaster models by hand. He continually reshaped them until he achieved the desired form before construction. This allowed organic geometrics that

conventional drafting methods could not easily produce, to be seen in full scale before attaching to the facade. The process of construction became a part of the design process with trial and error (*Sagrada Familia*).

6. Integrated Structural Attic

   Gaudi managed to transform the attic into one of the building's most iconic interior spaces and one of engineering's greatest achievements. The attic acts as a thermal buffer, roof support, and structural spine. Traditionally attic space is used for personal storage and a collection of structural systems. Its brick arches efficiently distribute the roof's weight while also improving insulation (Fuentes & Huerta 324-39). Today, when it's not open to tourists, the attic is used for private events such as yoga classes and art exhibitions.

7. Ergonomic furniture

   Gaudi and his craftsmen designed furniture to fit the Batlló family comfortably. This meant building pieces around the exact measurements of Mr. Batlló rather than relying on standard measurements. This gives the house an additional layer of personalization. For example, wooden pieces were sculpted with curved surfaces. A wooden door knob located on the second floor and attached to the closet for Mr. Batlló was designed to fit the natural curves of his hands (Collections). Additionally the furniture mirrors the lines of the home itself. This demonstrates that furniture became an extension of not only the customer but a part of the architecture.

All these ideas developed from Gaudi's broader architectural philosophy. His early studies of geometry, gothic construction, natural science, and traditional Catalan craftsmanship gradually evolved into an approach that rejected rigid historical imitation without abandoning older techniques altogether. Instead of imitating natural forms directly, he investigated their mathematical and structural elements. Projects such as Park Guell and later Sagrada Familia expanded Gaudi's ideas using his carefully-developed key principles. In all these projects, light, structure, craftsmanship and symbolism operate together instead of independently, creating an architecture that feels dramatic and technologically advanced. In this sense, Casa Batlló illustrates why Catalan Modernism cannot be understood simply as a regional variation of Art Nouveau. The building transforms architecture into an immersive environment where engineering, handcrafted materials, and cultural imagination function as a single system. Rather than separating tangible technological innovations from intangible artistic expression, Gaudi demonstrates that innovation can be made by observation of natural processes and traditions and use of hand-crafted elements. The end result is a building that appears to be in motion, a work of architecture that continues to change with light and season and human experience, making it feel alive even a century after construction.

International Modernism

Casa Batlló demonstrates that modern architecture did not develop along a single historical path. While the twentieth century is regarded as a time of geometric forms, industrial materials, and the universal ideals of the international modern movement, Catalan Modernism offers a different view of modern standards. Rather than rejecting history and craftsmanship in pursuit of technological progress, Gaudi demonstrated that innovation could incorporate regional traditions and artistic imagination. Casa Batlló therefore represents more than an extraordinary residence. It embodies a modern architecture in which new engineering techniques and mythology can coexist.

Following the First World War, a new definition of modern architecture emerged that differed from the ideals of Catalan Modernism. Gaudi and his contemporaries had viewed modernity as a balance between technological innovation and regional identity, but a new generation of architects associated with the Bauhaus, De Stijl, and international styles argued that architecture should reflect the realities of an increasingly industrial society. As scholar Andreas Huyssen explains, international modernism developed through a series of networks circulating ideas across Europe and the Americas, creating a movement that valued universality over local traditions. Architecture became a language intended to transcend national boundaries.

This shift reflected the historical conditions of the early twentieth century. Unlike Catalonia at the turn of the century, Europe after 1918 faced widespread destruction, housing shortages, and rapid industrial expansion driven by the imperative to rebuild after the devastation of the war.

Architects such as Walter Gropius, Ludwig Mies van der Rohe, and Le Corbusier sought practical solutions that could be reproduced efficiently using “modern” materials and established construction techniques. Larger buildings could be completed faster and more cheaply using steel, reinforced concrete, and glass. Ornamentation and craftsmanship were rejected in favor of structural clarity and functional planning (Jaillant & Martin). In this context, modern architecture became closely associated with efficiency and mass production.

This difference reflects two competing interpretations of innovation. International modernists believed that architecture should respond to the demands of an increasingly industrial and interconnected world through standardization and universal design. Catalan Modernism, on the other hand, argued that modernization should remain rooted in regional traditions. Gaudi believed that cultural identity was not an obstacle to progress but instead that it could be integrated.

Although international modernism ultimately became the dominant architectural movement of the twentieth century, many of its principles have since been reconsidered. Contemporary architecture increasingly values sustainability alongside environmental responsiveness and human centered design. Casa Batlló therefore challenges the idea that modern architecture followed a single path.

Conclusion

Casa Batlló initially appears to be a building defined by fantasy: a facade that reflects a story alongside its beauty, and a roof that mimics the scales of a dragon. At first these features can be dismissed as decorative elements inspired by the imagination of Antoni Gaudi. However, examining Casa Batlló through the lens of history reveals that its unusual appearance represents something much larger. The building emerged from a period when Barcelona was modernizing and transforming. In response to these changes, artists like Gaudi created Catalan Modernism to reflect a unique, regionally-inspired definition of what it meant to be "modern" (Metropolitan Museum of Art). Casa Batlló embodies this definition by bringing together the past and the future rather than choosing between them.

Casa Batlló's influence extends beyond its historical period or its preservation as a famous building. It continues to engage people because of its allusion to nature, use technology without becoming dependent upon it, incorporate regional identity, and create meaningful experiences for the people who occupy them. Ideas such as biomimicry, passive environmental design, sustainability, and human centered architecture are now central to contemporary architectural discussions. Although the technological circumstances have changed, the underlying questions are remarkably similar to those Gaudi explored at the beginning of the twentieth century.

Casa Batlló therefore demonstrates that modern architecture should not be understood as a single progression toward industrial production. Catalan Modernism represents another path, one in which innovation developed through the continued relationship between history and change. By examining Casa Batlló through the political context of Catalan nationalism, the artistic context of Art Nouveau and its contrast with international modernism, this paper has shown that Gaudi's architecture should not be seen solely as a decorative precursor to modernism. Instead, Casa Batlló was a compelling modern project in its own right. Its integration of nature, technology, craftsmanship, mythology, and cultural identity created an alternative understanding of architectural progress. The building therefore survives not only as evidence of what Barcelona was, but as an example of how architecture can carry the past forward. More than a century later, Casa Batlló continues to ask the same question that shaped its influence: How can architecture embrace change without losing identity, memory, and imagination that gives a location its purpose?

Works Cited

"1888 Universal Exposition (1885-1892)." Arxiu Municipal Contemporani de Barcelona, 19 Dec. 2023, Accessed June 15, 2026. https://ajuntament.barcelona.cat/arxiumunicipal/arxiucontemporani/en/fonds-and-collections/institutional-fonds/1888-universal-exposition-1885-1892.

"Barcelona and Modernity." The Metropolitan Museum of Art, 7 Mar. 2007, Accessed Feb 25, 2026. https://www.metmuseum.org/exhibitions/listings/2007/barcelona-and-modernity.

"Casa Batlló." Dosde.Com, 2020, Accessed Feb 25, 2026. https://www.dosde.com/discover/en/casa-batllo/?utm_.

"Casa Batlló Gaudi Barcelona." Casa Batlló, 2026, Accessed May 23, 2026. https://www.casabatllo.es/en/category/antoni-gaudi/https:/www.casabatllo.es/en/category/antoni-gaudi/.

Casa Batlló. (n.d.). Antoni Gaudí archives. Accessed November 2nd, 2026. https://www.casabatllo.es/en/category/antoni-gaudi/

Coral Hinojosa, Sebastián. "The Light Well as Element of Luminous Control inside the Architectural Spaces, Case of Study: Batlló House." Estoa, vol. 7, no. 13, June 2018, pp. 135–43. Accessed June 15, 2026. https://doi.org/10.18537/est.v007.n013.a11.

Grady, James. "Nature and the Art Nouveau." The Art Bulletin, vol. 37, no. 3, Sept. 1955, pp. 187–92. Accessed May 23, 2026. https://doi.org/10.1080/00043079.1955.11408304.

Gómez-Moriana, Rafael. "Revisit: Casa Batlló by Antoni Gaudí in Barcelona, Spain - the Architectural Review." The Architectural Review, 8 Jan. Accessed Feb 25, 2026. https://www.architectural-review.com/essays/revisit/revisit-casa-batllo-by-antoni-gaudi-in -barcelona-spain.

H. Robinson, William, et al. Barcelona and Modernity. 2006. Yale University Press, pp. 34–40, Accessed Feb 25, 2026. https://books.google.com.pr/books?hl=en&lr=&id=6EvIx6zOuqgC&oi=fnd&pg=PR8&dq=Catalan+Modernism+history&ots=eixVp-35QY&sig=20AQ6MlLx_Qg4Us53oejQ9eNo4&redir_esc=y#v=onepage&q=Catalan%20Modernism%20history&f=false.

Heritage Assets - Sagrada Familia. Unesco: World Heritage, 2025, Accessed July 18, 2026. https://sagradafamilia.org/en/heritage-assets.

Howard, Jeremy. Art Nouveau. Manchester University Press, 1996, pp. (52–56), Accessed May 23, 2026. https://books.google.com.pr/books?hl=en&lr=&id=r-DRh4O0qtMC&oi=fnd&pg=PR8&dq=How+does+Casa+Battlo+(particularly+its+merger+of+the+organic+and+fantastical)+typify+Catalan+Modernism+(and+contrast+with+International+Modernism%3F+What+is+the+origin+and+importance+of+this+style.

Fuentes, Paula, and Santiago Huerta. "Geometry, Construction and Structural Analysis of the Crossed-Arch Vault of the Chapel of Villaviciosa, in the Mosque of Córdoba." International Journal of Architectural Heritage, vol. 10, no. 5, May 2015, pp. 589–603.

Huerta, Santiago. "Structural Design in the Work of Gaudí." Architectural Science Review, vol.

49, no. 4, Dec. 2006, pp. 324–39, https://doi.org/10.3763/asre.2006.4943. Accessed 2 Nov. 2026.
Huyssen, Andreas. “Geographies of Modernism in a Globalizing World.” Geographies of Modernism, May 2007, pp. 16–28, Accessed Mar 15, 2026. https://doi.org/10.4324/9780203392362-8.
Jaillant, Lise, and Alison E. Martin. “Introduction: Global Modernism.” Modernist Cultures, vol.13, no. 1, Feb. 2018, pp. 1–13.
Mackay, David. Modern Architecture in Barcelona. 1985, Accessed May 23, 2026. https://www.anglo-catalan.org/downloads/acsop-monographs/issue03.pdf.
McCarthy, M. J. “Catalan ‘Modernisme,’ Messianism and Nationalist Myths.” Bulletin of Hispanic Studies, vol. 52, no. 4, Oct. 1975, p. 379.
Orman, Ben. “Art Nouveau & Gaudí: The Way of Nature.” JCCC Honors Journal, vol. 4, no. 1, Feb. 2013.
Sen, Furkan. “Casa Batlló: Gaudí’s Organic Modernisme Masterpiece.” Illustrarch, Accessed Mar 18, 2026. https://illustrarch.com/articles/73976-casa-batllo-gaudi.html?utm_#google_vignette.
“The Façade | Casa Batlló.” Casa Batlló, 2017, Accessed Mar 17, 2026. https://www.casabatllo.es/en/antoni-gaudi/casa-batllo/facade/.
Universitat Politècnica de Catalunya. Collections. Universitat Politècnica de Catalunya, Accessed June 16, 2026. https://catedragaudi.upc.edu/en/batches.
Vélez, Pilar. “‘Santiago Rusiñol.’” Transfer: Journal of the International Association of Catalan Writers, no. No. 3, 2008, Institut Ramon Llull / Llull.cat, pp. 98–111.

Investigating the Relationship Between Personality Traits, Social Media Use, and Mental Health Outcomes: A Literature Review By Pak Ying Shannon Tsang

Abstract

This review paper examines the relationship between the Big Five personality traits, social media use, and associated mental health outcomes. The main sources of data used in this review were peer-reviewed academic journal articles located using Google Scholar and ResearchGate, published between 2016 and 2026. The findings from this review demonstrate that certain personality traits are associated with certain patterns of social media use (SMU) and different levels of vulnerability to the mental health outcomes of SMU. Specifically, it finds that SMU is related to various mental health outcomes, such as depression and anxiety. Furthermore, it finds that personality traits are linked to different patterns of SMU, where extraversion and agreeableness were found to be associated with SMU for social interaction, openness was linked to SMU for information access, and neuroticism was linked to passive and problematic SMU. Finally, it was found that certain personality traits exhibited a greater vulnerability to negative mental health outcomes of SMU, where low conscientiousness and high neuroticism were consistently linked to problematic social media use (PSMU) and social media addiction (SMA). These findings highlight the need for mental health policies, targeted interventions, and future research to implement personality-specific responses that focus on at-risk users, aiming to prevent and treat mental health outcomes associated with specific personality traits and usage patterns, as well as define more causal relationships through longitudinal studies.

Keywords: Social Media Usage, Big Five Personality Traits, Mental HealthOutcomes, Vulnerability, Problematic social media use

Introduction

Social media is becoming increasingly ubiquitous, with up to 95% of youth (ages 13–17) reporting social media use (SMU), and more than a third saying they use social media "almost constantly" (Vogels et al.). It serves many beneficial purposes, including acting as a platform for people to maintain relationships, boost social interaction, and promote learning through access to various resources (Khan et al.). However, it has also been connected to increased mental health issues, namely, heightened risk for social anxiety, depression, and stress (Ventriglio et al.).

These effects, however, are not experienced equally, suggesting that individual differences may influence how social media affects psychological wellbeing. Some research suggests that personality traits play a significant role in shaping the types of platforms people use, the way they engage with them, and the impact they have on their mental health (Kircaburun et al.; Whaite et al.). Furthermore, different personality traits and patterns of SMU are associated with a range of mental health outcomes (Kircaburun et al.; Hawi and Samaha). Understanding how personality traits may mediate the relationship between SMU and mental health could help identify individuals who are more vulnerable to the negative mental health effects of social media.

This review paper will investigate how personality traits influence social media's impact on mental health outcomes. The existing literature exhibits a range of complex and conflicting findings relating to this topic e.g., Gil de Zúñiga et al.; Andrews et al.). Therefore, this review aims to synthesise and critically evaluate existing evidence on the relationship between SMU and mental health outcomes, personality traits and SMU, and personality traits and vulnerability to SMU, in order to create a more coherent picture of the findings. In doing so, this review aims to inform the development of policies and targeted interventions to improve the long-term mental health of social media users. The paper investigates whether some personality types are more vulnerable to the negative mental health outcomes of SMU.

The main body of this paper will first outline what social media is, including the different ways usage can be defined and measured. It will then move on to defining personality, focusing specifically on the Five Factor Model, and the way it is assessed. Then, the paper will proceed to address the three key research questions:

1. How does SMU relate to mental health outcomes?
2. How are personality traits related to patterns of SMU?
3. How do personality traits relate to vulnerability in response to SMU?

Methods

This narrative review analyses existing literature that explores the relationship between Big Five personality traits, SMU, and mental health outcomes. The main sources of data were peer-reviewed academic journal articles located using Google Scholar and ResearchGate. The keywords used for the literature search included "social media use", "mental health", and "personality traits". Articles were selected based on their relevance to the research question and their date of publication. Papers were only selected if they had been published within the last decade (2016-2026) and mentioned the Big Five personality traits. This is due to the evolving nature of SMU, which has changed noticeably in recent years in terms of the diversity of its platforms, personalised feeds, and means of communication (Dhingra and Mudgal). Therefore, this review prioritises recent studies to ensure findings are relevant and are able to reflect the current situation.

Main Body

Social media

Social media is made up of various user-driven platforms that facilitate the diffusion of compelling content, dialogue creation, and communication to a broader audience (Kapoor et al.). It has evolved, with earlier forms focusing mainly on communication and networking, and more recent forms focusing on content creation and sharing (Faverio and Sidoti). Common platforms include Instagram, TikTok, YouTube, and Facebook, which differ in the types of interaction they encourage and the ways users engage with them (Kietzmann et al.). While the general term "social media" encapsulates this broad range of platforms and activities, various studies in the literature have conceptualised and measured social media in distinct ways.

Different measures of SMU are employed in the literature. Many examine *time spent*

on social media, which refers to the self-reported total time per day spent on social media platforms for personal use (Lin et al.; Primack et al.; Berryman et al.). Others examine *frequency of SMU*, which refers to how often one visits social media platforms (Shensa et al.).

Moreover, others examine different types of engagement with social media, for example, passive vs. active use, and problematic use. Passive and active SMU are measured using self-report questionnaires about the frequency of the following items: for passive use, participants reported frequency of ‘‘reading discussions’’, ‘‘reading comments/reviews’’, and ‘‘watching videos or viewing pictures’’, and for active use, participants reported frequency of ‘liking/favoriting/voting’’, ‘‘sharing others’ content’’, ‘‘commenting on or responding to someone else’s content’’, and ‘‘posting your own content’’ (Escobar-Viera et al.).

Problematic SMU (PSMU) is characterized by excessive SMU that significantly disrupts daily functioning and is measured using various self-report scales such as the Social Media Use Questionnaire (SMUQ) in the Kircaburun et al. study, Bergen Social Media Addiction Scale (BSMAS) in the Ahmed et al. study, Social Media Addiction Questionnaire (SMAQ) in the Hawi and Samaha study, and finally Problematic Facebook Use (PFU) in the Marino et al. study.

Personality

Personality, defined as "the relatively enduring patterns of thoughts, feelings, and behaviors that reflect the tendency to respond in certain ways under certain circumstances" (Roberts), can shape one’s social behaviour and the way people interact with the world around them (Back et al.). Personality is commonly assessed using the five-factor model of personality (FFM), commonly referred to as the “Big Five”, which includes Extraversion, Agreeableness, Conscientiousness, Neuroticism, and Openness to Experience (Soto and Jackson). It has been continuously developed throughout the years by different researchers, namely Goldberg (1993), McCrae and John (1992), and John (2008). Each trait is characterised by different facets: extraverted individuals are more assertive and sociable, agreeable individuals are more cooperative and polite, conscientious individuals are more task-focused and orderly, neurotic individuals are more prone to experiencing negative emotions, and open individuals tend to have a broader range of interests. This model was chosen for this review, primarily due to the fact that it is the most widely used and empirically supported model of personality, providing a clearer rationale for focusing on these traits.

Q1 - How does SMU relate to mental health outcomes?

SMU has been associated with a range of detrimental mental health outcomes, including depression, suicidal ideation, loneliness, and social anxiety (Lin et al.; Berryman et al.). These associations are not consistent across all studies, where some report positive correlations, and some report negative correlations. The different usage measures used in these studies include *time spent* on social media, *frequency of use*, and *passive vs. active use* (Lin et al.; Primack et al.; Shensa et al.; Escobar-Viera et al.).

Several studies report a positive association between *time spent* on social media and depressive symptoms (Lin et al.; Primack et al.). Lin et al. found that *total time per day spent* on social media was significantly associated with increased depression, with participants in

the highest quartile of total time per day spent on SM having significantly higher odds of depression compared to those in the lowest quartile. This is supported by Primack et al., who built on the findings of Lin et al. by employing a longitudinal design to examine the temporal relationship between these two factors. The study found that baseline SMU, measured by average total daily *time spent* on social media for personal use, predicts later increases in depressive symptoms, as evidenced by the fact that those in the highest quartile of SMU had 3.41 times the odds of developing depression compared with those in the lowest quartile. Furthermore, the lack of reciprocity in this relationship suggests that increased SMU may precede increases in depressive symptoms.

Primack et al. proposed several mechanisms that may explain the relationship between *time spent* on social media and depression. They identified that time occupied by SMU meant less time available for forming in-person relationships, achieving goals, and self-reflection. Additionally, social comparison on SMU, through exposure to unattainable expectations and images, may evoke comparison and depressive cognitions, particularly in adolescents who are at a critical time of their development. Lastly, increased exposure to social media may interfere with neurocognitive processes, as normal developmental neurocognitive processes, such as social cognition and social reward processing, may be disrupted by social media.

However, there are inconsistencies in the literature that weaken the claim that social media usage is a strong predictor of negative mental health outcomes in certain cases. While Lin et al. and Primack et al. collectively found that increased SMU may predict increases in depressive symptoms, Berryman et al., who employed a similar measure of *time spent* online, reported that social media variables including *time spent* online and social media importance (the degree to which participants were emotionally connected to SMU) was not linked to outcomes such as suicidal ideation, loneliness, social anxiety and decreased empathy.

Methodological differences between the studies may account for the discrepancies in the findings. Studies used different sampling techniques, namely, Lin et al. used random sampling, whereas Berryman et al. focused on a population of undergraduate students attending a university in the southeastern region of the United States. This may limit the generalisability of the findings, as university students are not representative of the wider population and may differ in terms of age, education, and current life circumstances.

The different mental health outcome measures used in the studies may also contribute to inconsistent results. Specifically, Lin et al. and Primack et al. both focused on depression, assessing it using the Patient-Reported Outcomes Measurement Information System (PROMIS) Depression Scale Short Form and the 9-Item Patient Health Questionnaire, respectively. However, Berryman et al. focused on different constructs relating to mental health, such as suicidal ideation, loneliness, social anxiety, and decreased empathy, using different measures such as the Brief Symptoms Inventory-18 to assess general mental health symptoms, the Liebowitz Social Anxiety Scale-Self-Report Version to assess social anxiety, and the UCLA Loneliness Scale-3 to assess loneliness. Due to the different constructs captured in each of the scales, for example, depression scales capture a different construct from measures of loneliness or social anxiety, this makes direct comparisons between studies more difficult, as they may be impacted differently by social media.

In addition to differences in mental health outcome measures, Shensa et al. suggest that the inconsistencies in the literature may also stem from the way SMU is conceptualised and measured. This study assessed *frequency of visits* (to eleven social media platforms: Facebook, Twitter, Google, YouTube, LinkedIn, Instagram, Pinterest, Tumblr, Vine, Snapchat, and Reddit) in addition to total *time spent* on platforms per day for personal use. It concluded that while *frequency of visits* to these platforms was significantly associated with depression, total *time spent* on SM was not significantly associated with depression. This distinction suggests that the relationship between SMU and depression is more multifaceted than just the duration of *time spent* on sites, but rather the way people engage with social media platforms.

Moreover, passive versus active SMU offers a different conceptualisation of engagement on social media (Escobar-Viera et al.). The study by Escobar-Viera et al. analysed passive and active SMU and looked at their relationship with depression, concluding that each one-point increase in passive SMU was associated with a 33 percent increase in depressive symptoms. In contrast, each one-point increase in active SMU was associated with a 15 percent decrease in depressive symptoms (Escobar-Viera et al.). This discrepancy with previous studies could be explained by the fact that passive lurking, scrolling, and low engagement with others are linked to decreased well-being, increased feelings of exclusion, marginalisation, and increased social anxiety; whereas active SMU involving interactions and sharing of content can instead improve well-being and promote self-expression.

Ultimately, methodological differences (measures of SMU, samples, mental health outcome measures) contribute to conflicting findings in the literature. The link between SMU and depression varies, with some studies reporting a strong positive correlation (Lin et al.), and others reporting non-significant correlations between the two variables (Shensa et al.). Considering the studies used different measures, this indicates that how SMU is measured and conceptualised is just as important as how much it is used, as SMU is multifaceted. Specifically, some measures of SMU are more subjective than others; for example, time spent can be objectively measured, whereas the way people interact with comments takes a more personal reflection and could potentially be subject to more self-reporting bias. Furthermore, most studies present correlational evidence that SMU and mental health outcomes such as depression are linked, whereas longitudinal studies that provide stronger evidence for temporal relationships are lacking in the literature. Therefore, it is crucial that future research implements more consistent measures to better represent the relationship between SMU and mental health outcomes, and apply longitudinal designs to clarify causal relationships.

Q2 - How are personality traits related to patterns of SMU?

Research on the relationship between personality and SMU suggests that personality traits are associated with different patterns of SMU. Specifically, studies have examined how personality correlates with overall *frequency of use*, *time spent online*, *motives for SMU*, and different types of interaction, including *social engagement* and *news consumption* (Gil de Zúñiga et al.; Andrews et al.; Choi et al.; Kircaburun et al.).

Personality traits have been shown to be correlated with *frequency of use* throughout the literature, where all five traits predict specific patterns of SMU. Gil de Zúñiga et al. found

that the *frequency of social media site visits* was positively correlated with extraversion, agreeableness, conscientiousness, and openness to experience, whereas emotional stability was negatively correlated with the *frequency of social media site visits*. Gil de Zúñiga et al. offered a range of possible explanations for this pattern of associations. Specifically, individuals with extraverted traits who tend to feel more comfortable talking to others and are keen to meet new people with shared interests may approach SMU with a more socially engaging outlook. Similarly, individuals high in openness are more prone to searching for and pursuing new content and connections with others. Furthermore, agreeable and conscientious individuals may use social media more to socialise, due to their tendency to sympathise with others and comfort others. Conversely, emotionally stable individuals tend to use social media for socialisation purposes less, as they feel secure in their real-life relationships, thereby explaining why they visit sites less frequently (Gil de Zúñiga et al.).

However, due to the cross-sectional design of the study, it fails to address the directionality of the relationship between personality and SMU. This means that it cannot be established what is driving this relationship, and whether SMU influences personality or if personality is influencing SMU. This places limitations on the findings, as the sequence of events is unclear; therefore, the study cannot determine the temporal direction of the relationship.

Extending on the previous finding, Andrews et al. implement a longitudinal design, helping it overcome some of the limitations of the cross-sectional study by providing stronger evidence about the direction of the relationship between SMU and personality traits. In this case, however, he focused on total *time spent* on social media, as opposed to *frequency of site visits*. The study found that more *time spent* on social media predicted greater neuroticism, and similarly, increased neuroticism predicted more *time spent* on social media. This bidirectional relationship highlights the vicious cycle of the relationship between neurotic individuals and their usage of social media. While this study found a clear reciprocal association between neuroticism and SMU, it found no significant relationships between the other Big Five personality traits (extraversion, openness, agreeableness, conscientiousness) and *time spent* on social media.

This difference in findings may be the result of Andrews' and Gil's different measures of social media usage, where one focuses on total *time spent* and the other on *frequency of social media site visits*. These different measures likely capture different types and intensities of engagement with the platforms. This suggests that distinguishing between these different types of engagement may be crucial in strengthening understanding around the link between personality and SMU.

Personality not only relates to total *time spent* or *frequency of site visits* on social media, but also expands to a larger spectrum of uses, such as specific types of interaction on social media. A study by Choi et al. specifically looked at engagement around the posting of 'selfies' on social media. It reported that all Big Five traits, except for extraversion, were significantly associated with different degrees of concern about others' reactions to one's selfies. Specifically, individuals high in agreeableness show a greater degree of concern about others' responses to their selfies, as they tend to prioritise positive relationships online. Furthermore, those high in

conscientiousness and neuroticism were associated with higher degrees of concern about others' reactions to their selfies. This could be explained by conscientious individuals being more cautious about their online self-presentation, and neurotic individuals having a greater sensitivity to rejection, prompting them to seek validation and feedback online. Lastly, those low in openness were associated with a higher degree of concern for others' responses to their selfies, as they may spend more time seeking validation rather than exploring topics online. In addition, extraversion and agreeableness predicted the tendency to comment on or like others' selfies, where extraverts specifically were more likely to leave comments and like others' selfies.

Furthermore, personality traits are also associated with different motives for using social media, as evidenced in two studies that observed this relationship (Gil de Zúñiga et al.; Kircaburun et al.). Between the studies, it was noted that extraversion was linked to higher levels of SMU for social interaction, which lines up with extraverted individuals' main uses of social media being to maintain relationships. Moreover, Choi et al. found that high agreeableness was positively associated with all forms of online social connection, whereas low openness was negatively associated with online social connection. This suggests that more agreeable individuals may use social media in more interactive and relational ways, whereas those lower in openness may be less engaged in social interaction. Supporting this, Gil de Zúñiga et al. also found that high agreeableness was linked to higher levels of SMU for social interaction, whereas low openness was linked to lower levels of SMU for the equivalent activities.

In contrast, other traits show more nuanced patterns of use. Kircaburun et al. found that those high in openness in fact used social media for more educational and informational purposes, providing an explanation as to why they may interact less on social media. Additionally, high neuroticism was linked to both reduced SMU for social interaction and increased use for passing time. Lastly, conscientiousness was linked to various patterns of SMU, where high conscientiousness was linked to greater use of social media for news consumption and social interaction, and low conscientiousness was linked to greater use of social media for expressing or presenting a more popular self (Gil de Zúñiga et al.; Kircaburun et al.).

Ultimately, different personality traits are associated with different patterns of SMU, where frequency of site visits is positively correlated with extraversion, agreeableness, conscientiousness, and openness, but negatively correlated with emotional stability (Gil de Zúñiga et al.), and total time spent on SMU presented a bidirectional relationship with neuroticism (Andrews et al.); there was also no significant relationship found between personality traits and total time spent. These conflicting findings could be explained by the use of different SMU measures, which focused on different types of engagement with social media. Furthermore, specific behaviors such as concern about others' reactions to individuals' selfies were associated with all Big Five traits excluding extraversion, and specific motives for SMU were associated with each personality trait. This highlights the need for future research, particularly studies with longitudinal designs to provide stronger evidence for causal relationships between patterns of SMU and personality traits, allowing us to better identify and implement social media management strategies catered to specific personality types, as SMU is

not uniform.

Q3 - How do personality traits relate to vulnerability in response to SMU?

In addition to different personality traits being associated with different patterns of SMU, specific personality traits also appear to influence individuals' vulnerability to the negative outcomes associated with SMU, particularly PSMU, which has been associated with detrimental outcomes such as addiction, social isolation, and depression.

PSMU and addiction

Several studies in the literature have investigated the relationship between personality traits, problematic social media use (PSMU), and social media addiction (SMA). While some personality traits have been consistently associated with an increased or reduced risk of PSMU and SMA, findings for others remain mixed. PSMU is characterized by excessive SMU that significantly disrupts daily functioning (Kircaburun et al.), and SMA is characterised by obsessive and excessive SMU (Hawi and Samaha).

Across these studies, low conscientiousness was found to be consistently linked to PSMU and SMA (Kircaburun et al.; Ahmed et al.; Hawi and Samaha). It has been suggested by Kircaburun et al. that this could be explained by the fact that less conscientious people tend to be less organised, ambitious, and consistent. Due to the portrayal of these undesirable traits in real life, it may prompt less conscientious individuals to excessively self-promote a better self online.

Similarly, high neuroticism appeared to be frequently linked to PSMU and SMA, acting as a significant predictor in three of the four studies (Kircaburun et al.; Hawi and Samaha; Marino et al.). This could be attributed to neurotic individuals' tendency to avoid real-life social interactions due to feelings of social anxiety and fear of rejection, turning to social media as an escape. Impulsivity, a trait of neurotic individuals, exacerbates this, making it more difficult to resist checking social media platforms and consequently spending time on apps, thereby leading to excessive SMU.

The evidence for the relationship between mental health outcomes and the remaining personality traits is less consistent across studies. Low extraversion was linked to PSMU in two studies. This could be explained by the fact that people low in extraversion tend to have lower self-esteem and higher loneliness, both of which are positively associated with passive engagement, which can lead to obsessive checking of others' profiles to avoid real-life struggles and social interaction. This pattern of behaviour can easily become excessive as individuals low in extraversion may become reliant on the lower levels of stress and social pressure it offers (Kircaburun et al.; Marino et al.). Additionally, they may prefer online interactions and maintain these relationships online, resulting in more *time spent* engaging in PSMU. However, the remaining studies found no significant association between low extraversion and PSMU.

Openness to experience was also associated with PSMU in two studies, where low openness to experience was identified as a predictor of problematic use, specifically SMA and PFU (Hawi and Samaha; Marino et al.). This could be explained by their disinclination to engage in real-world activities, which are likely to result in a greater variety of experiences than spending time online. However, the remaining studies found no significant association between

openness and PSMU.

The relationship between agreeableness and PSMU presents the most surprising contradiction, where only three of the four studies reported that low agreeableness predicted PSMU (Ahmed et al.; Hawi and Samaha; Marino et al.). This can be explained by the fact that agreeable individuals may struggle with maintaining real-life relationships, turning to social interaction and validation online. However, Kircaburun et al. directly contradict this, finding that higher agreeableness was linked to PSMU.

This discrepancy in findings could possibly be explained by differences in measures of PSMU, samples, and platforms across studies. While some studies used addiction-focused measures, such as the Bergen Social Media Addiction Scale, Internet Addiction Test (IAT), and Social Media Addiction Questionnaire (SMAQ), others used broader measures of PSMU, including the Social Media Use Questionnaire (SMUQ) and PFU scales (Ahmed et al.; Hawi and Samaha; Marino et al.; Kircaburun et al.). The use of different measures may have influenced the findings, as they target different aspects of mental health outcomes that may not be consistently observed in all the studies, making it difficult to generalise findings. Furthermore, while most of the studies looked at a range of social media platforms, Marino et al. only observed PSMU on Facebook, which may be too narrow and possibly outdated to represent the general relationship between traits and vulnerability accurately.

Social isolation

Some studies have examined the relationship between personality and more specific mental health outcomes associated with SMU, such as social isolation. Different levels of conscientiousness appear to influence individuals' vulnerability to social isolation (SI) as a result of SMU. A cross-sectional study found that individuals who acted less conscientiously experienced a greater risk of SI due to SMU, as opposed to those who acted more conscientiously (Whaite et al.). Whaite et al. provide an explanation for this, suggesting that conscientious individuals tend to be more organised, diligent, and have good impulse control, possibly using social media as a way to organise real-life interactions instead of scrolling obsessively and isolating oneself from society. In this case, they use social media in a way that supports the maintenance of good face-to-face interactions, thereby preventing SI. Furthermore, individuals higher in conscientiousness tend to interact more carefully on social media, which could protect themselves from negative interactions that may contribute to SI.

Furthermore, a study by Chi et al. looks at another aspect of social isolation, investigating the relationship between personality traits and phubbing behavior (the act of ignoring one's immediate social interactions in favor of engaging with a smartphone). It found that conscientiousness was most inversely correlated with phubbing. This could be explained by individuals who are less conscientious exhibiting poorer discipline and organisation, resulting in communication disturbances and phubbing behaviours in social settings.

Depression

While personality appears to predict vulnerability for certain outcomes such as PSMU and SI, it does not appear to predict vulnerability to depression relating to SMU as consistently.

A longitudinal study investigating the development of depression over six months found that the strength of the association between SMU and depression did not significantly vary among people with or without specific personality characteristics (Merrill et al.).

Differences in the conceptualisation of SMU may account for this discrepancy in findings. It is important to note that personality may be more important for behavioural outcomes such as PSMU, addiction, phubbing and SI, rather than for long-term mood disorders such as depression. This can be explained by the fact that behavioral outcomes are more closely related to personality traits; for instance, individuals low in conscientiousness are more likely to use social media impulsively, whereas a person high in neuroticism may use social media more passively. In contrast, depression is shaped by a wide range of psychological, biological, and social factors, arising not only from SMU; therefore, the relationship is more obscure. While they may appear to be related, these outcomes are distinct constructs; therefore, it is unsurprising that certain personality traits predict behavioural outcomes more consistently than a complex mood disorder like depression.

Furthermore, methodological differences may have also contributed to conflicting findings. Particularly, the study by Whaite et al., which found that the occurrence of SI differed in specific personality types, was based on a cross-sectional design, whereas the longitudinal study examined the development of depression over a period of 6 months. Cross-sectional studies can identify associations; however, longitudinal designs are better able to examine temporal relationships and changes over time. Additionally, the use of different measures of SMU (daily time spent on SM versus frequency of use) may have also affected results.

Overall, it can be concluded that certain personality traits predict levels of vulnerability to mental health outcomes relating to SMU. Evidence for this relationship is not consistent across personality traits, where low conscientiousness and openness, as well as high neuroticism and low extraversion, appear to increase vulnerability to both PSMU and SMA (Ahmed et al.; Hawi and Samaha; Marino et al.; Kircaburun et al.), whereas some findings remain mixed, particularly the relationship between low agreeableness and PSMU, which is directly contradicted among the studies, and the finding that all five personality traits did not appear to predict vulnerability to depression entirely (Merrill et al.). This suggests that certain personality traits may present a higher risk factor in relation to SMU and mental health outcomes, including PSMU, SMA, and SI, highlighting traits that may require more attention to prevent these detrimental mental health outcomes. These findings highlight the importance of conducting further research and longitudinal studies to better understand the causal relationship between personality traits and vulnerability to negative mental health outcomes, so that better interventions, prevention measures, and treatment can be implemented.

Conclusion

This review aimed to synthesise findings in the literature relating to the relationship between personality traits, SMU, and associated mental health outcomes. It focused on analysing how SMU relates to mental health outcomes, how personality traits can predict those patterns of SMU, and finally how traits may relate to vulnerability in response to SMU.

Across the literature, it was observed that SMU is not inherently harmful, but rather its relationship with negative mental health outcomes depends on patterns of usage, the outcomes that are measured, and which personality traits an individual shows. Firstly, it was found that time spent on social media was associated with depressive symptoms in some studies, but was insignificant in others. This was due to differences in the measures used, including frequency of visits and passive vs. active use, which resulted in varying outcomes, highlighting the multifaceted nature of SMU. Secondly, it was found that personality traits were linked to different patterns of SMU, particularly extraversion and agreeableness, which were generally associated with increased SMU for social interaction; openness, which was linked to access to information; and neuroticism, which was linked to passive use and problematic forms of SMU. Lastly, certain personality traits showed greater vulnerability to negative mental health outcomes relating to SMU. Specifically, low conscientiousness and high neuroticism were repeatedly linked to PSMU and SMA, while findings for the other traits were less consistent.

While it is evident that SMU is associated with negative mental health outcomes, especially when it is characterised by passive, problematic, and compulsive use, and that neuroticism and conscientiousness show consistent links with vulnerability, there are still areas where findings remain inconclusive. Namely, it is unclear whether time spent accurately predicts depression, whether agreeableness predicts risk of PSMU, and the directionality of these relationships.

The findings of this review present several implications. It suggests that interventions, prevention measures, and mental health policies that view SMU broadly may be less effective in targeting specific personality types. Therefore, it is crucial that they consider specific patterns of use and differentiate problematic use when creating targeted interventions. This will benefit individuals who are low in conscientiousness and high in neuroticism, who are especially vulnerable to PSMU and SMA, as well as individuals who tend to use social media passively and impulsively. By further conducting research and longitudinal studies to define causal relationships, it will allow at-risk individuals to be identified earlier and personalised interventions to be made to better protect them from detrimental outcomes.

Strengths of this review paper comprise the inclusion of various measures of SMU, including time spent, frequency of visits to social media sites, passive vs. active use, motives, and specific behaviours displayed while using social media, as well as multiple mental health outcomes, including depression, social isolation, problematic use, and addiction. Moreover, the inclusion of longitudinal studies provides stronger evidence for the temporal relationships between the variables. Despite this inclusion, most of the studies are cross-sectional; therefore, it is hard to establish the directionality of the relationship between personality and SMU.

Furthermore, limitations stem from the fact that most of the evidence is correlational, so causal relationships cannot be extracted from the studies. While the use of different measures of SMU and mental health outcomes is effective in covering a larger range of situations, it makes comparisons between studies more difficult. Lastly, samples for most of the studies often come from adolescents, university students, or young adults, restricting the ability to make generalisations representative of a larger population.

Potential gaps in the literature include the lack of longitudinal research examining the

relationship between SMU and mental health outcomes, limited evidence proving causality, and inconsistent definitions and measures of SMU. There is also limited research on certain outcomes, such as social isolation, whereas outcomes such as depression are more heavily studied, creating an imbalance. Therefore, it is important for future research to implement longitudinal designs, create clear distinctions between measures of SMU, and examine more diverse samples using consistent measures.

Overall, the literature presents the mental health consequences of SMU and how they are shaped by different patterns of usage that vary between personality traits. It highlights how certain personality traits are more vulnerable to certain negative mental health outcomes and may require more attention and awareness.

Works Cited

Ahmed, Oli, et al. "Problematic Social Media Use, Personality Traits, and Mental Health Among Bangladeshi University Students." Journal of Technology in Behavioral Science, vol. 7, no. 2, June 2022, pp. 183–91. Springer Link, https://doi.org/10.1007/s41347-021-00235-1.

Allport, G. W. Personality: A Psychological Interpretation. Holt, 1937, pp. xiv, 588. Personality: A Psychological Interpretation.

Andreassen, Cecilie Schou, et al. "The Relationship between Addictive Use of Social Media, Narcissism, and Self-Esteem: Findings from a Large National Survey." Addictive Behaviors, vol. 64, Jan. 2017, pp. 287–93. ScienceDirect, https://doi.org/10.1016/j.addbeh.2016.03.006.

Andrews, Nadia P., et al. "Is Social Media Use Changing Who We Are? Examining the Bidirectional Relationship Between Personality and Social Media Use." Cyberpsychology, Behavior, and Social Networking, vol. 23, no. 11, Nov. 2020, pp. 752–60. SAGE Journals, https://doi.org/10.1089/cyber.2019.0744.

Back, Mitja D., et al. "Personality and Social Relationships: What Do We Know and Where Do We Go." Personality Science, vol. 4, no. 1, Jan. 2023, p. e7505. SAGE Journals, https://doi.org/10.5964/ps.7505.

Berryman, Chloe, et al. "Social Media Use and Mental Health among Young Adults." Psychiatric Quarterly, vol. 89, no. 2, June 2018, pp. 307–14. Springer Link, https://doi.org/10.1007/s11126-017-9535-6.

Chi, Li-Chiu, et al. "The Phubbing Phenomenon: A Cross-Sectional Study on the Relationships among Social Media Addiction, Fear of Missing out, Personality Traits, and Phubbing Behavior." Current Psychology, vol. 41, no. 2, Feb. 2022, pp. 1112–23. Springer Link, https://doi.org/10.1007/s12144-021-02468-y.

Choi, Tae Rang, et al. "Get behind My Selfies: The Big Five Traits and Social Networking Behaviors through Selfies." Personality and Individual Differences, vol. 109, Apr. 2017, pp. 98–101. ScienceDirect, https://doi.org/10.1016/j.paid.2016.12.057.

Dhingra, Manish, and Rakesh K. Mudgal. "Historical Evolution of Social Media: An Overview." SSRN Scholarly Paper no. 3395665, Social Science Research Network, 15 Mar. 2019. papers.ssrn.com, https://doi.org/10.2139/ssrn.3395665.

Escobar-Viera, César G., et al. "Passive and Active Social Media Use and Depressive Symptoms Among United States Adults." Cyberpsychology, Behavior, and Social Networking, vol. 21, no. 7, July 2018, pp. 437–43. SAGE Journals, https://doi.org/10.1089/cyber.2017.0668.

Faverio, Michelle, and Olivia Sidoti. "Teens, Social Media and Technology 2024." Pew Research Center, 12 Dec. 2024, https://www.pewresearch.org/internet/2024/12/12/teens-social-media-and-technology 2024/. Teens & Tech.

Gil de Zúñiga, Homero, et al. "Personality Traits and Social Media Use in 20 Countries: How Personality Relates to Frequency of Social Media Use, Social Media News Use, and

Social Media Use for Social Interaction." Cyberpsychology, Behavior, and Social Networking, vol. 20, no. 9, Sept. 2017, pp. 540–52. SAGE Journals, https://doi.org/10.1089/cyber.2017.0295.

Hawi, Nazir, and Maya Samaha. "Identifying Commonalities and Differences in Personality Characteristics of Internet and Social Media Addiction Profiles: Traits, Self-Esteem, and Self-Construal." Behaviour & Information Technology, vol. 38, no. 2, Feb. 2019, pp. 110–19. Taylor and Francis+NEJM, https://doi.org/10.1080/0144929X.2018.1515984.

Jaradat, Mohammed-Issa, and Anas jh. Do Personality Traits Play a Role in Social Media Addiction? Key Considerations for Successful Optimized Model to Avoid Social Networking Sites Addiction: A Developing Country Perspective. Aug. 2017.

Kapoor, Kawaljeet Kaur, et al. "Advances in Social Media Research: Past, Present and Future." Information Systems Frontiers, vol. 20, no. 3, June 2018, pp. 531–58. Springer Link, https://doi.org/10.1007/s10796-017-9810-y.

Khan, Gohar Feroz, et al. "Social Media Risks and Benefits: A Public Sector Perspective." Social Science Computer Review, vol. 32, no. 5, Oct. 2014, pp. 606–27. SAGE Journals, https://doi.org/10.1177/0894439314524701.

Kietzmann, Jan H., et al. "Social Media? Get Serious! Understanding the Functional Building Blocks of Social Media." Business Horizons, SPECIAL ISSUE: SOCIAL MEDIA, vol. 54, no. 3, May 2011, pp. 241–51. ScienceDirect, https://doi.org/10.1016/j.bushor.2011.01.005.

Kircaburun, Kagan. "Effects of Gender and Personality Differences on Twitter Addiction among Turkish Undergraduates." Journal of Education and Practice, vol. 7, no. 24, 2016, pp. 33–42.

———. "Uses and Gratifications of Problematic Social Media Use Among University Students: A Simultaneous Examination of the Big Five of Personality Traits, Social Media Platforms, and Social Media Use Motives." International Journal of Mental Health and Addiction, vol. 18, no. 3, June 2020, pp. 525–47. Springer Link, https://doi.org/10.1007/s11469-018-9940-6.

Lin, Liu Yi, et al. "Association Between Social Media Use and Depression Among U.s. Young Adults." Depression and Anxiety, vol. 33, no. 4, 2016, pp. 323–31. Wiley Online Library, https://doi.org/10.1002/da.22466.

Marino, Claudia, et al. "Modeling the Contribution of Personality, Social Identity and Social Norms to Problematic Facebook Use in Adolescents." Addictive Behaviors, vol. 63, Dec. 2016, pp. 51–56. ScienceDirect, https://doi.org/10.1016/j.addbeh.2016.07.001.

Merrill, Renae A., et al. "Associations between Social Media Use, Personality Structure, and Development of Depression." Journal of Affective Disorders Reports, vol. 10, Dec. 2022, p. 100385. ScienceDirect, https://doi.org/10.1016/j.jadr.2022.100385.

Perrin, Andrew. "Social Media Usage: 2005-2015." Pew Research Center, 8 Oct. 2015, https://www.pewresearch.org/internet/2015/10/08/social-networking-usage-2005-201 5/.

Social Media.
Primack, Brian A., et al. “Temporal Associations Between Social Media Use and Depression.” American Journal of Preventive Medicine, vol. 60, no. 2, Feb. 2021, pp. 179–88. ScienceDirect, https://doi.org/10.1016/j.amepre.2020.09.014.
Roberts, Brent W. “Back to the Future: Personality and Assessment and Personality Development.” Journal of Research in Personality, Personality and Assessment at Age 40: Reflections on the Past Person–Situation Debate and Emerging Directions of Future Person-Situation Integration, vol. 43, no. 2, Apr. 2009, pp. 137–45. ScienceDirect, https://doi.org/10.1016/j.jrp.2008.12.015.
Shensa, Ariel, et al. “Problematic Social Media Use and Depressive Symptoms among U.S. Young Adults: A Nationally-Representative Study.” Social Science & Medicine, vol. 182, June 2017, pp. 150–57. ScienceDirect, https://doi.org/10.1016/j.socscimed.2017.03.061.
Soto, Christopher, and Joshua Jackson. Five-Factor Model of Personality. 2020. ResearchGate, https://doi.org/10.1093/obo/9780199828340-0120.
Umegaki, Yusuke, and Ayaka Higuchi. “Personality Traits and Mental Health of Social Networking Service Users: A Cross-Sectional Exploratory Study among Japanese
Undergraduates.” Computers in Human Behavior Reports, vol. 6, May 2022, p. 100177. ScienceDirect, https://doi.org/10.1016/j.chbr.2022.100177.
Ventriglio, Antonio, et al. “Social Media Use and Emerging Mental Health Issues.” Industrial Psychiatry Journal, vol. 33, no. Suppl 1, Aug. 2024, pp. S261–64. PubMed Central, https://doi.org/10.4103/ipj.ipj_45_24.
Vogels, Emily A., et al. “Teens, Social Media and Technology 2022.” Pew Research Center, 10 Aug. 2022, https://www.pewresearch.org/internet/2022/08/10/teens-social-media-and-technology 2022/. Social Media.
Whaite, Erin O., et al. “Social Media Use, Personality Characteristics, and Social Isolation among Young Adults in the United States.” Personality and Individual Differences, vol. 124, Apr. 2018, pp. 45–50. ScienceDirect, https://doi.org/10.1016/j.paid.2017.10.030.

Consumer Behaviour and Policy Preferences in response to Rising Fuel Prices: Evidence from an Online Survey in India By Danish Pereira Kalia

Abstract

Rising fuel prices can have a significant effect on household transportation expenditure, travel behaviour and perceptions of cost of other goods and services. This study examined respondents' reported experiences of rising fuel prices in May 2026 in India and explored their preferences for measures to reduce the associated financial burden. An anonymous online questionnaire was distributed using convenience sampling generating 196 responses. Descriptive statistics were used to summarise responses, and chi square tests were used to examine associations between household income, transportation mode and behavioral responses. Nearly three fourths of the respondents reported that rising fuel prices had affected their daily lives. Household income was significantly associated with perceived impact and changes in travel behaviour, with respondents from lower-income households reporting greater effects. A large majority expressed willingness to consider electric vehicles or alternative transportation. Improved public transport was the most frequently selected policy measure. The findings suggest that lower income respondents may experience a greater perceived burden from rising fuel prices and highlight the importance of affordable public transportation and diversified mobility options in reducing this burden.

1. Introduction

1.1 Background

Fuel prices play an important role in household expenditure because transportation is closely linked to employment, education, access to services and the movement of goods. Changes in fuel prices can therefore extend beyond the direct cost of transportation and influence household budgets and the prices of other goods and services. These effects may not be experienced equally across households, particularly when transportation represents a larger share of household expenditure. **(Bhuvandas and Gundimeda; Priya and Sharma)**

India is particularly exposed to changes in international crude oil prices because a large proportion of its crude oil requirements is met through imports **(Petroleum Planning and Analysis Cell)**. Consequently , international oil- price fluctuations, exchange rate movements and domestic taxes and charges can influence retail fuel prices and broader economic conditions **(Canara Bank)**. International crude oil prices have experienced substantial fluctuations in recent decades, exposing oil- importing economies such as India to external price shocks **(U. S. Energy Information Administration, "Spot Prices")**.

In May 2026, escalating conflict in the Middle East and disruptions affecting the Strait of Hormuz contributed to heightened concerns regarding global oil supply and fuel prices. During the study period, retail petrol prices also varied across Indian cities. The prices reported for Delhi, Mumbai, Bhopal, Hyderabad, Bengaluru, Kolkata and Chennai ranged from ₹ 102.72 to

₹118.07 per litre, with differences reflecting state taxes, freight, transportation charges, dealer commissions and other local charges **(Ministry of Petroleum and Natural Gas)**.

While changes in fuel prices are widely recognized as an important economic issue, their effects may differ according to household income, transportation choices and individual travel requirements. Understanding how people perceive the effects of fuel price increases and whether they modify their travel behaviour can provide insight into the household- level consequences of such changes. Examining preferences for measures to reduce the associated financial burden may also provide useful information about transportation policies that respondents consider most beneficial.

Against this background, the present study examined respondents' reported experiences following the rise in fuel prices in May 2026. It assessed the perceived effects on daily life and household expenditure, examined whether responses differed according to household income and primary mode of transportation, explored changes in travel behaviour, investigated willingness to consider electric vehicles or alternative modes of transportation, and examined respondents' preferences for measures to reduce the burden associated with rising fuel prices.

1.2 Research Objectives

This research was conducted to determine

1. Assess the perceived impact of rising fuel prices on respondents' daily lives and transportation expenditure.
2. Examine whether the perceived impact of rising fuel costs differs according to household income.
3. Examine whether respondents' primary mode of transportation was associated with changes in monthly transportation expenditure and travel behaviour.
4. Assess respondents' willingness to consider electric vehicles or alternative modes of transportation in response to rising fuel prices.
5. Examine respondents' preferences for measures and government policies aimed at reducing the financial burden associated with rising fuel prices.

2. Literature Review

2.1 Fuel Prices and Household Welfare

Fuel-price increases can affect households through both direct and indirect channels. The direct effect arises from higher expenditure on petrol and diesel, while indirect effects may occur through changes in transportation costs and the prices of goods and services. Because transportation is an important component of household expenditure, increases in fuel prices can alter the allocation of household resources and affect overall welfare **(Bhuvandas and Gundimeda; Priya and Sharma)**.

Research examining the Indian context has found that the burden of transport-fuel price increases is not distributed equally across households. Bhuvandas and Gundimeda examined the

welfare effects of transport-fuel price changes in India using household expenditure data and found that lower- and middle-income households were relatively more vulnerable to such increases **(Bhuvandas and Gundimeda).** Their findings also highlighted the importance of affordable and efficient public transportation as a means of reducing the welfare burden associated with fuel-price increases.

More recent research has examined the transmission of energy-price shocks to household consumption in India. Priya and Sharma found that energy-price shocks adversely affected household consumption and that the effects varied across categories of expenditure.[8] These findings suggest that the consequences of energy-price increases can extend beyond direct expenditure on fuel and influence broader household spending decisions **(Priya and Sharma).**

2.2 Income and the Distributional Effects of Energy-Price Shocks

Household income is an important consideration when examining the effects of changes in energy and transportation costs. Lower-income households may have fewer financial resources available to absorb increases in essential expenditure and may also have fewer opportunities to substitute away from private transportation.

Evidence from India supports the existence of distributional differences in the effects of energy-price shocks. Bhuvandas and Gundimeda found that lower- and middle-income households experienced greater relative welfare effects from transport-fuel price increases. Similarly, research by Bhattarai, Chatterjee and Udupa examined the distributional effects of global food and oil-price shocks on Indian households and reported heterogeneous effects across income groups. These findings provide a basis for examining whether respondents' perceptions of the effects of a recent fuel-price increase also vary according to household income **(Bhuvandas and Gundimeda; Bhattarai et al.).**

2.3 Fuel-Price Elasticity and Transportation Demand

The relationship between fuel prices and transportation demand is also important from an economic perspective. Price elasticity describes the extent to which demand changes in response to a change in price. Transportation fuel may be relatively difficult to substitute in the short term because individuals may depend on private vehicles for essential travel **(Bhuvandas and Gundimeda)**. Bhuvandas and Gundimeda's analysis provides evidence on this relationship in India, estimating transport-fuel demand and welfare effects across rural and urban households and income groups. Their results demonstrate the importance of considering both fuel-demand responses and the welfare consequences of price changes **(Bhuvandas and Gundimeda)**.

Yadav, Kanjilal, Dutta and Ghosh examined gasoline, high-speed diesel and electric-vehicle demand in India's road-transport sector using data from November 2014 to April 2022. They found that gasoline and high-speed diesel demand were price inelastic in their analysis, suggesting that increases in fuel prices or taxation may not produce proportionate reductions in conventional fuel demand. Their study also found that policies affecting the relative cost of electricity and conventional fuels could influence EV demand **(Yadav et al.)**.

This literature provides an economic basis for examining whether respondents in the present study reported changes in transportation expenditure or travel behaviour following a rise in fuel prices.

2.4 Fuel Prices and Travel Behaviour

Fuel prices can also influence transportation decisions. When the cost of private transportation increases, individuals may reduce discretionary journeys, combine trips, change transportation modes or consider alternative forms of mobility. However, the extent to which individuals can change their behaviour may depend on the availability, affordability and accessibility of alternative transportation **(Bhuvandas and Gundimeda; Yadav et al.)**.

Recent research in India provides evidence of behavioural responses to fuel-price shocks. **Giri and Shukla** examined household mobility responses to a fuel-price shock in a Tier-II Indian city and reported 25% reductions in vehicle-based trips, increased trip chaining and reductions in discretionary travel. Their findings also indicated that structural constraints could limit households' ability to shift to alternative modes of transportation.

These findings are relevant to the present study because the survey examines both reported changes in travel behaviour and the relationship between transportation choices and the perceived financial effects of rising fuel prices.

The existing evidence therefore suggests that fuel-price increases can produce behavioural adaptations, but that such adaptations are not necessarily uniform. The availability of public transportation and other alternatives may determine whether households can meaningfully reduce their dependence on conventional fuel.

2.5 Public Transportation as a Policy Response

The availability of affordable public transportation is an important consideration when evaluating policy responses to fuel-price increases. **Bhuvandas and Gundimeda** argued that expansion of sustainable and efficient public transportation could help reduce household vulnerability to transport-fuel price increases, particularly where households lack affordable substitutes for petrol and diesel.

The evidence from Gorakhpur similarly suggests that simply increasing fuel costs does not automatically result in substantial shifts toward public transportation. **Giri and Shukla** found that public-transport use changed relatively little despite the fuel-price shock, highlighting the role of structural and accessibility constraints.

This has implications for policy. If households cannot easily substitute away from private vehicles, policies aimed solely at increasing the cost of conventional fuels may place additional financial pressure on households without producing an equivalent reduction in fuel demand. Investment in reliable, affordable and accessible public transportation may therefore be an important complementary policy.

2.6 Electric Vehicles and Alternative Transportation

Electric vehicles represent another potential alternative to conventional petrol- and diesel-powered transportation. However, willingness to adopt an EV does not necessarily translate into actual adoption, as consumers may face financial, infrastructural and informational barriers **(Munshi et al; Yadav and Yadav).**

Munshi et al. examined potential EV buyers in Hyderabad using stated-preference data and discrete-choice models. They found general willingness to purchase EVs among the middle-income working population but also identified important barriers related to awareness, running costs and vehicle specifications. Their findings suggested that financial and policy incentives could positively influence EV choice.

Yadav and Yadav examined barriers and motivators associated with EV adoption among Indian consumers using a mixed-method approach. Their study identified subsidies, environmental concerns and personal and social values as factors supporting EV adoption, while lack of infrastructure, range anxiety, trust and high initial costs were identified as barriers.

These findings are relevant to the present study because the survey examines respondents' willingness to consider electric vehicles or alternative transportation. Importantly, the present study measures stated willingness rather than actual EV adoption and therefore should not interpret such responses as evidence that respondents will necessarily purchase an EV.

2.7 Fuel Pricing, Taxation and Government Policy in India

Fuel prices in India are influenced by international crude-oil prices as well as domestic taxation, pricing mechanisms and government policy. Consequently, the policy response to fuel-price increases involves a balance between protecting consumers, maintaining government revenues and ensuring the financial viability of petroleum-sector companies **(Jain)**.

Jain examined India's petroleum-subsidy reforms and highlighted the interaction between three major policy instruments: retail prices, petroleum taxation and subsidies The study showed that petroleum-price reform involves more than simply increasing or decreasing retail fuel prices and that policy decisions must balance fiscal, commercial and consumer interests.

The policy literature therefore suggests that there is no single intervention capable of addressing all consequences of fuel-price increases. Consumer protection, public transportation, taxation, subsidies and alternative-energy policies can interact in complex ways **(Jain).**

This provides a useful context for examining the measures that respondents themselves consider most effective for reducing the burden associated with rising fuel prices.

2.8 Research Gap

Existing research has established that fuel- and energy-price increases can affect household welfare, consumption and transportation behaviour, with potentially greater effects among lower-income households **(Bhuvandas and Gundimeda; Priya and Sharma; Bhattarai et al.; Yadav et al.**). Research has also examined behavioural responses to fuel-price shocks and the barriers and incentives associated with alternative transportation and electric-vehicle adoption **(Giri and Shukla; Munshi et al.; Yadav and Yadav)**. In addition, studies of India's

petroleum policy demonstrate the complex interaction between retail prices, taxation and subsidies **(Jain)**.

Much of the existing Indian literature has examined these relationships using household expenditure datasets, econometric models, stated-preference studies or focused investigations of particular cities and policy reforms. The present study complements this literature by examining respondents' self-reported experiences following a recent fuel-price increase in India, with particular attention to perceived effects on daily life and transportation expenditure, changes in travel behaviour, willingness to consider electric vehicles or alternative transportation, and preferences for policy measures. By examining these dimensions together, the study provides a descriptive and exploratory perspective on how respondents perceive and respond to a contemporary fuel-price increase**.**

3. Methodology

3.1 Study Design

This study used a cross-sectional, questionnaire-based survey design to examine respondents' reported experiences and perceptions following the rise in fuel prices in India in May 2026. The survey focused on the perceived effects of rising fuel prices on daily life and transportation expenditure, changes in travel behaviour, willingness to consider alternative modes of transportation, and preferences for measures to reduce the associated financial burden.

3.2 Study Population and Sampling

An anonymous online questionnaire was developed using Google Forms based on the research objectives of the study. The questionnaire was distributed electronically through WhatsApp groups that included family members, friends and other individuals residing in India. Convenience sampling was used, with participants recruited through the researcher's existing social networks.

The survey was open for 20 days. A total of 196 responses were received and included in the analysis. Because convenience sampling was used, the sample was not intended to be nationally representative of the Indian population.

3.3 Questionnaire Design

The questionnaire collected information on respondents' demographic and household characteristics and their experiences following the increase in fuel prices.

The questionnaire included questions relating to:

- age of the respondent;
- annual household income;
- primary mode of transportation;
- the area of daily life perceived to be most affected by rising fuel prices;
- the reported increase in monthly transportation expenditure;

- changes in travel behaviour following the increase in fuel prices;
- perceived effects of rising fuel prices on the prices of other goods and services;
- willingness to consider electric vehicles or alternative modes of transportation; and
- respondents' preferred measures for reducing the burden associated with rising fuel prices.

  Some questions permitted respondents to select more than one response, allowing the study to capture multiple perceived effects and preferred policy measures.

3.4 Data Collection and Confidentiality

Before completing the questionnaire, respondents were informed that the survey formed part of a school research project examining the recent increase in fuel prices in India. Responses were collected anonymously, and no personally identifying information was collected as part of the survey.

The Google Form was configured to allow one response per participant in order to reduce the possibility of duplicate submissions. Following completion of the data-collection period, the responses were compiled and prepared for statistical analysis.

3.5 Statistical Analysis

The collected data were initially entered and organised using Microsoft Excel and subsequently analysed using IBM SPSS Statistics, Version 30.0. Descriptive statistics were used to summarise the characteristics of the respondents and their responses. Results were presented primarily as frequencies and percentages.

For questions permitting multiple responses, multiple-response analysis was performed to determine the frequency and percentage of respondents selecting each available option.

Associations between categorical variables were examined using the Pearson chi-square (χ^2) test. Where the assumptions for the Pearson's chi-square test were not satisfied because of low expected cell counts, the Fisher–Freeman–Halton exact test with Monte Carlo estimation was used. A p-value < 0.05 was considered statistically significant.

4. Results

Annual Family Income	**Frequency (n)**	**Percentage (%)**
Less than ₹5 lakh	10	5.1
₹5–10 lakh	26	13.3
₹10–25 lakh	81	41.3
More than ₹25 lakh	79	40.3
Total	**196**	**100.0**

Table 1. Distribution of Respondents According to Annual Family Income (N = 196)

A total of 196 responses were included in the analysis, of which 195 respondents (99.5%) were 18 years and above. Only 1 respondent (0.5%) was below 18 years. The sample was predominantly composed of respondents from higher-income households, with 41.3% reporting an annual family income of ₹10-25 lakh and 40.3% reporting more than ₹25 lakh. Respondents reporting less than ₹5 lakh and ₹5-10 lakh accounted for 5.1% and 13.3%, respectively. (Table 1)

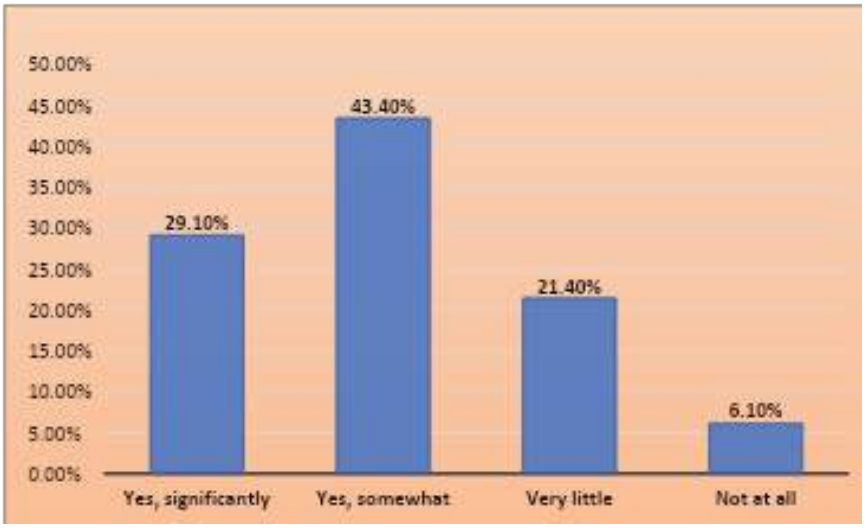

Figure 1. Perceived Impact of Rising Fuel Prices on Daily Life (N = 196)

Figure 1 shows that 72.5% of respondents reported that the rise in fuel prices had affected their daily lives, with 29.10% reporting a significant impact and 43.4% reporting that they were affected somewhat. A further 21.4% reported very little impact, while 6.1% reported no impact.

Annual Family Income	Yes, significantly n (%)	Yes, somewhat n (%)	Very little n (%)	Not at all n (%)	Total	Pearson χ^2 (df)	Fisher–Freeman–Halton Exact Test (Monte Carlo)	Exact p-value
Less than ₹5 lakh	7 (70.0)	3 (30.0)	0 (0.0)	0 (0.0)	10	**20.613 (9)**	**19.187**	**0.014***
₹5–10 lakh	11 (42.3)	9 (34.6)	4 (15.4)	2 (7.7)	26			
₹10–25 lakh	27 (33.3)	35 (43.2)	15 (18.5)	4 (4.9)	81			
More than ₹25 lakh	12 (15.2)	38 (48.1)	23 (29.1)	6 (7.6)	79			
Total	**57 (29.1)**	**85 (43.4)**	**42 (21.4)**	**12 (6.1)**	**196**			

*Statistically significant, Pearson chi-square assumptions were violated; therefore, the Fisher–Freeman–Halton Exact Test with Monte Carlo estimation was used. Exact p = 0.014; 43.8% of cells had expected counts <5 (minimum = 0.61).

Table 2. Association between Annual Family Income and the Perceived Impact of Rising Fuel Prices (N = 196)

A statistically significant association was observed between annual family income and the perceived impact of rising fuel prices (Table 2). Since the assumptions for the Pearson chi-square test were not satisfied, with 43.8% of cells having expected counts below 5 and minimum expected count of 0.61, the Fisher–Freeman–Halton Exact Test with Monte Carlo estimation was used for inference (p=0.014; 99% Cl: 0.011-0.017). Within this sample, respondents from lower-income families were more likely to report a significant impact: 70.0% of those earning less that ₹5 lakh reported a significant impact , compared with 15.2% of those reporting an annual income of more than ₹25 lakh. Conversely, higher income respondents more frequently reported that the impact was somewhat or very little.

Mode of Transport	**Frequency (n)**	**Percentage (%)**
Walking	3	1.5
Bicycle	1	0.5
Auto/Cab	6	3.1
Two-Wheeler	30	15.3
Four-Wheeler	153	78.1
Public Transport	3	1.5
Total	**196**	**100.0**

Table 3. Primary Mode of Transportation Used by Respondents (N = 196)

Table 3 depicts four- wheelers were the predominant primary mode of transportation, reported by 153 respondents (78.1%), followed by two- wheelers, reported by 30 respondents (15.3%). The remaining respondents primarily used auto/cab services (3.1%), walking (1.5%), public transportation (1.5%), or bicycles (0.5%). Overall, 93.4% of respondents primarily used either a four- wheeler or two- wheeler.

Primary Mode of Transport	**< ₹1000 n (%)**	**₹1000–₹5000 n (%)**	**> ₹5000 n (%)**	**Not sure n (%)**	**Total**	**Pearson χ^2 (df)**	**Fisher–Freeman–Halton Exact statistic**	**Exact p-value**
Walking	2 (66.7)	1 (33.3)	0 (0.0)	0 (0.0)	3	**25.543 (15)**	**25.686**	**0.012***
Bicycle	0 (0.0)	1 (100.0)	0 (0.0)	0 (0.0)	1			
Auto/Cab	1 (16.7)	3 (50.0)	1 (16.7)	1 (16.7)	6			
Two-Wheeler	11 (36.7)	14 (46.7)	3 (10.0)	2 (6.7)	30			
Four-Wheeler	24 (15.7)	92 (60.1)	33 (21.6)	4 (2.6)	153			
Public Transport	1 (33.3)	0 (0.0)	1 (33.3)	1 (33.3)	3			
Total	**39 (19.9)**	**111 (56.6)**	**38 (19.4)**	**8 (4.1)**	**196**			

***Statistically significant,** Fisher–Freeman–Halton Exact Test with Monte Carlo estimation (10,000 samples); exact p = 0.012, 99% CI: 0.009–0.015. Pearson assumptions were violated: 17/24 cells (70.8%) had expected counts <5; minimum expected count = 0.04.

Table 4. Association between Primary Mode of Transportation and Increase in Monthly Transportation Expenditure (N = 196)

More than half of the respondents (56.6%) reported an increase of ₹1,000–₹5,000 in their monthly transportation expenditure following the rise in fuel prices, while 19.9% reported an increase of less than ₹1000, 19.4% reported an increase of more than ₹5000 and 4.1% were unsure of the amount. A statistically significant association was observed between the primary mode of transportation and the reported increase in monthly transportation expenditure (Pearson $\chi^2 = 25.543$, df = 15). Since the Pearson chi-square assumptions were not satisfied, with 17 of the 24 cells (70.8%) having expected counts below 5 and a minimum expected count of 0.04, the Fisher–Freeman–Halton Exact Test with Monte Carlo estimation was used for inference (p = 0.012; 99% CI: 0.009–0.015). Among respondents using four-wheelers, the majority (60.1%) reported an increase in monthly transport expenditure of ₹1,000–₹5,000, while 21.6% reported an increase of more than ₹5,000. Similarly, most two-wheeler users (46.7%) also experienced an increase of ₹1,000–₹5,000. (Table 4)

Primary Mode of Transport	No, n (%)	Yes, n (%)	Total	Pearson χ^2 (df)	Fisher–Freeman–Halton Exact Statistic	Exact p-value
Walking	2 (66.7)	1 (33.3)	3	**2.430 (5)**	**2.666**	**0.843 (NS)**
Bicycle	0 (0.0)	1 (100.0)	1			
Auto/Cab	3 (50.0)	3 (50.0)	6			
Two-Wheeler	17 (56.7)	13 (43.3)	30			
Four-Wheeler	90 (58.8)	63 (41.2)	153			
Public Transport	1 (33.3)	2 (66.7)	3			
Total	**113 (57.7)**	**83 (42.3)**	**196**			

Fisher–Freeman–Halton exact test was used because of low expected cell counts. **NS = Not statistically significant.**

Table 5. Association between Primary Mode of Transportation and Change in Travel Habits Due to Rising Fuel Prices (N = 196)

When asked if they had implemented any changes in their travel habits due to rising fuel prices nearly two-fifths (42.3%) of respondents reported changing their travel habits, whereas 57.7% reported no change. No statistically significant association was found between the primary mode of transportation and changes in travel habits due to rising fuel prices (Fisher–Freeman–Halton Exact Test = 2.666, exact p = 0.843). Among four-wheeler users, 41.2% (63/153) reported changing their travel habits, compared with 43.3% (13/30) of two-wheeler users. The proportions varied across the smaller transport categories; however, there was no statistically significant evidence of an association in this sample. (Table 5)

Change in Travel Habit	Frequency (n)	Percent of Respondents (%)*
Reducing unnecessary travel	86	72.9
Walking/Cycling more	33	28.0
Switching to fuel-efficient vehicles	25	21.2
Carpooling	21	17.8
Using public transport more often	17	14.4

*Multiple- response question; respondents could select more than one option. Percentages are calculated using 118 respondents who reported at least one **change** and therefore do not total 100%

Table 6. Changes in Travel Habits Adopted Due to Rising Fuel Prices: Multiple Response Analysis (N = 118)

Among the 118 respondents who reported at least one travel habit change, the most commonly reported change was reducing unnecessary travel, reported by 86 of the 118 of respondents (72.9%). This was followed by walking or cycling more (33 respondents, 28.0%), switching to fuel-efficient vehicles (25 respondents, 21.2%), carpooling (21 respondents, 17.8%), and using public transport more often (17 respondents, 14.4%). As respondents could select more than one option, the percentages do not total 100%. These findings suggest that respondents primarily adopted cost-saving measures, with reducing unnecessary travel being the most commonly reported behavioural response to rising fuel prices. (Table 6)

Annual Family Income	**No n (%)**	**Yes n (%)**	**Total**	**Chi-square (χ^2)**	**df**	**p-value**
Less than ₹5 lakh	4 (40.0)	6 (60.0)	10	**7.858**	3	**0.049***
₹5–10 lakh	11 (42.3)	15 (57.7)	26			
₹10–25 lakh	44 (54.3)	37 (45.7)	81			
More than ₹25 lakh	54 (68.4)	25 (31.6)	79			
Total	**113 (57.7)**	**83 (42.3)**	**196**			

***p < 0.05; Statistically significant**

Table 7. Association Between Annual Family Income and Change in Travel Habits Due to Rising Fuel Prices (N = 196)

A statistically significant association was observed between annual family income and changes in travel habits due to rising fuel prices ($\chi^2 = 7.858$, df = 3, p = 0.049). The proportion of respondents reporting changes in their travel habits was higher among lower-income groups. Specifically, 60.0% of respondents with an annual family income of less than ₹5 lakh and 57.7% of those earning ₹5–10 lakh reported changing their travel habits, compared with 45.7% of those earning more than ₹25 lakh. These findings indicate an association between household income and reported changes in travel behaviour, with a greater proportion of respondents in the lower-income groups reporting changes in their travel habits. (Table 7)

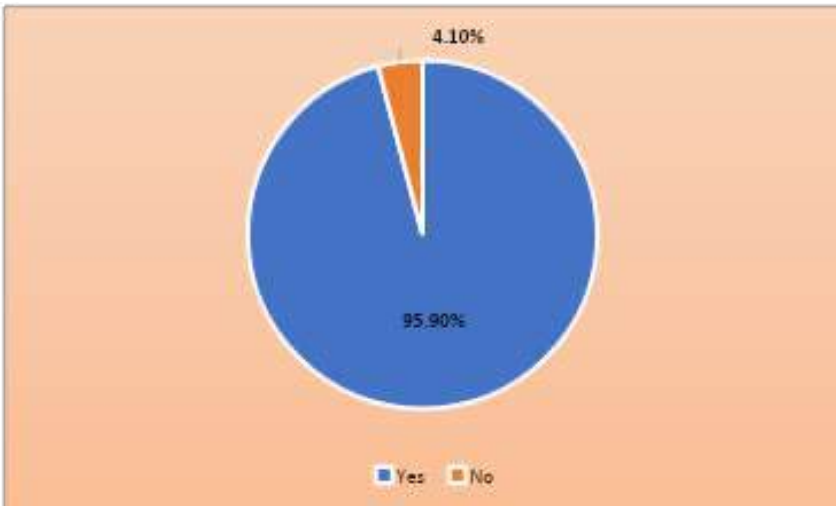

Figure 2. Perceived Effect of Rising Fuel Prices on the prices of Other Goods and Services (N=196)

An overwhelming majority of respondents (95.9%) perceived that rising fuel prices had increased the prices of other goods and services, while 4.1% did not report such an effect (Figure 2).

Response	Frequency (n)	Percentage (%)
Yes	148	75.5
No	48	24.5
Total	**196**	**100.0**

Table 8. Willingness to Switch to Electric Vehicles or Alternative Modes of Transport (N = 196)

Three-fourths of respondents (75.5%) expressed willingness to switch to electric vehicles or alternative modes of transport in the future**,** while 24.5% were not willing to do so (Table 8)

Annual Family Income	No n (%)	Yes n (%)	Total	Chi-square (χ^2)	df	p-value
Less than ₹5 lakh	3 (30.0)	7 (70.0)	10	**2.864**	3	**0.413 (NS)**
₹5–10 lakh	3 (11.5)	23 (88.5)	26			
₹10–25 lakh	22 (27.2)	59 (72.8)	81			
More than ₹25 lakh	20 (25.3)	59 (74.7)	79			
Total	**48 (24.5)**	**148 (75.5)**	**196**			

Table 9. Association between Annual Family Income and Willingness to Switch to Electric Vehicles or Alternative Transport (N = 196)

No statistically significant association was found between annual family income and willingness to switch to electric vehicles or alternative modes of transport ($\chi^2 = 2.864$, df = 3, p =

0.413). Willingness to adopt electric vehicles or alternative transport was reported across all income groups, ranging from 70% among respondents with an annual family income of less than ₹5 lakh to 88.5% among those earning ₹5-10 lakh. Among respondents earning ₹10-25 lakhs and more than ₹25 lakh, 72.8% and 74.7% respectively expressed willingness to switch. Overall, the findings indicate that willingness to consider electric vehicles or alternative modes of transport did not differ significantly according to annual family income**.** (Table 9)

Change in Travel Habits	No n (%)	Yes n (%)	Total	Chi-square (χ^2)	df	p-value
No	30 (26.5)	83 (73.5)	113	**0.612**	1	**0.434 (NS)**
Yes	18 (21.7)	65 (78.3)	83			
Total	**48 (24.5)**	**148 (75.5)**	**196**			

Table 10. Association between Change in Travel Habits and Willingness to Switch to Electric Vehicles or Alternative Transport (N = 196)

No statistically significant association was found between changes in travel habits due to rising fuel prices and willingness to switch to electric vehicles or alternative modes of transport ($\chi^2 = 0.612$, df = 1, p = 0.434). Among respondents who did not change their travel habits, 73.5% expressed willingness to switch to electric vehicles or alternative transport, compared with 78.3% among those who did change their travel habits. Although willingness to consider alternative transport was slightly higher among respondents who had changed their travel habits, the difference was not statistically significant (Table 10).

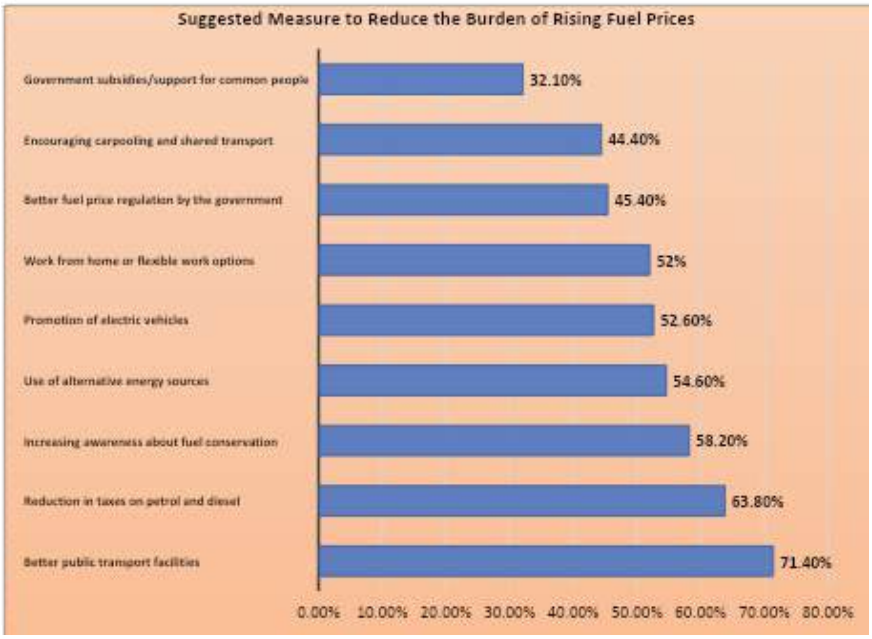

***Multiple responses were permitted; therefore percentage do not add to 100%**

Figure 3. Suggested Measures to Reduce the Burden of Rising Fuel Prices Among Respondents (N=196)

Among the 196 respondents, the most frequently selected measure for reducing the burden of rising fuel prices was improving public transport facilities (71.4%), followed by reducing taxes on petrol and diesel (63.8%) and increasing awareness about fuel conservation (58.2%). More than half of the respondents also selected the use of alternative energy sources (54.6%), promotion of electric vehicles (52.6%), and work-from-home or flexible work options (52.0%). Other suggested measures included better government regulation of fuel prices (45.4%), encouraging carpooling and shared transport (44.4%), and providing government subsidies or support for the general public (32.1%). Overall, respondents identified a combination of infrastructure, fiscal, energy transition, and behavioural measures as potential ways of reducing the burden associated with rising fuel prices (Figure 3).

Variable	Key Finding	Result
Age	Respondents above 18 years	99.5%
Annual family income	₹10–25 lakh	41.3%
Primary mode of transport	Four-wheeler	78.1%
Impact of rising fuel prices	Yes, somewhat	43.4%
Most affected area*	Daily travel expenses	64.1%
Increase in monthly transport spending	₹1,000–₹5,000	56.6%
Changed travel habits	No	57.7%
Most common travel habit change*	Reducing unnecessary travel	72.9%
Believe fuel prices increased prices of other goods	Yes	95.9%
Willing to switch to electric vehicles/alternative transport	Yes	75.5%
Most preferred measure to reduce fuel burden*	Better public transport facilities	71.4%

*Multiple response questions; therefore, percentages do not total 100%.

Table 11: Summary of Key Findings of the Study (N = 196)

The summary of the key findings demonstrates the broad economic and behavioural impact of rising fuel prices among 196 respondents (Table 11). Four wheelers were the predominant mode of transport (78.1%), while 43.4% of respondents reported that rising fuel prices had affected them “somewhat”. Daily travel expenses were the most frequently reported area of impact (64.1%), and 56.6.% reported an increase of ₹1000- ₹5000 in monthly transport expenditure. Although 57.7% reported no change in their travel habits, among those who made changes, reducing unnecessary travel was the most common response (72.9%). A substantial majority (95.9%) perceived that rising fuel prices had increased the prices of other goods and services. Furthermore, 75.5% expressed willingness to switch to electric vehicles or alternative modes of transport, while improved public transport facilities emerged as the most preferred measure for reducing the burden of rising fuel prices (71.4%). Overall, the findings indicate that rising fuel prices have affected household expenditure and have generated considerable interest in more sustainable and accessible transport alternatives.

Table	Variables Compared	Pearson χ^2 (df)	Fisher–Freeman–Halton Exact Statistic	Exact/Monte Carlo p-value	Result
2	Annual family income × Impact of fuel prices on daily life	20.613 (9)	19.187	**0.014***	Significant
4	Primary mode of transport × Increase in monthly transport spending	25.543 (15)	25.686	**0.012***	Significant
5	Annual family income × Change in travel habits	7.858 (3)	—	**0.049***	Significant
9	Annual family income × Willingness to switch to EV/alternative transport	2.864 (3)	—	0.413	Not significant
10=	Change in travel habits × Willingness to switch to EV/alternative transport	0.612 (1)	—	0.434	Not significant

$p < 0.05$. Fisher–Freeman–Halton exact test with Monte Carlo estimation was used for Tables 12–14 because of low expected cell counts. Pearson χ^2 values are shown for reference. Tables 15–17 met the assumptions for Pearson's χ^2 test.

Table 12: Summary of Chi-square Analysis

The summary of the Chi- square analysis (Table 12) gives us inferential statistics which indicated a statistically significant association between annual family income and perceived impact of rising fuel prices (Fisher- Freeman- Halton exact test, Monte carlo p = 0.014). Similarly, primary mode of transportation was significantly associated with increased monthly transportation expenditure (Monte Carlo p= 0.012), while annual family incomes was significantly associated with changes in travel habits (Pearson's $X^2 = 7.858$, df = 3, p= 0.049).

Respondents also provided several additional recommendations through the open-ended question. These included promoting research on hydrogen fuel and flex-fuel technologies as sustainable alternatives to conventional fuels. Some respondents felt that flex-fuel vehicles could help address the driving range concerns associated with electric vehicles. Other suggestions included improving public transport systems and last-mile connectivity, encouraging the use of bicycles for short-distance travel, linking domestic fuel prices more closely to global crude oil prices, consolidating service providers to reduce unnecessary travel, expanding electric vehicle charging infrastructure, and improving governance, tax collection, and the effective utilization of public funds to reduce the financial burden of rising fuel prices on citizens.

Overall, the findings indicate that rising fuel prices have significantly affected the daily lives of respondents, particularly through increased transportation expenses and effects on household budgets. While 42.3% of respondents reported changing their travel habits, 75.5% expressed willingness to switch to electric vehicles or alternative modes of transport. Improving public transport, reducing taxes on petrol and diesel, and promoting alternative energy sources were amongst the most commonly suggested measures to reduce the burden of rising fuel prices. Overall, the findings highlight the need for effective policy measures and sustainable transportation solutions to minimize the impact of increasing fuel costs.

5. Discussion

This study examined the perceived effects of rising fuel prices on households in India, with particular emphasis on household income, transportation expenditure, travel behaviour, willingness to adopt alternative transportation, and preferences for measures to reduce the burden of rising fuel prices.

The findings indicate that rising fuel prices were perceived to have a substantial effect on respondents' daily lives, with the magnitude of the reported impact varying significantly across income groups. The largest proportion of respondents belonged to households with an annual income of ₹10–25 lakh (41.3%), followed by those with an annual income of more than ₹25 lakh (40.3%). The sample therefore consisted predominantly of higher-income households. The study also had a substantially greater proportion of four-wheeler users than may be expected in the broader Indian population. These characteristics may partly reflect the convenience sampling approach, as respondents were recruited through the researcher's family and social networks.

One of the most important findings was the differential perceived impact of rising fuel prices across income groups. Lower-income respondents were considerably more likely to report that rising fuel prices had significantly affected their daily lives and were also more likely to report changes in their travel behaviour. This pattern suggests that lower-income households may have less financial capacity to absorb increases in transportation costs and may therefore be more likely to modify their behaviour in response to rising fuel prices. However, because this was a cross-sectional survey using convenience sampling, the findings should be interpreted as an observed association rather than evidence of a causal relationship.

Income was significantly associated with changes in travel behaviour, but not with willingness to adopt electric vehicles or alternative transportation. This distinction is important. It suggests that economic differences may influence immediate behavioural responses to fuel price increases, whereas willingness to consider alternative transportation may be relatively widespread across income groups.

Although 75.5% of respondents expressed willingness to consider switching to electric vehicles or alternative modes of transport, this finding should not be interpreted as evidence that they will necessarily adopt such options. Stated willingness and actual behavioural adoption may differ because of financial constraints, vehicle costs, charging infrastructure, availability of public transportation, travel requirements, and other practical barriers. The contrast between the relatively high stated willingness to consider alternative transportation and the fact that 57.7% of respondents reported no change in their travel habits highlights the difference between intention and actual behaviour.

The strong preference for improved public transportation, selected by 71.4% of respondents, is particularly noteworthy given that only a small proportion of respondents reported public transportation as their primary mode of transport. This may indicate that respondents perceive limitations in existing public transportation systems as a barrier to reducing dependence on private vehicles. Improving the availability, reliability, affordability, and connectivity of public transportation could therefore provide households with a practical alternative to private vehicle use.

Several respondents suggested flex-fuel technology as an alternative to conventional petrol and diesel. This reflects an interest in diversifying transportation fuels; however, wider adoption of flex-fuel vehicles would depend on vehicle compatibility, fuel availability, supporting infrastructure, and consumer acceptance. Similarly, suggestions relating to hydrogen fuel and electric vehicles indicate an interest in longer-term alternatives to conventional fossil fuels.

The suggestion regarding improvement of last-mile connectivity is also relevant. Better connectivity between homes, workplaces, and major public transportation hubs could make public transportation more practical and potentially reduce dependence on private vehicles. Such measures could be particularly important where the lack of convenient connections discourages people from using public transportation.

The study has several limitations. First, convenience sampling was used, with respondents recruited through family, friends, and WhatsApp groups; therefore, the sample is not nationally representative. Second, the sample contained a relatively high proportion of higher-income households and four-wheeler users, which may limit the generalizability of the findings to lower-income households and populations with different transportation patterns. Third, the study relied on self-reported perceptions and behaviour rather than objectively recorded household expenditure or fuel consumption. Finally, the cross-sectional nature of the study means that associations observed between variables should not be interpreted as causal relationships.

Overall, the findings suggest that rising fuel prices do not affect households uniformly. Lower-income respondents were more likely to report significant effects on their daily lives and changes in travel behaviour, suggesting that fuel-price increases may place a relatively greater burden on households with fewer financial resources. At the same time, the high willingness to consider alternative transportation across income groups suggests an openness to behavioural change if suitable, accessible, and affordable alternatives are available. The strong preference for improved public transportation therefore has important policy implications, particularly if India seeks to reduce dependence on private vehicles while limiting the transportation-related financial burden on households.

6. Conclusion

The present study found that rising fuel prices were perceived to have a substantial impact on the daily lives and household finances of respondents, particularly through increased transportation expenditure. Lower- income respondents were more likely to report a significant impact on their daily lives and changes in travel behaviour, indicating that the burden of rising fuel prices may differ across income groups.

The findings also showed that 75.5% of the respondents were willing to consider switching to electric vehicles or alternative modes of transport. However, this willingness was not significantly associated with annual family income or with whether respondents had previously changed their travel habits. This suggests that willingness to consider alternative

transportation was relatively widespread, although stated willingness may not necessarily translate into immediate behavioural change.

Respondents favoured a combination of policy measures and infrastructure improvements to reduce the burden of rising fuel prices. Improved public transportation facilities were the most frequently suggested measure, followed by reductions in taxes on petrol and diesel, greater awareness of fuel conservation, promotion of alternative energy sources, electric vehicles, and flexible work arrangements.

Overall, the findings indicate that the effects of rising fuel prices extend beyond the direct cost of fuel and may influence household expenditure, travel behaviour, and perceptions of alternative transportation. Addressing these effects may require a combination of accessible and affordable public transportation, appropriate fiscal measures, improved alternative-transport infrastructure, and greater adoption of sustainable energy and mobility options. Further research using larger and more representative samples, together with objective measures of household expenditure and transportation use, would help provide a more comprehensive understanding of the impact of fuel- price increases on Indian households.

Works Cited

Bhuvandas, Dhanyashree and Haripriya Gundimeda. "Welfare impacts of transport fuel price changes on Indian households: An application of LA-AIDS model." Energy Policy, vol. 144, 2020, article 111583. Accessed on 18th July 2026.https://doi.org/10.1016/j.enol.2020.111583

Bhattarai, Saroj, Arpita Chatterjee, and Gautham Udupa. "Food, fuel, and facts: Distributional effects of global price shocks." Working paper, no. 6888, Indian Institute of Management Bangalore, 2023. Accessed on 20th July 2026 https://www.iimb.ac.in/node/12200

Canara Bank. "Oil Price Sensitivity of the Indian Economy." Canara Bank, 26 June 2025. Accessed on July 17 2026. https://www:canarabank.bank.in/documents/d/guest/20-oil-price-sensitivity-of-indian-economy.

Giri Shantanu, and Bharti Shukla. "Adaptive urban mobility under fuel price shocks: Empirical evidence from a Tier- II Indian city." Case studies on transport policy, vol. 24, 2026, article 101805. Accessed on 22th July 2026.https://doi.org/10.1016/j.cstp.2026.101805

Jain, Anil K. "A fine balance: Lessons from India's experience with petroleum subsidy reforms." Energy Policy, vol. 119, 2018, pp. 242-249. https://doi/10.1016/j.enpol.2018.04.050

Munshi, Talat, Subhash Dhar, and Jyoti Painuly. "Understanding barriers to electric vehicle adoption for personal mobility: A case study of middle-income in-service residents in Hyderabad city, India." Energy Policy, vol. 167, 2022, article 112956. Accessed on 22th July 2026. https://doi.org/10.1016/j.enpol.2022.112956

Petroleum Planning & Analysis Cell. "Import/Export of Crude Oil and Petroleum Products." Ministry of Petroleum and Natural Gas, Government of India, 2026. Accessed on July 16 2026 https://ppac.gov.in/import-export.

Petroleum Planning & Analysis Cell. "Price Build-up and State RSP of Petrol and Diesal." Ministry of Petroleum and Natural Gas, Government of India, 2026. Accessed on 14th July 2026. https://ppac.gov.in/retail-selling-price-rsp-of-petrol-diesal-and-domestic-lpg/price-build-up-of-petrol-and-diesel.

Priya, Pragati, and Chandan Sharma. "On transmission channels of energy prices and monetary policy shocks to household consumption: Evidence from India." Energy economics, vol. 136, 2024, article 107723. Accessed on 20th July 2026. https://doi.org/10.1016/j.eneco.2024.107723

U.S. Energy Information Administration. "Refining Crude oil: The Refining Process." U.S. Energy Information Administration, 2024. Accessed on July 17 2026 https://www.eia.gov/energyexplained/oil-and-petroleum-products/refining-crude-oil the-refining-process.php

Yadav, Purshottam, Kakali Kanjilal, Anupam Dutta, and Sajal Ghosh. "Fuel Demand, Carbon Tax and Electric Vehicle Adoption in India's Road Transport." Transportation Research

Part D: Transport and Environment, vol. 127, 2024, article 104010. Accessed on 21th July 2026. https://doi.org/10.1016/j.trd.2023.104010

Yadav, Rambalak, and Radha Yadav. "Are consumers ready to adopt electric vehicles? Analyzing the barriers and motivators associated with electric vehicle adoption in India: Policy implications for various stakeholders." Energy Policy, vol. 190, 2024, article 114173. https://doi.org/10.1016/j.enpol.2024.114173

Why Quantum Computing is Not Powerful Enough to Break AES By Seaena Kim

Abstract

The 21st century technological movement in encryption has contributed to the elevation of daily life through private communications, online banking, and digital commerce. Accordingly, strong cybersecurity is pivotal as new technology is introduced. Currently, methods like Elliptic Curve Cryptography and RSA are at risk in the face of quantum development. Quantum computers, using Shor's algorithm, are able to evaluate prime factors exponentially faster than classical computers. With the introduction of quantum computers, a shift towards new cryptographic standards ensued, prompted by NIST's first Post-Quantum Cryptography standards. Prior to quantum advancement, database search algorithms had a linear probability of being guessed randomly. Grover's algorithm performs search with quadratic speedup, which halves a cipher's security strength. Therefore, the next steps to combat the degradation of traditional encryption is for researchers to understand how and why certain algorithms are quantum-resistant. Potential factors include symmetric encryption, utilizing the same shared key to encrypt and decrypt data. Moreover, the size of these keys contributes to a method's security. Additionally, the process of encryption and decryption, which can be block based or stream, affects a method's strength. This paper focuses on Advanced Encryption Standard and how correlation of symmetry, size, and block-based nature impacts its invulnerability.

Introduction

The first discussion of quantum cryptography in the early 1970s, prompted by physicist Stephen Weisner, established the foundation of modern cybersecurity. One of the earliest forms of cipher, Caesar cipher was invented by Julius Caesar around 58 BCE. A Caesar cipher is a type of substitution cipher which changes each letter in plaintext according to numerical values and its alphabetical position. Invented in 1553 by Italian cryptographer Giovan Battista Bellaso, Vigenere cipher utilizes Caesar cipher by using a key word that prompts sequences of different Caesar shifts. For around 300 years, the Vigenere cipher was impenetrable until it was broken by a technique determining the length of the key word and dissecting the cipher into individual Caesar ciphers. Although modern encryption and older encryption methods are similar in that they synthesize mathematical operations and a key to encrypt and decrypt data, modern encryption standards are much more complex, relying on functions such as factoring prime numbers and logarithms.

As traditional encryption standards such as RSA and Elliptic Curve Cryptography (ECC) become increasingly vulnerable to being broken by quantum cryptography, the need for quantum cybersecurity has in turn increased. Specifically, in 2024, NIST announced its first set of post-quantum cryptography standards, including Crystals-Kyber, a quantum-resistant Key Encapsulation Mechanism. In 2016, NIST commenced the Post-Quantum Cryptography Standardization project, which aimed to protect systems from quantum threat by systematizing

public key cryptography. However, prior to this announcement researchers have articulated that the quantum threat is not yet applicable to all classical cryptography methods, notably AES.

In 1978, mathematician Robert McEliece created the McEliece cryptosystem which pioneered the implementation of randomization in public-key encryption schemes. Classic McEliece is a post-quantum key encapsulation mechanism (KEM) and is presently deemed as the most secure public-key encryption scheme, as it has been quantum-resistant for over 40 years. This system functions to incorporate errors into an encrypted message to portray as gibberish or static to any eavesdroppers without the key. The private key grants intended receivers to take out the errors and uncover the message. McEliece uses Syndrome Decoding Problem (SDP) to bolster its strength. SDP is a math puzzle that has variables matrix H, vector s, error limit t, and vector error e. Communicating a codeword c over a noisy line turns it into…

$$r = c + e$$

To calculate s, multiply the matrix H by r to the power of t. The results are the errors since the correct codewords turn into zero once multiplied by H. The original message is the error with the least mistakes or lowest Hamming weight.

Prior to the development of AES, Data Encryption Standard (DES), utilized a vulnerable 56-bit key length that soon became prone to brute-force cracking. Therefore, in the late 1990s, NIST launched a global competition which spurred the establishment of Advanced Encryption Standard (AES) in 2021, supporting key sizes of 128, 192, and 256 bits. Although it was not developed with the intention to be quantum-resistant, it does withstand modern quantum computers. Why is that?

Discussion

Using the synthesis of existing papers on quantum encryption such as Buchanan and Woodward's “Will quantum computers be the end of public key encryption” and experimental setup of the 2021 study “Surviving the Quantum Cryptocolypse” and its impact on traditional encryption systems, it can be concluded which factors impact the strength of today’s cybersecurity. Understanding quantum resilience will aid scholars in the quantum and cryptography field by prioritizing the causes and characteristics of specific functions of strong classical cryptography. Classical algorithms have developed significantly in recent years due to improvements in encryption and decryption. Furthermore, as quantum computers continue to match and exceed the most powerful classical supercomputers, the effort to combat quantum advancement has become more and more pivotal.

Symmetry

Symmetric encryption utilizes the same shared secret key in order to encrypt and decrypt data. On the other hand, asymmetric encryption uses a public key to encrypt data and a separate private key to decrypt data. Rivest-Shamir-Adleman (RSA), first published in 1977, is an example of an asymmetric encryption algorithm which uses mathematics to calculate public and private keys. RSA generates a key by multiplying two secret large prime numbers p and q to

create modulus N. Taking decryption exponent D and encryption exponent E, E can be determined by the equation:

$$E \times D = 1 mod (p-1) \times (q-1)$$

To encrypt, take the message M, raise it to the power of E and mod N to create ciphertext C. To decrypt, take C to the power of D and mod N to attain M again. As shown, the public key is both N and E while the private key is D, making RSA encryption asymmetric.

Consequently, Advanced Encryption Standard (AES), first published in 2001, is a symmetric encryption algorithm, which recently has taken center stage for its resistance against quantum attacks. Unlike RSA, AES uses the same key to convert plaintext to ciphertext and vice versa. A sequence of operations are repeatedly acted on the data to be encrypted. Based on the length of the key, 128, 192, or 256 bits, encryption is reiterated a certain number of times. The key is created through a pseudo-random number generator. It is then used in operations called SubBytes, ShiftRows, MixColumns, and AddRoundKey. Firstly, AES uses a 4 by 4 array of 16 bytes to process the data. Each byte consists of blocks of 128 bits. SubBytes, otherwise Substitute Bytes, is intended to enhance security by substituting bytes to reduce similarities between ciphertext and plaintext. S-box substitution, used in symmetric key encryption, utilizes the s-box as a lookup table to replace bytes. This step is pivotal because it is the only non-linear step of the AES round. Next, ShiftRows shifts the second row of the 4 by 4 array of bytes to the left where the leftmost byte wraps around to the right. The third row shifts two times to the left and the fourth row shifts three times to the left. Then, MixColumns uses matrix operations to change the position of each byte in a column. Finally, the AddRoundKey step applies the exclusive OR function with the round key to the 128 bits. The rounds are repeated until all the data is completely encrypted. During the last round, the MixColumns step is skipped.

To decrypt, the stages of AES are reversed: AddRoundKey, Inverse MixColumns, ShiftRows, and Inverse SubByte. The AddRoundKey function in decryption acts exactly the same as in encryption. On the other hand, InvMixColumns reverses the column rearrangement with different coefficients. Additionally, instead of left shifts, ShiftRows, rotates the data grid rows to the right. Finally, InvSubByte traces the original value by tracing back and referencing the inverse S-box. Again, for the 128 bit AES algorithm, 10 rounds are necessary to fully encrypt or fully decrypt, while for 192 bits it is 12 rounds and 256 bits it is 14 rounds. In the first round of decryption, InvMixColumns is omitted, opposite of encryption. Public key encryption lacks speed and efficiency that symmetric-key encryption has according to implementation comparisons.

Quantum Key Distribution (QKD) encryption keys can be utilized in Advanced Encryption Standard and other symmetric cipher schemes to improve security. Alternatively, one-time-pad (OTP) encryption schemes, which generate random one-use keys with lengths as long as a directed message, confirm absolute security. Key generation with QKD occurs when two parties, referred to as Alice and Bob, exchange photons through a quantum channel, simultaneously creating a shared string of random bits. Eavesdroppering can then easily be identified since these bits called qubits are altered or collapsed when it is observed. While QKD

generates the keys, the keys are stored in classical encryptors. This is where AES can be applied. Since both sides have access to the key generated by QKD, it can be implemented into the different steps of AES rounds for encryption and decryption.

Size

When the Advanced Encryption Standard was released, three standard sizes (AES-128, AES-192, and AES-256) were published. Unlike AES-128 which has 64 bits of security and is vulnerable to quantum brute-force, AES-256 has 128 bits of security, ensuring a strengthened quantum resiliency. Grover's algorithm, invented by Lov Grover in 1996, is a quantum method which beats classical computers' search speed exponentially. An unsorted database of n entries only needs around square root(n) steps using Grover's algorithm to search. On the other hand, classical computers require around n/2 steps to solve. Additionally, the effective quantum security is determined by the key size in bits divided by 2. Therefore, the size of AES-256 significantly impacts the quantum immunity of the classical cryptographic algorithm.

Although McEliece scheme may be unrealistic for daily internet use, its public key length benefits the system's security against quantum computers. Due to its extensive research and real world testing, Classic McEliece has resisted all penetration since its existence. Additionally, since the encrypted messages are full of errors, it is near impossible to decode the original message without having a private key. Quantum computers are still unable to break McEliece because of the system's lack of repetition and patterns detected. The downside to McEliece is that its public key size can be over 1 MB. This means that key generation would take longer and is impractical for common use in websites and apps.

Unlike systems like Grover's algorithm, both McEliece and AES rely on linear algebra mathematics such as matrix and vector multiplication instead of complex factoring problems. Key differences between the two schemes is that McEliece is an asymmetric cipher, requiring both a public and private key, while AES is a symmetric cipher. Moreover, McEliece has larger key sizes. Public keys range from 261 KB to over 1 MB. Since the actual ciphertext sent by Classic McEliece is generally very small, it is suitable for long-term static server keys rather than back and forth communication which can cause network traffic. McEliece thus prioritizes security over optimization. AES, in comparison, uses smaller keys for speed and convenience. While AES key is 32 bytes for AES-256, McEliece public key is around 1,000,000 bytes. AES-256 key can fit into a CPU cache and other registers, allowing modern processors to encrypt large amounts of data by the second. In order to maintain its quantum security, AES scales 128 bits to 256 bits and 16 bytes to 32 bytes which squares the work needed by a hacker to crack the message.

Furthermore, RSA, an asymmetric encryption algorithm uses mathematics for its private and public keys. A sender is given a recipient's public key to encrypt data. After the ciphered text is sent, only the intended recipient, with access to the private key, is able to decrypt the data. RSA key sizes are between 2048 and 4096 bits. Although RSA is unable to encrypt data that is larger than its key size, its significant size contributes to its strong security. Like McEliece, it

also requires a large amount of power and time to process. ECC, on the other hand, utilizes much smaller keys, ranging from 224-521 bits, around 90% smaller than RSA keys. AES uses keys that are smaller than both ECC and RSA, suggesting that it is not the size that weakens ECC and RSA, but their fundamental mathematical structure.

Block-Based Nature

Since block ciphers rely on the key space, which is the magnitude of impracticality of brute testing all possible keys, in order to minimize the shortcuts that can decrease the cracking process, substantial key length is pivotal. Block-based ciphers are especially useful for files, emails, and databases due to their security strength and versatility. Specifically, these ciphers have a high diffusion, which means that small modifications drastically impact the encrypted message. Additionally, it has a strong tamper resistance, meaning the design is extremely secure against unauthorized activity and side channel attacks. Side channel leaks include physical vulnerabilities such as power, heat, and electromagnetic leaks. Since block ciphers have access to many operations such as Electronic Codebook (ECB) and Cipher Block Chaining (CBC), it is very versatile. ECB is one of the most basic block cipher modes where it divides plaintext into chunks and encrypts each chunk separately using the same key.

AES is a symmetric block cipher that frequently uses Galois/Counter Mode (GCM). First proposed in 2007, the GCM algorithm provides data encryption and authentication. A key k, plaintext p, and data ad allows GCM to encrypt p to create ciphertext C and authentication tag t. Commonly used in AES-128, GCM uses block sizes of 128 bits. Ciphertext blocks are evaluated at point H using a finite field arithmetic and is then encrypted to create the authentication tag. Galois field multiplication uses the polynomial:

$$x^{128}+x^{7}+x^{2}+x+1$$

Then, the blocks of data are fed into GHASH, an operation that mixes and multiplies the data. GHASH is used to generate the authentication tag. GCM is the most secure when paired with a block cipher that strays from random permutation. Additionally, the strength of authentication primarily relies on the length of the authentication tag which varies among the five values, 128, 120, 112, 104, or 96. By scrambling large blocks of data, AES effectively protects against quantum computers that use Grover's algorithm to search unsorted databases at a quadratic rate.

The Twofish cipher, introduced in 1998, similarly is a symmetric block cipher that is respected for its robust security. Originally intended to replace Blowfish, which had limited block size, Bruce Schneier, John Kelsey, Chris Hall, Niels Ferguson, David Wagner, and Doug Whiting, otherwise named Counterpane Labs, entered the NIST multi-year contest to replace DES. Although they lost to AES in the 2000 competition, its complex cryptographic keys and S-boxes, like AES, have made its algorithm unbroken to this day. There are still no complete attacks against the cipher that have been successful.

Conclusion

The Rijndael cipher, 1997 winner of NIST's competition in search of a replacement for Data Encryption Standard, became the foundation for AES. Initially created with the intention of developing an encryption algorithm that would protect US government data from classical brute force, AES has exceeded by being extremely resilient against quantum computing. Through an analysis of research papers of AES, ECC, RSA, Classic McEliece, and other encryption schemes, it can be concluded that AES's symmetry benefits its performance against quantum threat. By utilizing the same key for multiple rounds of substitution and permutation, AES achieves significant mathematical complexity. Moreover, AES incorporates key sizes of 128, 192, and 256 bits. Increasing encryption rounds and broadening the number of possible key combinations contribute to AES's strength. Finally, the block-based scrambling of AES's algorithm further strengthens its protection against brute force hacking and quantum computers, as each round of scrambling exponentially increases the mathematical difficulty of decoding the ciphertext.

In order to prevent cybersecurity flaws that come with "harvest now, decrypt later", it is pivotal for researchers to explore ways to combat quantum threats by studying the current encryption schemes that are invulnerable. Uncontained quantum development fosters threats to digital finances, identity theft, and even infrastructure. This especially impacts communities in rural areas, low-income individuals, seniors, and other marginalized groups.

Works Cited

Bauer, C. P. Secret History: The Story of Cryptology, Second edition.; Chapman & Hall/CRC cryptography and network security series; CRC Press, Taylor & Francis Group: Boca Raton, 2021. https://doi.org/10.1201/9781315162539.

Mendelsohn, C. J. Blaise de Vigenère and the "Chiffre Carré."

Sharpless, K. E.; Avila, R. L.; Bruce, S. S.; Chang, W. L.; Covahey, V.; Dimeo, R. M.; Dohne, A. K.; Evans, H. M.; Fein, A. P.; Kimball, D. J.; Medina-Smith, A. N.; Munter, A. E.; Olthoff, J. K.; Poster, D. L.; Roberts, K. M.; Schiller, S. B.; Scott, J. H. J.; Silcox, B. P.; St. Pierre, J. A.; Stiles, M. D. The NIST Plan for Providing Public Access to Results of Federally Funded Research; NIST IR 8084; National Institute of Standards and Technology, 2015; p NIST IR 8084. https://doi.org/10.6028/NIST.IR.8084.

Lindsay, J. R. Surviving the Quantum Cryptocalypse.

Secure Integrated Circuits and Systems; Verbauwhede, I. M. R., Ed.; Integrated Circuits and Systems; Springer US: Boston, MA, 2010. https://doi.org/10.1007/978-0-387-71829-3.

Singh, H. Code Based Cryptography: Classic McEliece. arXiv May 29, 2020. https://doi.org/10.48550/arXiv.1907.12754.

Buchanan, W.; Woodward, A. Will Quantum Computers Be the End of Public Key Encryption? J. Cyber Secur. Technol. **2017**, 1 (1), 1–22. https://doi.org/10.1080/23742917.2016.1226650.

O'Brien, J. L. Optical Quantum Computing. Science **2007**, 318 (5856), 1567–1570. https://doi.org/10.1126/science.1142892.

Lavor, C.; Manssur, L. R. U.; Portugal, R. Grover's Algorithm: Quantum Database Search. arXiv January 16, 2003. https://doi.org/10.48550/arXiv.quant-ph/0301079.

Boneh, D. Twenty Years of Attacks on the RSA Cryptosy

Street and Hidden Homelessness in Japan - A Study of Economic Pathways, Mutual Aid, and Policy Recognition By Yudai Tanaka

Abstract

This paper examines the distinction between street homelessness and hidden homelessness in Japan through economic change, loss of mutual aid, and government policy. Japan has reduced officially counted street homelessness since the early 2000s. Nevertheless, many people without stable housing who are accommodated in internet cafés, employee dormitories, or acquaintances' homes remain beyond the scope of the official definition. This paper argues that hidden homelessness is not an attenuated version of street homelessness, but a distinct form of housing insecurity. Street homelessness was shaped largely by the collapse of older labor structures after the 1990 asset bubble burst, especially construction day labor and *yoseba* communities. Hidden homelessness, by contrast, emerged from newer forms of precarious work after the expansion of non-regular employment, especially following the 2008 global financial crisis. The paper illustrates that while street homelessness reflects the significant dissolution of mutual aid networks, hidden homelessness is characterized by the 'exhaustion' of such ties, in which social connections remain but are too fragile to provide the support indispensable for long-term housing stability. Finally, this paper contends that Japan's policy system responds more effectively once homelessness becomes visible, though it struggles to identify hidden homelessness before it eventuates into street homelessness.

Introduction

In Japan, homelessness became a visible and pressing social issue after the 1990 asset bubble collapse, which contributed to economic stagnation and employment instability (Shirahase 76). The officially recorded homeless population peaked in 2003 at about 25,000, but subsequently declined sharply to fewer than 4,000 in 2020 (Goto et al.). However, this success partly depends on how homelessness is defined domestically rather than an actual decrease in homelessness (Kakita et al.).

The Japanese government defines homelessness narrowly as only people living in public spaces such as streets, parks, riverbanks, and railway stations (Kakita et al.). Beyond this official classification, however, a sizable population lacking stable housing resides temporarily in places such as internet cafes, low-cost lodging facilities, or acquaintances' homes (Kakita et al.).

This paper designates such individuals as the hidden homeless population. Although they are not counted as "homeless" under the Japanese definition, they would be included inside broader international frameworks, such as the ETHOS (European Typology of Homelessness and Housing Exclusion) definition, which recognizes insecure and inadequate housing as forms of homelessness (Kakita et al.).

The significance of this study lies in the fact that it does not treat street homelessness and hidden homelessness as merely different degrees of the same problem. This study argues that these two forms of homelessness are shaped by different social and economic causes. Although

prior research has been done on specific aspects of street and hidden homelessness individually (Kiener, “Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan”; Okura), there are relatively few papers that investigate the conceptual divide since many perceive them as points on a single continuum. Acknowledging this difference is crucial because it may enable policy intervention before housing insecurity leads to street homelessness. By comparing the economic, relational, and policy mechanisms behind street and hidden homelessness, this paper contributes to a more preventive understanding of homelessness in Japan. Therefore the research question is “How do street and hidden homelessness in Japan differ in their economic, relational, and policy pathways, and what are the implications of these differences for homelessness policy?”

This paper compares existing research, government data, and policy materials on street and hidden homelessness in Japan across three dimensions: economic pathways, mutual aid, and policy recognition.

Background and History

The concept of hidden homelessness is closely tied to stigma and social visibility. As street homeless emerged in streets and parks, meanwhile, they were recognized as being in a situation of poverty; they were negatively labeled as "unhygienic" or “disruptive”. Such societal perception helped reinforce an emphasis on individual responsibility, frequently framing homelessness as a consequence of personal shortcomings or insufficient individual effort, rather than as an outcome of broader structural or economic constraints (Kiener, “Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan” 85).

This tension was reflected in the ambiguous provision of the Act on Special Measures concerning Assistance in Self-Support of Homeless People (hereafter the Homeless Self-Reliance Support Act) enacted in 2002 (Okuda). Although the law stressed the human rights and self-determination of homeless individuals, Article 11 also permitted “necessary measures” when homeless people’s use of public space interfered with its normative use. In practice, this imprecision enabled public authorities to implement exclusionary measures, such as the removal of makeshift encampments from parks in Osaka. As a result, street homelessness carried a strong stigma: it meant not only the loss of housing, but also becoming publicly visible as a socially and politically defined problem.

There are also marked differences between the demographics of the hidden homeless population and the street homeless population. Street homelessness in Japan is disproportionately concentrated among older men. More specifically, around 70% of the officially counted homeless population is over the age of 60, and women account for only approximately 4% of the population. Hidden homelessness, by contrast, encompasses a larger share of younger people and women. In the internet survey across 14 prefectures analyzed by Kakita et al., the hidden homeless population had a much lower share of the population over 60, while the share of women rose to around 47% (Kakita et al.; Ministry of Health, Labour and Welfare, ホームレスの実態に関する全国調査（生活実態調査）の結果（詳細版）2-3).

Prior research suggests that street homelessness in Japan is made up predominantly of older men partly because it grew out of the collapse of traditional, male-dominated jobs (Kiener, "Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan"). Many of these men remained homeless because advancing aging, loss of physical capacity, weak social ties, and unstable work histories limited their ability to obtain employment or regular rental housing (Furugori and Li 105). Hidden homelessness looks different. They tend to be younger and include more women because it often comes from newer forms of unstable work, insecure housing, and situations that keep women out of public view (Fujita). Women's homelessness is often inconspicuous. Many low-income females are inclined to live in acquaintances' houses, women's shelters, or domestic-violence support, thus not visible on the streets.

Causes of Street Homelessness

This paper argues that street homelessness in Japan was not produced by a single cause but by the interplay of economic change, the loss of mutual aid, and government policy responses to poverty. These determinants worked together: economic restructuring removed unstable occupations, weakened social ties made it harder for people to rely on family or community scaffolding, and welfare policy often became effective only after homelessness was already visible. This section therefore examines these three causes separately in order to show why some people entered visible street homelessness rather than remaining in hidden forms of housing insecurity.

Economic change

Goto et al. argue that the economic change that in part precipitated current homelessness in Japan can largely be traced back to the collapse of the asset bubble in the early 1990s (Goto et al.). Japan's asset price bubble, fueled by monetary easing and rapid credit expansion that inflated stock and real estate prices, collapsed when the Bank of Japan tightened monetary policy and imposed lending restrictions, triggering a sharp decline in asset values. In its aftermath, the unemployment rate increased from 2% to 5% over 10 years, from 1991 to 2002 (Agnese and Sala 16). To dampen economic volatility, the government implemented labor market reforms aimed at increasing flexibility.

Japan gradually expanded the legal use of dispatch labor, initially through the Worker Dispatching Act of 1985 and later revisions. Dispatch workers are formally employed by one company but assigned to work for a client, often on short-term or unstable contracts. This leaves them more vulnerable to dismissal than regular full-time employees, also within non-regular employment.

Okamoto argues that, however, dispatch work did not immediately become a major cause of street homelessness in the early 1990s, because it was initially limited to circumscribed industries (Yoshiro Okamoto 19). Therefore, in the framework of this paper, dispatch labor is better understood as a background to later hidden homelessness, while the earlier rise of street

homelessness was rather more directly intertwined with the expansion of other non-regular employment. Non-regular employment in Japan encompassed multiple forms of work, including part-time and *arubaito* work, fixed-term contract work, temporary agency work, dispatch work, and day labor. These forms of employment, however, did not contribute to homelessness in the same way. In the case of visible street homelessness in the 1990s and early 2000s, the most consequential forms were not conventional part-time jobs, but construction day labor and other forms of casual work embedded in *yoseba*, company dormitories, or live-in work arrangements. This was because the loss of such work could entail not only the loss of income, but also the loss of access to accommodation, job information, and worker communities (Kiener, "Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan" 84). Driven by imperatives of global competitiveness and labor cost rationalization, firms increasingly institutionalized precarious employment. The initial increase in street homelessness was driven not by the popularization of dispatch work itself, but rather the strategic expansion of other non-regular work by firms. Companies not only reduced regular employment but also began dismantling what Japanese companies traditionally emphasized, such as company-provided welfare benefits, employee dormitories, housing-related benefits, and housing allowances. As a result employees had weaker ties with the company they work in (Yoshihiro Okamoto 52).

This paper argues that this precarious form of employment formed an important pathway into street homelessness. Between 40 and 60% of people who experienced street homelessness during the 1990s had worked as day laborers. In addition, 75% of those who became homeless after moving into insecure employment were contingent upon company dormitories (Kiener, "Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan" 84, 86). For precarious workers, job loss frequently entailed an absence of housing and an inability to find a new one. While housing-lost workers in Tokyo earned only about 114,000–120,000 yen per month; 62.8% reported difficulty saving initial move-in costs, 33.3% lacked stable income to continue paying rent, and 30.9% had difficulty procuring a guarantor (Yoshihiro Okamoto 55; Tokyo Metropolitan Government Bureau of Social Welfare and Public Health 39, 48).

The situation was further exacerbated by the contraction in the construction and manufacturing industries, which had previously served as a key safety net for employment. In 1967 construction covered around 40% of the day-labor work-force in Japan, while manufacturing employed 27%. However, the rise in street homelessness cannot simply be explained by the further decline in total construction employment. Although the sector continued to employ a large number of workers, employment opportunities available through *yoseba* decreased, owing to intensifying global competitiveness (Banerjee). Later on, construction employment contracted markedly after its late-1990s peak: the number of construction workers peaked at 6.85 million in 1997 and had fallen by approximately 1.8 million by 2010 (Furugori and Li 105).

Kurasawa argues that the erosion of *yoseba*, a notion peculiar to Japan, was an important factor in causing the unemployed to move directly into public spaces when the bubble burst. A *yoseba* is a place in Japan where day laborers (often from the construction industry) gather to find provisional work, associated with socially marginalized communities. These communities were important to the population in poverty, since they were not exclusively places to work but also places that offered food, housing, and a sense of community. Although the construction industry's decline mentioned above led to *yoseba*s disappearing from the early 2000s, *yoseba* workers were only able to secure 6-7 days of work per month. In fact, in Osaka, more than half of the homeless population has experienced living in a *yoseba* (Kurasawa). This concentration of work, accommodation, and social ties helps explain why the erosion of yoseba could affect housing status more directly than a loss of wages alone.

In addition to these compounding structural changes, this bubble burst occurred when the "Dankai Sedai", the post-war baby-boom generation (approximately 2.6 million births per year), was in their mid-40s, when they were most susceptible to such economic changes (Yoshiro Okamoto 20). Employers often avoided hiring workers over 50 because of concerns about physical capacity and workplace accidents. As a result, older day laborers who lost *yoseba*-based work had few alternative employment opportunities. This made them more vulnerable than younger workers, who could still enter low-wage non-regular employment during the 1990s.

Taken together, this paper argues that a major pathway into street homelessness in Japan involved the collapse of older labor structures, especially construction day labor and yoseba, which removed both employment and the community-based support that had previously kept poor workers from immediately entering the streets.

Mutual aid

Loss of mutual aid denotes the weakening of social embeddedness among friends, family, colleagues, and the local community that traditionally provided provisional housing, limited financial assistance, and emotional scaffolding. Okamoto argues that it is not a stand-alone cause of homelessness, but rather an amplifier that increases the likelihood of already vulnerable people becoming street homeless (Yoshiro Okamoto 23).

Prevailing Japanese legal frameworks such as the Civil Code and the Public Assistance Law formally impose a duty or prioritization of support on direct kin and siblings (Kiener, "Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan" 83-84). In practice, however, these familial support systems are often weakened or nonfunctional, reducing families' capacity to absorb members in crisis (Yoshiro Okamoto 23). The size of households has markedly contracted in the past 60 years, from around four to two people. In fact, among people experiencing homelessness, more than 20% were raised by single-parent households or without parental care (Kiener, "Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan" 85), with 70% not having had contact with a family member for the past year (Bureau of Social Welfare and Public Health, Tokyo Metropolitan Government 65). Thus, people who are homeless

on the street scarcely acquire mutual aid, the most reliable source of tie being mutual aid from another homeless person or an outreach worker.

Mutual aid is important not only because it provides direct support but also because it mediates private support and public aid. In many cases, people living on the streets have difficulty navigating information about aid and may be discouraged from seeking assistance because of the stigma. In those situations, the accompaniment of someone they trust, such as a family member, an acquaintance, or an outreach worker, often helps the application process (Hayashi). In fact, 51% of people receiving Seikatsuhogo first learned about the system through acquaintances and family members in a public survey in Odawara City (Odawara City 6). Here, Seikatsuhogo refers to Japan's last-resort public assistance system, which provides minimum living, housing, and medical assistance to people who cannot maintain a minimum standard of living through their own resources.

The erosion of mutual aid also creates serious barriers in the housing market. Once individuals lose accommodation, securing new accommodation can be difficult, especially in a system where guarantees remain customary (Yoshihiro Okamoto 52). In addition to the substantial initial move-in costs, a personal guarantor (rent guarantor) is usually required - someone, often a family member, who is legally responsible for paying the rent if the tenant cannot. In this sense, the absence of mutual aid not only increases the risk of homelessness but also prolongs its duration. Therefore street homelessness in Japan is a situation where relationships that can be relied on have collapsed, where there is barely anyone that the homeless person can ask for accommodation, ask to be a guarantor, and introduce, or explain aids.

Policy

The public-sector interventions aimed at reducing poverty can be conceptualized through three layers (Kiener, *Homeless Policies from Day Laborer Ghettos to the Entrepreneurial Welfare City: An Account on Osaka City's Changing Geographies of Public Assistance* 13): The three layers are social insurance, employment and self-reliance support such as the Self-Reliance Support System for the Needy, and lastly, Seikatsuhogo.

This section centers primarily on Seikatsuhogo, which is the policy most directly relevant to homelessness. The Seikatsuhogo is designed for individuals who "still face financial hardship despite making full use of all their assets and abilities". By contrast, the other two layers of policy only provide what would be regarded as temporary aid to the homeless population rather than consistent housing, income, and employment support. As a result, the homeless population oftentimes cannot receive sufficient aid from it.

Following the official enumeration of the homeless population reaching its highest number in 2003, Seikatsuhogo became one of the central policies used to reduce visible street homelessness. Although Seikatsuhogo has existed, its application to the homeless population was restricted. Not having an address, or even being slightly incapable of working, sometimes became an obstacle to receiving aid. Once the government treated the growing homeless population as a policy issue requiring direct intervention, it largely clarified and liberalized the

requirements for receiving Seikatsuhogo, playing a part in a sharp reduction in homelessness. Goto et al. argues that the decline in visible homelessness after the early 2000s cannot be explained only by economic recovery; it also reflected a change in how public assistance was applied to people living on the streets (Goto et al.).

In practice, Seikatsuhogo serves as a last-resort safety net for people in poverty, as it addresses dimensions of poverty simultaneously. It does not simply provide a single cash payment; rather, it combines different modalities of assistance based on the recipient's situation, including daily living expenses, housing assistance, medical assistance, and other forms of support, such as education or job-training aid. For people experiencing homelessness who are struggling to procure indispensable items to even live on the streets, these forms of assistance can help them exit the severe condition of street homelessness.

The process of Seikatsuhogo starts when an individual applies to a welfare office. Then, a caseworker investigates the individual's income, assets, health, and work capacity, providing the difference between the individual's income and the country's minimum living cost.

Goto et al. (2021) describe a case in Shinjuku where all 924 Seikatsuhogo applications submitted by people living on the streets were approved. Among these applicants, 829 moved into single-room units, 33 entered welfare facilities, and 62 were admitted to hospitals.

Goto attributes this success to its three main characteristics: generality, comprehensiveness, and expeditiousness (Goto et al.). First, in terms of generality, Seikatsuhogo is not limited to any specific type of poverty. It is formally available to anyone who meets the criteria. Second, its comprehensiveness lies in its structure, which combines all cash benefits, housing assistance, and other aids, if needed, into a single system rather than dividing them into separate programs. For a male in Tokyo, a combined maximum benefit of 126% of the relative poverty line could be earned from the Seikatsuhogo (Goto et al.; Melbourne Microfinance Initiative). Third, the system is expeditious: The application and certification process is relatively swift. Principally, it is completed within 14 days.

For street homelessness, public legibility is both a disadvantage and a policy advantage. Living in parks, stations, riverbanks, or streets exposes people to stigma, surveillance, and exclusion, but it also makes their poverty recognizable to the state as an urgent public problem. Therefore, the main policy problem for people experiencing street homelessness is not whether the government can identify them, but whether they can access Seikatsuhogo once they have become visible. In this sense, street homelessness is primarily a problem of access to a strong last-resort safety net.

Causes of hidden homelessness

This paper argues that hidden homelessness and street homelessness in Japan stem from the same three fundamental vulnerabilities, although these factors work in different ways and create different outcomes. Street homelessness is more closely related to the destruction of older labor structures and *yoseba*s. Meanwhile, hidden homelessness is more likely to be created from newer labor forms, such as low-wage irregular employment and fragile mutual ties. This

comparison matters because hidden homeless individuals seem to have more capacity to secure housing, though they still lack the ability. Therefore, hidden homelessness should not be treated simply as a less severe version of street homelessness, but it should be regarded as a separate form of housing insecurity that requires earlier recognition and preventive intervention.

Economic change

Economic change likewise constitutes an important dimension of hidden homelessness. The average income for people experiencing hidden homelessness in Tokyo stands at around 100,000 yen per month, which is marginally higher than that of people experiencing street homelessness. Nonetheless, it remains inadequate for secured housing in Japan, considering the barriers when moving in, such as rent, initial move-in costs, difficulty finding a guarantor (Tokyo Metropolitan Government Bureau of Social Welfare and Public Health 39). Entering the 2000s, the national population that earns less than 2 million yen per year increased to 30%. This did not automatically produce hidden homelessness, but it left a substantial percentage of the workforce living close to housing insecurity.

In contrast to street homelessness, hidden homelessness was shaped by a different economic process. While street homelessness was closely connected to the collapse of older labor structures after the burst of the asset bubble, hidden homelessness came into sharper view after the 2008 global financial crisis (Fujita; 特定非営利活動法人ビッグイシュー基金 1). By the time the crisis occurred, it was common for a large part of the country's population to depend on non-regular work. The global crisis, therefore, is not the singular cause; rather, it exposed the already vulnerable workers created by changes in labor forms.

The 2008 global financial crisis (Lehman Shock) refers to the several financial catastrophes provoked by the bankruptcy of the Lehman Brothers investment bank in September 2008. The crisis resulted in a worldwide credit crunch, sharp declines in financial markets, and wide-ranging consequences on employment and economic activity. This shock reduced credit availability and consumption in Europe and the US. When consumers and companies abroad decreased spending, Japanese companies, which extensively depended on exports and foreign customers, were compelled to reduce production. This decline in production meant firing workers for most companies, since non-regular employment was already common and layoffs were less complicated than before. As seen in The Big Issue Foundation reports, in the year after the crisis, approximately 240,000 non-regular workers experienced termination or non-renewal of employment, and that 2–3% of them also lost their housing (特定非営利活動法人ビッグイシュー基金 1; Iyetomi et al.).

The crisis drew particular attention to the young generation, ranging from 20 to 30 years old, slightly younger than the age range for street homelessness. Unlike older street homeless men, many of these younger workers had never experienced a stable period of work that could provide savings, confidence, or durable social ties. Accordingly, job loss not only signifies a loss of income but also makes it apparent that the individual lacks a stable job history sufficient to maintain paying capacity for the rent. This was particularly salient for younger workers because

by the late 2000s, many had already entered unstable forms of employment. This differed from the 1990s, when younger workers could often still find low-wage non-regular employment while middle-aged and older workers were more likely to descend into street homelessness after losing construction or day-labor work.

By the time of the Lehman Shock, however, non-regular employment had become much more widespread among younger workers. One example is the 1999/2004 revision of the Worker Dispatching Act, which greatly expanded the range of occupations in which dispatch work was legally permitted. As a result, more firms relied on flexible labor, and more young workers entered precarious employment.

In fact, non-regular workers increased from about 14 million in 1997 to 18 million in 2007, and while non-regular workers account for around 30% of the street homeless population (Ministry of Health, Labour and Welfare, ホームレスの実態に関する全国調査（生活実態調査）の結果（詳細版）80), the corresponding proportion is more than 70% for hidden homelessness (Ministry of Health, Labour and Welfare, ホームレスの実態に関する全国調査（生活実態調査）の結果（詳細版）; Tokyo Metropolitan Government Bureau of Social Welfare and Public Health 12). This helps explain why hidden homelessness is linked much more strongly to non-regular employment than street homelessness is.

Unlike street homelessness, hidden homelessness is less intertwined with collective labor communities such as *yoseba*. People experiencing street homelessness were often linked to day-labor markets, corporate dormitories, or welfare facilities, which gave them at least some institutional or social connection. In contrast, people experiencing hidden homelessness are more institutionally detached individuals without stable communities or support networks, and their work also tends to be more fragmented and informal. Kitagawa links this shift to the widespread use of smartphones (Kitagawa 40). The popularization of smartphones allowed people with difficulty in finding jobs to find one without the need to gather at a specific place. Consequently, workers who once gathered in one spot to find a job lost their valuable and scarce mutual ties. For instance, rather than construction or other day-labor jobs connected to *yoseba*, people who engage in a wider range of work, including simple, repetitive tasks, have emerged. This difference shows that hidden homelessness is not only a result of low income, but also of the disaffiliation of social and institutional places where poor workers could temporarily accommodate, find work, or receive support (Fujita). Hidden homelessness emerged from a different economic pathway compared to street homelessness - Young non-regular workers possess the capacity to work, though they live an unstable life because of inadequate income and lack of stable housing available.

Mutual aid

Mutual aid appears to play a particularly salient role in hidden homelessness because the issue is often not the complete absence of work but the insufficiency of usable social ties. Okura contends that these people should be understood as experiencing fluid, unstable living arrangements, supported neither by secure employment nor by family or public assistance.

Hidden homelessness in Japan should thus be understood not only as a labor-market problem but also as a problem of relational disconnection (Okura).

In this population, the loss of mutual aid is experienced in many ways. First, family ties are often formally present but practically unusable. Compared to street homelessness, the central issue is not always in the complete absence of family ties. Even though family ties exist, studies repeatedly report cases in which respondents had living parents or relatives yet were unable to return home because of conflict, disownment, family poverty, or strong norms of self-reliance. In fact, 13% of the hidden homeless population lost housing for family-related reasons (Tokyo Metropolitan Government Bureau of Social Welfare and Public Health 33). In this sense, the issue is not simply the absence of family, but the breakdown of family as a dependable safety net. Second, friendship and acquaintance networks are seldom robust to provide durable shelter, but they may still function as a safety net compared to street homelessness. 41% of hidden homeless respondents and around 60% of street homeless respondents, respectively, reported having no one or a few people they could consult about serious problems (Ministry of Health, Labour and Welfare, ホームレスの実態に関する全国調査（生活実態調査）の結果（詳細版）72; Tokyo Metropolitan Government Bureau of Social Welfare and Public Health 62). This rate is lower for people experiencing hidden homelessness, though it is still a serious problem. At the same time, in the Tokyo survey, around 20% of the hidden homeless population live in an acquaintance's house, suggesting that these ties still retain greater practical relevance than street homelessness (Kakita et al.).

In terms of the consequences that will be caused by this loss of mutual aid, people experiencing hidden homelessness face comparable issues to people experiencing street homelessness, notably with rent guarantors and aid access. However, the guarantor problem is less pronounced because of the stronger mutual aid among people experiencing hidden homelessness. Contrary to expectations, though, fewer people learn about assistance through acquaintances. This does not necessarily suggest that their social ties are weaker: rather, it may be because people experiencing hidden homelessness are often in a less complex situation and are applying to modalities of aid that are easier to access and carry less stigma than Seikatsuhogo. This paper argues that therefore, Hidden homelessness in Japan is not a condition in which mutual aid has been completely cut off, but rather a condition in which mutual aid has become fragile and inadequate. Social ties may still supply temporary shelter, although a guarantee, secured housing, and an appropriate explanation for aid are hard to obtain.

Policy

Policy constitutes one of the clearest domains in which hidden homelessness differs from street homelessness in Japan. The principal policy directed towards individuals who are at risk of falling into poverty or are already experiencing financial hardship is the Self-Reliance Support System for the Needy, which was launched nationwide in 2015 (生活困窮者自立支援のあり方等に関する論点整理のための検討会; Ministry of Health, Labour and Welfare, 生活困窮者自立支援制度における 横断的な課題について 9). By offering support at a stage prior to the

receipt of Seikatsuhogo, the system aims to promote self-sufficiency before problems become more complex and severe. With about 50 billion yen allocated to this system, it provides housing-related support, employment, and household financial management support. In contrast to Seikatsuhogo, which secures a minimum standard of living through a guaranteed package, this system is an accompaniment-based, case-management-oriented framework in which case workers formulate an individualized support plan with each applicant for a stable and sustainable livelihood.

In daily practice, caseworkers meet with individuals to understand their circumstances, collaborate with them to develop a plan for support, and help link them to services such as housing, job placement, financial advice, and daily life assistance. The program's reach is significant: in the past ten years, over 3 million people have received consultations, around 300,000 have achieved employment or increased income and 900,000 support plans have been established (Ministry of Health, Labour and Welfare, 生活困窮者自立支援制度における支援状況調査 集計結果 1).

A persistent limitation, however, is that it does not sufficiently address stigma. While outreach is considered essential, only a few local governments have developed effective outreach systems for people experiencing hidden homelessness. In fact, only 40% of the consultations were initiated by the applicant him/herself (Ministry of Health, Labour and Welfare, 生活困窮者自立支援制度における 横断的な課題について). A related issue is that local governments often do not collaborate sufficiently with local facilities and support organizations, which restricts their ability to mobilize resources. Ministry survey materials found that about half of municipalities said they knew local support resources only “to some extent,” about 30 percent said they could not grasp them well, and more than 60 percent regarded local support resources as insufficient (Ministry of Health, Labour and Welfare, モデル事業実施状況調査 集計結果について).

This paper argues that the policy contrast between street and hidden homelessness can be understood as one between access and recognition. For the street homeless population, public conspicuousness has two conflicting meanings. For an individual, it may mean being exposed to the public, therefore experiencing a strong stigma. On the other hand, it may signify being easier to be found by outreach workers, allowing a higher possibility of policy intervention. For most people experiencing hidden homelessness, the main question is whether the state can recognize them before they appear on the street. The Self-Reliance Support System for the Needy is designed for this earlier stage, but its effectiveness depends on outreach, local collaboration, and the ability to detect housing insecurity in private or semi-private spaces. In this sense, Japan’s welfare system is relatively strong after homelessness becomes visible but weaker at preventing invisible housing insecurity from becoming street homelessness.

The main differences are identified in Table 1.

Cause/Type	Street Homelessness	Hidden Homelessness

Economic Change	- Asset bubble burst - Destruction of older forms of employment and *yoseba*	- Lehman Shock - Prevalence of irregular employment
Mutual Aid	- Completely cut ties with family/acquaintance - Stronger ties with outreach worker	- Fragile but existing ties with family/acquaintances/outreach worker
Policy	- Stronger stigma, but better outreach - Mainly financial support	- Less stigma, but weaker outreach - Stronger accompany-based support

Table 1. Different causes of street and hidden homelessness.

Conclusion

This study examined how street and hidden homelessness in Japan differ in their economic, relational, and policy pathways. It argues that hidden homelessness should not be regarded as a prior stage of street homelessness. Rather, it is a distinct form of housing insecurity. Street homelessness was mainly shaped by the collapse of old labor structures after the 1990 asset bubble burst, especially construction day labor and *yoseba*, which had once provided poor workers with work, shelter, and community. Hidden homelessness, by contrast, emerged from a new form of economic insecurity following the expansion of non-regular and dispatch work, especially after the Lehman Shock. Similar differences can be seen in mutual aid. Street homelessness often reflects the collapse of usable mutual aid, whereas hidden homelessness is characterized by social ties that remain but are fragile. Even when able to obtain temporary housing, securing stable housing and receiving sufficient aid may be difficult. In policy terms, street homelessness is more visible to welfare and outreach systems, whereas hidden homelessness is more difficult to identify before housing insecurity becomes severe.

The central problem, therefore, is not only poverty itself but the timing of recognition of need. Japan's welfare system is relatively strong once homelessness becomes visible, but weaker when homelessness remains hidden inside internet cafés, employee dormitories, temporary workplaces, or acquaintances' homes.

One limitation is that the evidence comes from sources that differ in time, geography, and population. Future research should examine how to identify and reach people experiencing hidden homelessness before their housing insecurity becomes more severe. Because the sources differ in time, geography, and population, the comparisons are interpreted descriptively rather than causally.

Works Cited

Agnese, Pablo, and Hector Sala. “Unemployment in Japan: A Look at the ‘Lost Decade.’” Asia Pacific Journal of Economics and Business, vol. 12, no. 1, 2008, pp. 16–35.

Banerjee, Arpan. “Marginality in Japanese Society : Globalisation, Nojukusha, and Homelessness in Contemporary Japan.” Shodhshauryam International Scientific Refereed Research, vol. 6, no. 4, 2023.

Bureau of Social Welfare and Public Health、Tokyo Metropolitan Government. 令和3年度 ホームレスの実態に関する全国調査（東京23区の生活実態調査結果）. Survey Report. Tokyo Metropolitan Government、Apr. 2022.

Fujita, Hirohito. “新たなホームレス問題の展開.” 愛知県立大学教育福祉学部論集, vol. 61、2013、pp. 29–35. https://doi.org/10.15088/00001566.

Furugori, Tomoko, and Qingya Li. “路上生活者になる社会的背景とその決定要因の分析.” 商学論纂, vol. 55, no. 3, Mar. 2014, pp. 95–117. CiNii Research, https://cir.nii.ac.jp/crid/1050282677692349056.

Goto, Hiroshi, et al. “Why Homelessness Has Decreased in Japan: Comparison of Public Assistance in Japan and the U.S.” 2021, 15th European Research Conference on Homelessness, https://researchmap.jp/7000004790/presentations/33806609.

Hayashi, Shingo. “生活保護受給者が社会的孤立に至るプロセス : SCAT を用いたインタビュー分析からの考察.” 都市と社会 , vol. 9、2025, pp. 12–21、https://omu.repo.nii.ac.jp/records/2002776.

Iyetomi, Hiroshi, et al. “What Causes Business Cycles? Analysis of the Japanese Industrial Production Data.” Journal of the Japanese and International Economies, vol. 25, no. 3, Sept. 2011, pp. 246–72. ScienceDirect, https://doi.org/10.1016/j.jjie.2011.06.002.

Kakita, Yusuke, et al. “Homelessness and Housing Exclusion in Japan from an International Perspective.” Conference presentation / Presentation slides. 16th European Research Conference on Homelessness, Bergamo, Italy.

Kiener, Johannes. Homeless Policies from Day Laborer Ghettos to the Entrepreneurial Welfare City: An Account on Osaka City’s Changing Geographies of Public Assistance. 2019. Osaka City University, Doctoral Dissertation.

———. “Social Networks of Homeless People under the Influence of Homeless Self-Sufficiency Support Centres in Japan.” Vienna Journal of East Asian Studies, vol. 5, 2014, pp. 77–109, https://doi.org/10.2478/vjeas-2014-0004.

Kitagawa, Shinya. “「社会の総寄せ場化」における労働移植のロジスティクス——外国人技能実習制度、移動の自律性、流動的下層労働者.” 季刊経済理論, vol. 60, no. 2, July 2023, pp. 39–49, https://doi.org/10.20667/peq.60.2_039.

Kurasawa, Aiko. 都市下層の生活構造と移動ネットワーク. 明石書店, 2007.

Melbourne Microfinance Initiative. “Reducing Homelessness in Japan, a History of Systemic and Cultural Oversight.” Melbourne Microfinance Initiative, 8 May 2024, https://www.melbournemicrofinance.com/new-blog/2024/5/8/reducing-homelessness-in-japan-a-history-of-systemic-and-cultural-oversightnbsp.

Ministry of Health、Labour and Welfare. ホームレスの実態に関する全国調査（生活実態調査）の結果（詳細版）. Survey report. Ministry of Health、Labour and Welfare, 6 Apr. 2022, https://www.mhlw.go.jp/stf/newpage_25330.html.

———. 生活困窮者自立支援制度における 横断的な課題について. Meeting Document 資料1、Ministry of Health、Labour and Welfare, 3 Dec. 2021. 第1回生活困窮者自立支援のあり方等に関する論点整理のための検討会 横断的課題検討班.

Ministry of Health、Labour and Welfare. モデル事業実施状況調査 集計結果について. Meeting Document 資料8、Tokyo、Japan, 26 Sept. 2014.

———. 生活困窮者自立支援制度における支援状況調査 集計結果. Statistical Report. Ministry of Health、Labour and Welfare、2025.

Odawara City. 生活保護利用者アンケート調査の結果. Odawara City、2019、https://www.city.odawara.kanagawa.jp/field/welfare/p-cassi/p22170.html.

Okamoto, Yoshihiro. “Japanese Social Exclusion and Inclusion from a Housing Perspective.” Social Inclusion, vol. 4, no. 4, Oct. 2016, pp. 51–59. www.cogitatiopress.com, https://doi.org/10.17645/si.v4i4.628.

Okamoto, Yoshiro. “The necessity of comprehensive policy for homeless issues in Japan.” JOURNAL of POLICY STUDIES CHUKYO UNIVERSITY, vol. 10, Mar. 2019, pp. 17–28. CiNii Research, https://cir.nii.ac.jp/crid/1050001201666884992.

Okuda, Toki. “Developing Homelessness Support through Public-Private Partnerships.” 明治学院大学大学院社会福祉学, vol. 49, Mar. 2025.

Okura, Yuji. “「ネットカフェ難民」の存在は社会的に何を意味するか.” 寄せ場 : 日本寄せ場学会年報、no. 2010 May 2010、pp. 96–107. CiNii Research、https://cir.nii.ac.jp/crid/1523106604590771840.

Shirahase, Tatsuya. “あいりん地域における居住支援──ホームレス支援の新たな展開と課題──.” 理論と動態, no. 7, 2014, pp. 76–91, https://www.jstage.jst.go.jp/article/istd/7/0/7_05/_article/-char/ja.

Tokyo Metropolitan Government Bureau of Social Welfare and Public Health. 住居喪失不安定就労者等の実態に関する調査報告書. Tokyo Metropolitan Government, Jan. 2018.

特定非営利活動法人ビッグイシュー基金. 若者ホームレス白書. White Paper. 特定非営利活動法人ビッグイシュー基金、Dec. 2010.

生活困窮者自立支援のあり方等に関する論点整理のための検討会. 生活困窮者自立支援のあり方に関する論点整理. Government Committee Report. Ministry of Health、Labour and Welfare、2017.

Reduced-Order Modeling of Nano-Aluminum Additives in Ethanol/LOX Rocket Propellants By John Li

Abstract

This study seeks to investigate the effects of nano-aluminum particles on the performance of Ethanol/LOX propellants [8] in both idealized and physical environments, integrating in factors such as loading fraction and mixture ratio. RocketCEA was used to establish the baseline Ethanol/LOX performance, and a reduced-order analytical model was then used to estimate the effects of nano-aluminum loading. The baseline model used a chamber pressure of 15 bar, a nozzle expansion ratio of 40, and a reference mixture ratio of 1.85.The findings show a generally positive correlation as when loading fraction increases, combustion temperature and specific impulse also increase as well. The model predicts increasing performance across the tested 0–5 wt% loading range, although higher loading also increases condensed-particle mass and oxygen demand. At the reference mixture ratio of 1.85, the baseline chamber temperature and vacuum specific impulse were 3259.05 K and 342.87 s, respectively. At 3 wt% passivated nano-aluminum loading, the reference model predicts a temperature increase of approximately 193.3 K, a thermal specific-impulse increase of 10.02 s, and a momentum-adjusted increase of 7.49 s. Sensitivity studies show that the predicted performance depends strongly on particle size, oxide-shell thickness, burn-time assumptions, heat coupling, and particle/gas velocity slip. These results are only to be interpreted as sensitivity estimates, not experimentally validated engine forecasts.

Introduction

Rockets have a wide variety of uses, ranging from commercial to defense, with many being related to interplanetary exploration, satellite communications, and reconnaissance. In all rocket-related disciplines, there is a single issue that is the most common hindrance to progress: efficiency. According to Tsiolkovsky's equation [1], adding extra propellant requires even more propellant, so this is where efficiency comes into play. A typical expendable rocket is 85-95% propellant by mass, so even a small gain in specific impulse directly translates to a more deliverable payload. By adding reactive metallic powders, such as boron, magnesium and aluminum, researchers have found increases in specific impulse/density specific impulse, allowing for higher efficiency.

Past research on micro-sized aluminum particles has found that the trade-offs in performance caused by agglomeration and noncombustion cause it to fall behind on its theoretical deliveries in improvement. According to Yuan et al. [2], the agglomeration of aluminum particles leads to incomplete combustion, and for every 10% increase in unburnt aluminum there is approximately a 1% loss in specific impulse, while also causing other losses in performance. Additionally, a review by Deluca et al. [3] determined that because micron-sized aluminum particles burn far from the surface it experiences long ignition delays and a distributed

flame zone, causing a slower combustion time and only small increases in the burning rate, respectively.

Nano-Aluminum particles may alleviate these issues as it should burn faster and more completely, with less ignition delay. In a study by Jayaraman et al. [4], it was determined that nano-aluminum can increase the burning rate of AP/HTPB/Al propellants by up to 100% compared to micron-sized Al particles. This boost in burning rate could be attributed to the reduction in size of the particles as it dramatically increases the surface-area-to-volume ratio, meaning more reactive metal surface is exposed to the oxidizer. Nano-aluminum containing methanol and ethanol fuels also showed an increase in average flame-propagation speed by 48% in ethanol with 2 wt% aluminum, compared to pure ethanol [10]. A separate study by Galfetti et al. [5] found that the combustion products of nano-aluminum particles have lower agglomeration tendencies, meaning more efficient burning of aluminum during combustion [6]. Many of the articles mentioned contain solid propellants rather than Ethanol/LOX, although many characteristics of the additives are independent of the propellant type.

Previous investigations have shown, however, that size reduction is simply not enough to alleviate issues such as agglomeration or formation of oxide layers [7]. Current research is focused heavily on surface coatings and developing aluminum alloys to improve propellant energy performance by reducing two-phase losses and improving vacuum specific impulse (Isp). Nonetheless, a systematic analysis of the properties of nano-aluminum (e.g. loading fraction, in-chamber release fraction, etc.) on combustion temperature and specific impulse across varying combustion efficiency assumptions remains limited in the literature. Overall, we investigated whether nano-aluminum particles can improve the performance of Ethanol/LOX rocket propellant under specified assumptions regarding loading and in-chamber release fractions. It is predicted that, under assumed ideal conditions, as the loading fraction and in-chamber release fraction increase, the combustion temperature and specific impulse of the propellant will also increase. This study asks how passivated nano-aluminum may affect the modeled thermal and momentum performance of an Ethanol/LOX propellant. RocketCEA is first used to establish the unmetalized propellant baseline. Nano-aluminum is then treated analytically using particle core-shell geometry, finite combustion completion, aluminum oxidation, condensed-phase heat capacity, and particle/gas velocity slip. Sensitivity studies are used to examine the effects of powder loading, particle size, oxide-shell thickness, burn-time scaling, heat coupling, and velocity slip. The goal is to identify performance trends rather than predict the performance of a specific engine.

Methods and Baseline Model

The baseline RocketCEA model is used to generate data about the baseline performance of Ethanol/Liquid Oxygen (LOX) bipropellant, using the RocketCEA Python package, which interfaces with NASA CEA thermochemical equilibrium methods. We began by inputting the chamber pressure, nozzle expansion ratio, and mixture ratio, the oxidizer to fuel mass ratio, (MR) sweep bounds that will also carry over to future sections of the study. The chamber

pressure is set to 15.0 bar, converted into psia to match RocketCEA, the chosen nozzle expansion ratio (A_e/A_t) is 40.0, and the mixture ratio was swept from 1.0 to 2.6 in increments of 0.05. With these values, the initial model yields combustion temperature, vacuum specific impulse, characteristic velocity, and specific heat ratio (γ). After that baseline is established, the nano-Al calculation is performed with explicit analytical equations. This structure keeps computational equilibrium chemistry separate from the assumptions used to approximate the metal additive. We now explain the analytical equations.

Theoretical Framework

The oxidizer to fuel mixture ratio is defined as

$$MR = \frac{m_{LOX}}{m_{ethanol}} \tag{1}$$

where m_{LOX} is the LOX mass and $m_{ethanol}$ is the ethanol mass. The reference liquid mixture ratio selected from the RocketCEA sweep is MR=1.85. At this condition $T_0 = 3259.05\ K$ and $I_{sp,0} = 342.87s$. These values define the unmetalized reference state. All later changes in temperature and specific impulse are measured relative to this baseline.

The nozzle expansion ratio is

$$\epsilon = \frac{A_e}{A_t}, \tag{2}$$

where A_e is nozzle exit area and A_t is the nozzle throat area. Here, we set $\epsilon = 40$. The total incoming propellant stream is normalized to one kilogram. Let f_p be the mass fraction of total passivated nano-Al powder. Then $m_{powder} = f_p$ and $m_{liquid} = 1 - f_p$. Because the remaining liquid retains the reference mixture ratio, $m_{ethanol} + m_{LOX} = 1 - f_p$ and $m_{LOX} = MR\, m_{ethanol}$. This gives:

$$m_{ethanol} = \frac{1-f_p}{1+MR} \tag{3}$$

and

$$m_{LOX} = \frac{MR\left(1-f_p\right)}{1+MR}. \tag{4}$$

A simplified combustion reaction for ethanol is

$$C_2H_5OH + 3O_2 \rightarrow 2CO_2 + 3H_2O \tag{5}$$

The corresponding oxygen-to-fuel mass ratio is

$$MR_{stoich} = \frac{3M_{O_2}}{M_{ethanol}} \approx 2.08378. \tag{6}$$

Passivated nano-Al is not modeled as pure metallic aluminum. Each particle is represented as a spherical metallic-Al core surrounded by a native Al2O3 shell. Let the particle diameter be given by d. If δ is the native oxide-shell thickness, then the core radius is $r_c = r - \delta$. Therefore, $V_{core} = \frac{4}{3}\pi r_c^3$ and $V_{shell} = \frac{4}{3}\pi\left(r^3 - r_c^3\right)$. Using $m = \rho V$, the fraction of total powder mass that is active metallic aluminum becomes

$$x_{active} = \frac{\rho_{Al} r_c^3}{\rho_{Al} r_c^3 + \rho_{Al_2O_3}\left(r^3 - r_c^3\right)}. \tag{7}$$

Substituting $r_c = r - \delta$ we get

$$x_{active} = \frac{\rho_{Al}(r-\delta)^3}{\rho_{Al}(r-\delta)^3 + \rho_{Al_2O_3}\left[r^3 - (r-\delta)^3\right]}. \tag{8}$$

The initial metallic-Al mass is $m_{Al,initial} = f_p x_{active}$. The native oxide initially carried by the powder is $m_{oxide,native} = f_p\left(1 - x_{active}\right)$. There is a time dependence to the burning of active aluminum. We model this behavior as:

$$x_{burn} = 1 - exp\left(-\frac{\tau_{res}}{\tau_{burn}}\right). \tag{9}$$

The mass of active aluminum that burns is $m_{Al,burned} = f_p x_{active} x_{burn}$, while the remaining aluminum is $m_{Al,unburned} = f_p x_{active}\left(1 - x_{burn}\right)$. Here we use $\tau_{res} = 10$ ms and $\tau_{burn} = 5$ ms, which gives$x_{burn} \approx 0.8647$. Aluminum oxidation is represented by

$$4Al + 3O_2 \rightarrow 2Al_2O_3 \tag{10}$$

From stoichiometry, the oxygen demand is approximately 0.8894637 kg O2 per kilogram of burned Al, so that $m_{O_2,demand} = 0.8894637\, m_{Al,burned}$. Therefore, the newly formed alumina is $m_{Al_2O_3,new} = 1.8894637\, m_{Al,burned}$. and the total alumina is $m_{oxide,total} = m_{oxide,native} + m_{Al_2O_3,new}$. We can then calculate the total condensed-particle mass fraction as $\chi_p = m_{oxide,total} + m_{Al,unburned}$, and the remaining gas mass fraction as $m_g = 1 - \chi_p$. Only active aluminum that burns contributes to the modeled chemical-energy release. The incremental released energy is calculated using

$$q_{release} = \eta_{heat} m_{Al,burned} Q_{Al} \tag{11}$$

where Q_{Al} is the aluminum oxidation energy scale and η_{heat} is a heat-coupling factor. Here, we use $Q_{Al} \approx 31.05$ MJ/kg active Al and $\eta_{heat} = 1$. The released aluminum energy must heat the gaseous products, alumina particles, and unburned Al. The total incremental heat capacity is

$$C_{total} = m_g C_{p,g} + m_{oxide,total} C_{p,Al_2O_3} + m_{Al,unburned} C_{p,Al} \quad (12)$$

We can now calculate the temperature increase fusing $q = C\Delta T$. Next we calculate the thermal performance effect using a square-root temperature scaling:

$$I_{sp,thermal} = I_{sp,0}\sqrt{T_1/T_0} \quad (13)$$

Therefore, we can calculate the thermal specific-impulse increment as

$$\Delta I_{sp,thermal} = I_{sp,thermal} - I_{sp,0}. \quad (14)$$

Finally, because condensed particles may leave the nozzle more slowly than the gas, we can define $s = \frac{v_p}{v_g}$, where v_p is particle exhaust velocity and v_g is gas exhaust velocity. We can define a mixture momentum factor as

$$P_{slip} = (1 - \chi_p) + \chi_p s = 1 - \chi_p(1 - s) \quad (15)$$

The momentum-adjusted specific impulse is $I_{sp,momentum} = P_{slip} I_{sp,thermal}$, and the corresponding increment is

$$\Delta I_{sp,momentum} = I_{sp,momentum} - I_{sp,0} \quad (16)$$

Results

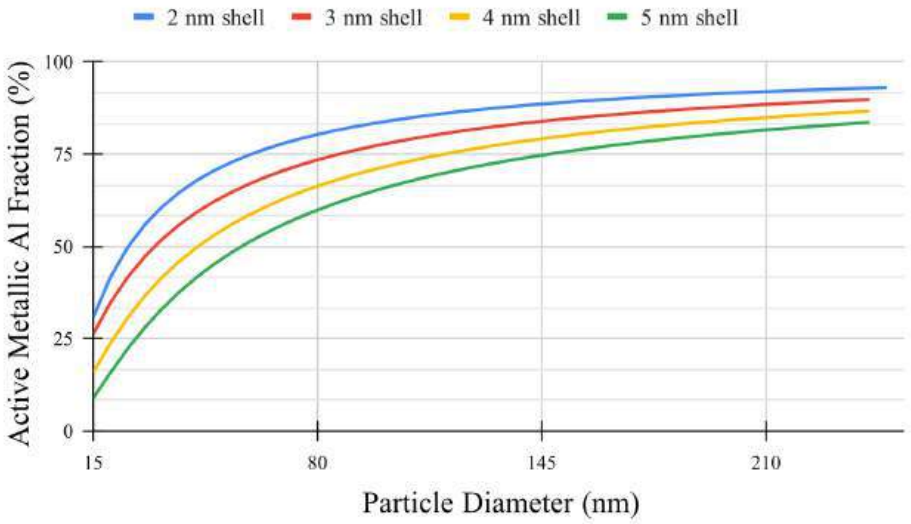

Figure 1. Active metallic aluminum fraction predicted by the spherical core-shell model for several native oxide-shell thicknesses.

Figure 1 shows that the active metallic-Al fraction increases with particle diameter for a fixed oxide-shell thickness. Smaller particles contain a greater oxide mass fraction because the

shell occupies a larger fraction of the total particle volume. For the reference 100 nm particle with a 3 nm shell, approximately 76.9% of the powder mass is active metallic Al.

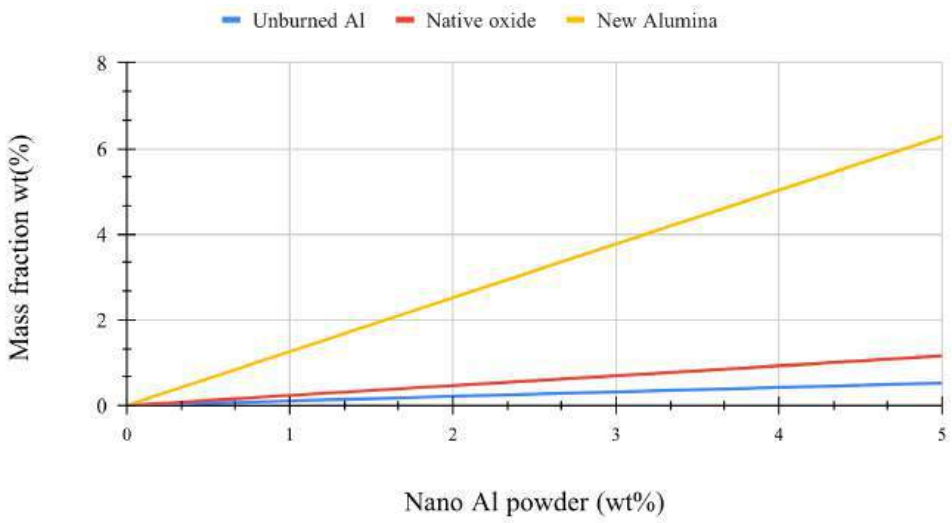

Figure 2. Condensed-particle mass partition across 0-5 wt% loading sweep.

Figure 2 plots mass fraction in wt% against nano aluminum powder also in wt%.For the 3 wt% reference case, the powder initially contains approximately 2.31 wt% active Al and 0.69 wt% native oxide. The burn-completion model consumes about 2.00 wt% Al and leaves about 0.31 wt% unburned Al. Formation of new alumina increases total Al2O3 to approximately 4.46 wt% of the one-kilogram stream. This shows that the condensed-product mass grows faster than the original aluminum loading because oxygen from the propellant becomes incorporated into newly formed alumina. Increasing powder loading therefore increases both available chemical energy and condensed mass.

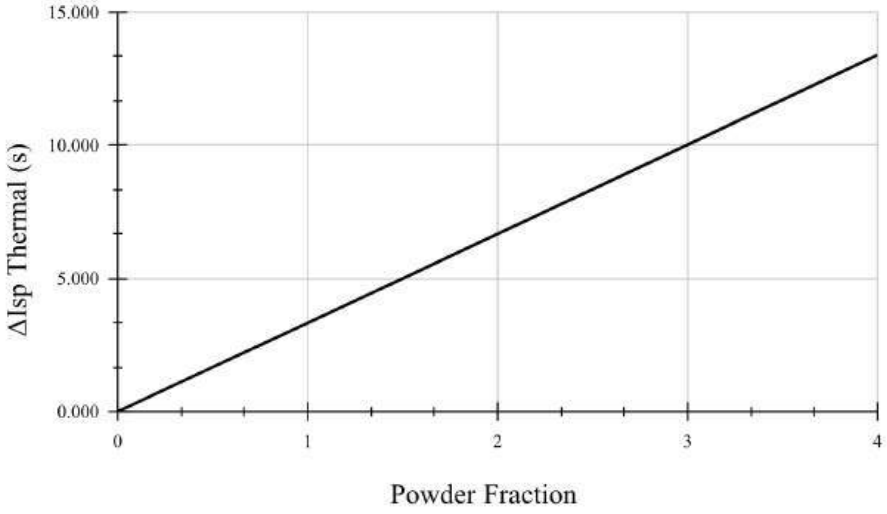

Figure 3. Temperature-only specific-impulse proxy versus passivated nano-Al powder loading.

Next we plot the thermal specific impulse (figure 3). The temperature-only specific-impulse proxy increases with powder loading, reaching approximately 10.02 s at 3 wt% and 16.78 s at 5 wt%. These values represent the thermal contribution before particle/gas momentum slip is applied. Because condensed particles may not accelerate to the same velocity as the gaseous exhaust, Figure 4 examines the effect of three assumed particle-to-gas velocity ratios on the momentum-adjusted specific-impulse increment. Figure 4 shows that lower particle velocity reduces the predicted performance gain.

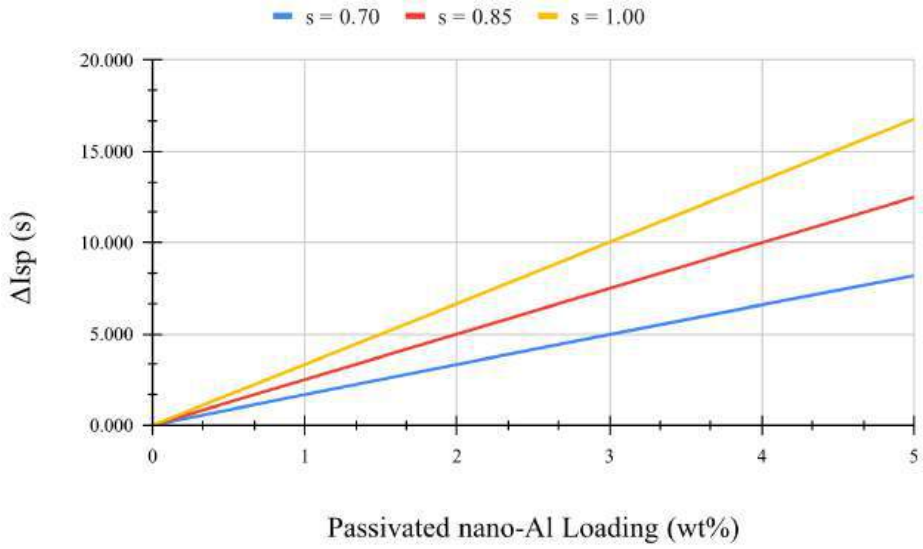

Figure 4. Momentum-adjusted specific-impulse proxy for three assumed particle/gas velocity ratios.

At 3 wt% powder:

$$s = 0.70 \rightarrow \Delta I_{sp} \approx 4.96\ s,$$

$$s = 0.85 \rightarrow \Delta I_{sp} \approx 7.49\ s,$$

and

$$s = 1.00 \rightarrow \Delta I_{sp} \approx 10.02\ s.$$

The spread between these cases shows that the assumed particle-to-gas velocity ratio has a strong effect on the predicted result. Particle size affects both the fraction of active metallic aluminum and the time available for combustion. Figure 5 therefore examines particle-size sensitivity using three assumed burn-time scaling exponents.

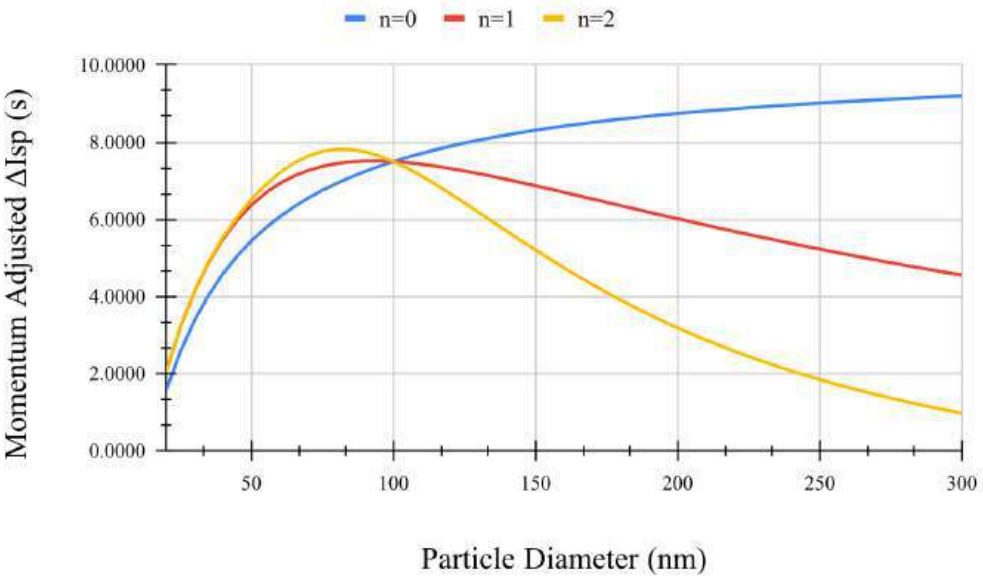

Figure 5. Particle-size sensitivity for three assumed burn-time exponents at 3 wt% loading.

When $n = 0$, burn time is held fixed and larger particles appear more beneficial because the oxide burden decreases. For 1, the model rises from about 6.36 s at 50 nm to 7.49 s at 100 nm and falls to about 6.01 s at 200 nm. For $n = 2$, the large-particle burn penalty is much stronger. These results show that the predicted effect of particle size depends on the assumed relationship between diameter and combustion time.The peaks in these curves should not be interpreted as proven optimum particle sizes because the values of n=0, 11, and 22 are sensitivity assumptions rather than experimentally calibrated combustion laws.

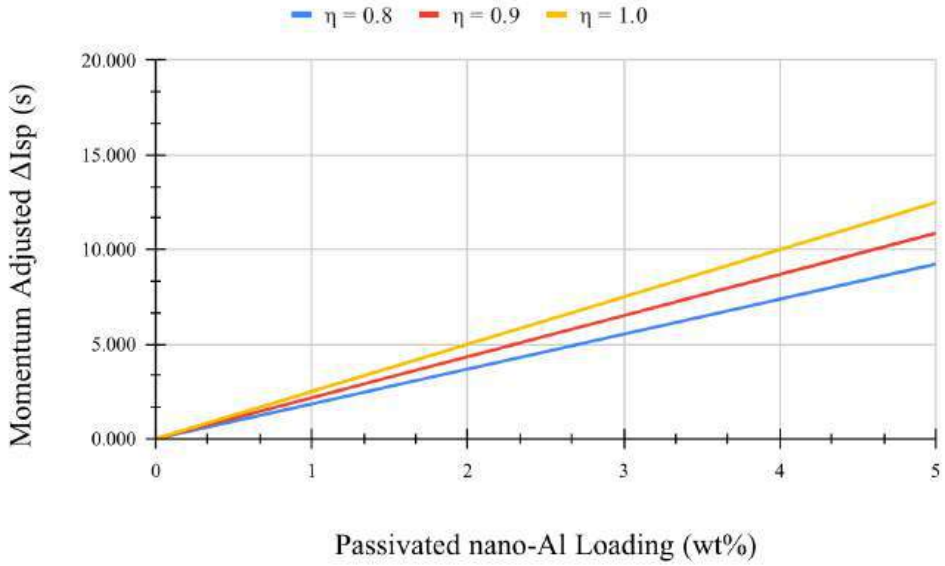

Figure 6. Post-release heat-coupling sensitivity for the 0-5 wt% loading sweep.

A final sensitivity analysis was performed to determine how the assumed heat-coupling factor, η_{heat}, affects predicted performance across the loading range. The results for η_{heat} = 0.8, 0.9, and 1.0 are shown in Figure 6. Figure 6 shows that reducing η_{ea} decreases the modeled performance gain at every loading fraction. At 3 wt% loading, the momentum-adjusted ΔI_{sp} is approximately 5.53 s for η_{heat} = 0.8, 6.51 s for η_{ea} = 0.9, and 7.49 s for η_{heat} = 1.0. Because the particle mass distribution is unchanged in this sensitivity study, the difference between the curves results from the fraction of released aluminum chemical energy that contributes to the modeled temperature increase.

Discussion and Limitations

The main result of this study is that increasing passivated nano-aluminum loading produces a positive modeled ΔT and ΔI_{sp} under the reference assumptions. Across the 0–5 wt% range, increasing powder loading increases the amount of active aluminum available to burn, which increases the modeled chemical-energy release. However, increasing loading also increases the amount of condensed Al_2O_3 and unburned Al in the exhaust. Because of this, the momentum-adjusted results are lower than the temperature-only estimates when particle velocity slip is included. Overall, the model shows a tradeoff between the additional energy provided by aluminum and the added condensed-particle mass produced during combustion.

Particle size also has a noticeable effect on the modeled performance. At a fixed oxide-shell thickness, smaller particles contain a larger oxide mass fraction because the shell takes up a greater fraction of the total particle volume. This reduces the amount of active metallic aluminum available to burn. However, smaller particles may also burn faster, so the overall particle-size trend depends on the assumed relationship between particle diameter and burn time. The burn-time sensitivity cases show that changing this relationship can shift the predicted performance considerably. For this reason, the results do not show that a certain particle size is truly optimal.

The heat-coupling and velocity-slip sensitivity studies also show that the predicted gains depend strongly on model assumptions. Lower values of η_{heat} reduce the amount of aluminum chemical energy that contributes to the modeled temperature increase, while lower values of the particle-to-gas velocity ratio, s, reduce the momentum contribution of the condensed particles. These effects show that a temperature increase alone does not fully describe the performance of a metalized exhaust. Even when the same amount of aluminum burns, the final ΔI_{sp} can change depending on how effectively the released energy heats the mixture and how closely the particles follow the gas velocity.

This study also experiences several limitations because of the reduced-order model and the constant parameters used. The calculated ΔT is an incremental post-baseline estimate rather than a full adiabatic flame-temperature calculation. The model begins with the RocketCEA Ethanol/LOX chamber temperature and then adds the estimated aluminum energy release

through a separate analytical energy balance. Because of this, it does not recompute a fully self-consistent equilibrium state after some of the liquid propellant is replaced by nano-aluminum. Aluminum oxidation is also not directly coupled to the Ethanol/LOX equilibrium calculation. The model calculates the O2 demand of burned aluminum separately, while a real fuel-rich chamber distributes oxygen among C, H, O, Al, and other high-temperature species.

Other physical effects are also simplified or not included. RocketCEA assumes equilibrium chemistry and does not directly model agglomeration, injector mixing, particle dispersion, or two-phase nozzle flow. The residence time and burn time are treated as assumed effective values, and the particle-to-gas velocity ratio is also imposed rather than calculated from a detailed flow model. In addition, the original particle diameter may not be the same as the final agglomerate size leaving the nozzle. These effects could change the actual performance of a physical engine.

A strength of the reduced-order approach is that these assumptions can be changed independently, which makes it useful for sensitivity analysis and for identifying which variables have the greatest effect on the result. The model therefore has more value for identifying general physical trends than for predicting exact engine performance. The qualitative trends are more reliable than the exact values of ΔTor ΔI_{sp} . Future studies should include coupled C/H/O/Al equilibrium chemistry, higher-fidelity combustion kinetics, two-phase nozzle modeling, CFD-based mixing and residence-time analysis, and experimental validation to determine whether the modeled performance gains can be achieved in a real propulsion system.

Conclusion

This study developed a reduced-order sensitivity model to evaluate how nano-aluminum additives may affect the performance of an Ethanol/LOX rocket propellant. RocketCEA was first used to establish the baseline performance of the unmetalized propellant, and the nano-aluminum effects were then modeled analytically using factors such as particle geometry, finite burn completion, aluminum oxidation, heat capacity, and particle velocity slip. Overall, the results showed that increasing nano-aluminum loading from 0–5 wt% increased the modeled ΔT and ΔI_{sp} under the reference assumptions. At 3 wt% loading, the model produced a ΔT of about 193.3 K, a thermal $\Delta I_{sp}$of 10.02 s, and a momentum-adjusted ΔI_{sp} of 7.49 s.The sensitivity studies also showed that the final performance depends heavily on the assumptions used in the model. Particle size, oxide-shell thickness, burn time, heat coupling, and the particle-to-gas velocity ratio all affected the predicted ΔI_{sp}. For example, smaller particles contain a greater relative amount of oxide, while larger particles may experience slower combustion depending on the assumed burn-time relationship. Lower heat coupling and greater particle velocity slip also reduced the modeled performance gain. However, these results should not be interpreted as proving an optimal loading fraction or particle size. The model does not fully account for factors such as coupled oxygen allocation, agglomeration, detailed injector mixing, or two-phase nozzle

flow. Therefore, the results are better understood as general trends and sensitivity estimates rather than final engine-performance predictions. Future studies should include more complete C/H/O/Al equilibrium chemistry, higher-fidelity two-phase flow modeling, CFD-based mixing and residence-time analysis, and experimental testing to determine whether these modeled gains can be achieved in a physical rocket engine. Comparing nano-Al with other nano-additives could also prove beneficial to improving rocket performance [9].

Acknowledgements

I thank Sammy Hassan and Kiranjyot Gill for peer-reviewing the paper.

Works Cited

1.NASA. Ideal rocket equation | Glenn Research Center | NASA [Internet]. Glenn Research Center | NASA. 2023 [cited 2026 Sept 17]. Available from: https://www1.grc.nasa.gov/beginners-guide-to-aeronautics/ideal-rocket-equation/

2.Yuan J, Liu J, Zhou Y, Wang J, Xu T. Aluminum agglomeration of AP/HTPB composite propellant. Acta Astronautica [Internet]. 2019 Mar 1 [cited 2026 Sept 17];156:14–22. Available from: https://www.sciencedirect.com/science/article/pii/S0094576518315492

3.DeLuca LT. Overview of Al-based nanoenergetic ingredients for solid rocket propulsion. Defence Technology. 2018 Oct;14(5):357–65.

4.Jayaraman K, Chakravarthy, Sarathy R. Behaviour of nano-aluminum in solid propellant combustion. 44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit [Internet]. 2008 July 21 [cited 2026 Sept 17];. Available from: https://www.researchgate.net/publication/268481444_Behaviour_of_Nano-Aluminum_in_Solid_Propellant_Combustion

5.Galfetti L, DeLuca, L.T, Severini F, Colombo G, Meda L, Marra G. Pre and post-burning analysis of nano-aluminized solid rocket propellants. Propellant combustion phenomena [Internet]. 2007 Jan 1 [cited 2026 Sept 17];11(1):26–32. Available from: https://www.sciencedirect.com/science/article/pii/S1270963806001349

6.Wang F, Wu Z, Shangguan X, Sun Y, Feng J, Li Z, et al. Preparation of mono-dispersed, high energy release, core/shell structure AL nanopowders and their application in HTPB propellant as combustion enhancers. Sci Rep [Internet]. 2017 July 12 [cited 2026 Sept 17];7(1):5228. Available from: https://pubmed.ncbi.nlm.nih.gov/28701741/

7.Jones M D, Li CH, Afjeh AA, Peterson GP. Experimental study of combustion characteristics of nanoscale metal and metal oxide additives in biofuel (ethanol). Nanoscale Research Letters. 2011 Mar 22;.

8.Liang D, Ke R, Wu Z, Jiang Y, Shen D, Li H, et al. Combustion characteristics of oxygenated slurry droplets of nano-Al/EtOH and nano-Al/TPGME blends. Energy [Internet]. 2021 Apr 1 [cited 2026 Sept 17];220:119693. Available from: https://www.sciencedirect.com/science/article/pii/S0360544220328000

9.Sim HS, Plascencia MA, Vargas A, Bennewitz JW, Smith OI, Karagozian AR. Effects of inert and energetic nanoparticles on burning liquid ethanol droplets. Combustion Science and Technology. 2018 Aug 21;191(7):1079–100.

10.Shi E, Yan X, Qi C, Lei C, Yu J. Study on the combustion characteristics and reaction mechanism of aluminum/alcohol-based nanofluid fuel. Combustion and Flame [Internet]. 2024 Nov 1 [cited 2026 Sept 17];269:113684. Available from: https://www.sciencedirect.com/science/article/pii/S0010218024003936

Perceptions of Learning Difficulties Across Cultures and Generations: A Cross Sectional Survey in Singapore and Indonesia By Michael Vincent Junaidi

Abstract

Learning difficulties are now increasingly understood as neurodevelopmental differences rather than simply a lack of effort, ability or intelligence. Yet public understanding and perceptions remain uneven across societies. This cross sectional survey aims to compare the difference in perceptions of learning difficulties between Singapore and Indonesia across different age groups. The survey has collected information from a total of 117 respondents, with 71 from Indonesia and 46 from Singapore across 3 different age groups. Likert scale questions on perceptions of learning difficulties in four domains (i.e., personal awareness, friends/family, school, and country/society) were used to obtain quantitative data and were combined with open-ended questions to provide qualitative aspects to the analysis. The results were then analyzed and compared to highlight and explain the differences in perception. Respondents from Singapore had higher overall learning difficulty perception scores than respondents from Indonesia, although differences varied across domains.. Respondents under the age of 19 had the highest average perception scores among the three age groups, although age group was not statistically significant in ANCOVA analyses. ANCOVA identified country as a significant predictor of perception sum scores in the country/society domain, while age group was not a significant predictor. The open-ended responses also showed that respondents most commonly attributed changes in perception to social media and public visibility, followed by advances in science and education.

Introduction

Learning difficulties are neurological differences that affect the brain and how a person processes, gains, stores, and recalls information (Peterson and Pennington, 2015; Butterworth et al., 2011). These conditions may not only affect how a person, even with average or above-average intelligence, performs in academic subjects but also in their daily lives (Landerl, Bevan, and Butterworth, 2004; Sharfi and Rosenblum, 2014). Learning difficulties impact many areas such as memory, attention, reading, writing, problem solving and organization (Smith-Spark et al., 2016; Willcutt et al., 2005). There are many types of learning difficulties such as ADHD, dyscalculia and dyslexia, each one impacting a different area (Peterson and Pennington, 2015; Butterworth et al., 2011; Willcutt et al., 2005; Landerl et al., 2009). One crucial aspect about learning difficulties is that they are not a sign that a person lacks intelligence (Landerl, Bevan, and Butterworth, 2004). However, people that have them require support to be successful in their lives (Swanson and Hoskyn, 1998; Dobson Waters and Torgerson, 2021).

Over the past several decades, there have been significant improvements in the awareness and level of general education about learning difficulties (Gelber, 2025). This may be due to several factors. Firstly, scientific developments have improved significantly. Over the past few decades, medicine, psychology, and education have started to pay more attention to a broader

spectrum of cognitive variations. This has led to a better understanding surrounding the issue (Peterson and Pennington, 2015; Butterworth et al., 2011). Secondly, the presence of social media has also allowed people to gain more information through both shared experiences and facts, unlike people in the past who mostly relied on doctors, teachers, and textbooks they came across (Zhao and Zhang, 2017). Thirdly, modern research on parenting has increasingly emphasized autonomy support and emotional understanding in parent-child relationships (McCurdy et al., 2020). Finally, as schools increasingly use online and digital learning environments, students with learning difficulties may face additional challenges related to accessibility, engagement, and learning support (Petretto et al., 2021). All of these factors may have contributed to reducing some stigma surrounding learning difficulties and increasing awareness surrounding this issue (Haft et al., 2023).

The way that people with learning difficulties are perceived, viewed, and handled by society, educators, and experts varies across different areas of the globe (Haft et al., 2023; Ahmed et al., 2022). This may be due to a number of factors. Firstly, it may be due to the culture of the specific country (Haft et al., 2023). For example, in countries with high-stakes exam cultures, if a person with a learning difficulty does not do well academically, they might be labeled as lazy, not giving enough effort, or not having enough discipline (Haft et al., 2023). In some communities with more traditional perspectives, learning difficulties might be misunderstood instead of being seen as neurodevelopmental conditions. This can lead to the child not being able to receive proper support, as the family might try to hide the child's struggles out of shame (Haft et al., 2023; Katz et al., 2022). Variation in perception of learning difficulties in different areas might also be due to the economic conditions of the specific place. Lower-income or developing countries might lack the necessary resources to support these conditions (Ahmed et al., 2022). For example, educators might not have the necessary training to handle conditions like dyslexia and dyscalculia (Ahmed et al., 2022). This is contrasted with higher-income countries, which may have stronger infrastructure and greater resources for early intervention of these conditions (Ahmed et al., 2022).

Singapore is a small city-state in Southeast Asia with a highly developed economy (Huff, 1995). Academically, Singapore has a highly rigorous and high-performing education system, with major national examinations such as the PSLE and later national examinations (Deng and Gopinathan, 2016; Ma, 2025). This places immense pressure on students, which can cause students with learning difficulties to suffer and struggle in their academic journey (Chua et al., 2026; Jamaludin et al., 2024). In terms of support for people with learning difficulties, Singapore has a structured system of special educational needs and inclusive education support (Thaver and Lim, 2014). Additionally, support for students with special educational needs has increasingly been incorporated into mainstream schools (Thaver and Lim, 2014).

In contrast to Singapore's highly developed economy, Indonesia has experienced economic growth but continues to face significant income inequality and uneven access to resources (Handayani and Hanifa, 2024). Indonesia's academic system is diverse across regions and school types, with differences in school resources and access to support (Wirandana and

Khoirunurrofik, 2024). However, disparities between schools can be significant, with differences in funding, infrastructure, and educational resources (James et al., 1996; Bedi and Garg, 2000). The government has attempted to provide support to people with disabilities and learning needs through inclusive schools and dedicated programs. However, implementation can be inconsistent, and inclusive education continues to face challenges (Andriana and Evans, 2021; Ahmed et al., 2022). Even with systems in place, there are still many challenges in Indonesia regarding the support of people with learning difficulties due to differences in resources, teacher training, and access to inclusive education (Ahmed et al., 2022; Andriana and Evans, 2021). Although the perceptions of learning difficulties have improved over time, there has been limited direct comparison between the levels of awareness in Indonesia and Singapore. This study aims to compare the levels of awareness and understanding between these two countries in relation to different age groups. So, the main research question is: How do perceptions of learning difficulties differ between Singaporean and Indonesian respondents, and how do these perceptions vary across age groups? After doing research about the context of both countries and the evolution of perception levels of learning difficulties over the past years, it is hypothesized that Singaporean respondents and respondents from younger age groups will show a generally greater level of awareness and understanding compared to Indonesian respondents and older age groups.

Methodology

1. Population & Recruitment

A convenience sampling method was employed. Participants were recruited through the researcher's existing social networks, including friends, acquaintances, and family members in both Singapore and Indonesia. The survey was distributed directly through messaging platforms and social media, as well as in person where possible. This recruitment strategy enabled efficient access to multiple age groups across both national contexts.

The study population consisted of members of the general public from Singapore and Indonesia across multiple age groups. Participants were categorized into the following age brackets: (1) 15 years and under, (2) 16–19 years, (3) 20–25 years, (4) 26–35 years, (5) 36–45 years, (6) 46–55 years, (7) 56 years and above. This categorization allowed for generational comparison within and between both countries. To improve the reliability of our results, we have decided to group the ages into 3 groups: (1) under 19 years and under, (2) 20–45 years, and (3) 46 years and above. This made it easier to compare the age groups with each other by making group sizes more comparable.

2. Research Design & Tools

This study adopted a mixed-methods, cross-sectional comparative design. It aimed to compare awareness of learning difficulties across four domains: personal knowledge, school, family, and country. Data were collected using a structured questionnaire created using Google Forms. Participants under the age of 15 were required to have permission from their parents/legal

guardians before answering the form. All participants were required to provide their consent before completing the survey. The survey consisted of the following:

- Demographic Information (20 questions)
- Likert scale questions (15 questions) with 5 options: Strongly Disagree (-2), Disagree (-1), Neutral (0), Agree (1), Strongly Agree (2)
- Time-based Likert scale questions (5 questions) with 5 options: Always (2), Often (1), Sometimes (0), Rarely (-1), Never (-2)
- Open-ended questions (2 questions)

The 20 Likert scale questions were divided into four domains: Personal Knowledge (7 questions), School (3 questions), Family (3 questions), and Society & Country (7 questions). For each subject, a sum score was calculated for each domain as the sum of the questions assigned to a given domain (further details below). Since each domain has a different number of questions, they each have different minimum and maximum values and therefore average sum scores are not comparable between the domains. Instead, comparisons were made within the same domains such as between Indonesia or Singapore or the different age groups.

The survey concluded with two open-ended questions allowing participants to elaborate on their views or experiences regarding learning difficulties. This qualitative component was included to capture nuanced perspectives beyond scaled responses. These two questions were: (1) In your opinion, what has caused changes in public perception of learning difficulties over the years? And (2) Is there anything else you'd like to share about your views on learning difficulties and how they are perceived?

The survey was originally developed in English. For Indonesian participants, the questionnaire was translated into Bahasa Indonesia to ensure accessibility. Responses provided in Bahasa Indonesia were subsequently translated back into English for analysis.

3. Preprocessing & Statistical Analysis

The data collected was compiled and analyzed in a Google Sheets spreadsheet. For the Likert scale questions, each response option was converted into a numerical value to allow the results to be analyzed quantitatively. A score of 0 represented a neutral response, while higher or positive values represented responses that indicated a higher level of awareness, discussion, or recognition of learning difficulties. Lower or negative values represented responses that indicated a lower level of awareness, discussion, or recognition. The numerical values ranged from -2 to 2. However, the specific response options that were assigned positive or negative values varied depending on the wording of each question, as some questions were phrased differently. For example, for the question “Please rate your level of agreement with each statement. [I have a clear understanding of what learning difficulties are.]” positive values were assigned to those who agreed with the statement as it would indicate higher levels of awareness. However, for the question “Please rate your level of agreement with each statement. [Learning

difficulties limit what an individual is capable of.],” negative values were assigned to those who agree with the statement as it would indicate lower levels of awareness.

The converted values were then summed up to create the total sum score within each domain to create sum scores for the four domains (i.e., school, family/friends, country and personal). Then, pivot tables were created based on country and age group using the values for the average total sum score of a respondent. With the help of bar charts, the sum scores were compared and the age group or country with higher sum score has greater levels of awareness about learning difficulties. These figures were generated using Google Sheets.

For the open ended questions, the responses were categorized into the following categories: (1) Increased Awareness & Information Access, (2) Advances in Science and Education, (3) Social Media & Public Visibility, (4) Changing Social Attitudes, (5) Others. The number of responses for each country were then counted and converted into the percentage of answers about each topic per country to analyze the trends in perceptions of learning difficulty for each country. Individual responses were also read to provide more specific analysis.

The questionnaire also included the questions “How many hours per day do you use social media?” and “Do you personally know anyone with a learning difficulty?" to further explore other factors affecting awareness.

ANCOVA analyses were used to examine whether country, age group, social media use, and knowing someone with a learning difficulty were related to perception scores. Each of the four domain sum scores (Friends and Family, School, Personal, and Country/Society) were analyzed separately. Country and age group were included as categorical variables, and the models also included the interaction between country and age group to examine whether age-related differences varied between the two countries. Additional covariates included: whether participants knew someone with a learning difficulty and daily social media use in hours per day (0, 0.5, 1.5, 3.5, or 6 hours). ANCOVA analyses were conducted in Google Colab using Python 3.12..

Results

After gathering all the data, the final analytical sample included slightly more respondents from Indonesia compared to Singapore (i.e., 71 Indonesian and 46 Singaporean respondents). In total, across both countries, the 20–45 age group had the greatest number of respondents, while the 46 and above age group had the least (**Table 1**). This shows that the sample is not evenly distributed across countries and age groups, which should be considered when interpreting the results.

Country	Age	Number
Indonesia	46 and above	14
	20-45	36

	19 and under	21
Indonesia total		71
Singapore	46 and above	17
	20-45	16
	19 and under	13
Singapore total		46
Grand total		117

Table 1: Distribution of Respondents by Country and Age Group

Country	**Age**	**Average of Friends and Family Sum Score**	**Average of Personal Sum Score**	**Average of School Sum Score**	**Average of Country Sum Score**	**Total Sum Score**
Indonesia	46 and above	-1.44	4.11	0.44	0.53	3.64
	20-45	-1.57	3.93	-0.07	0.64	2.93
	19 and under	-1.43	4.76	0.38	0.24	3.95
Indonesia Total		-1.46	4.27	0.32	0.46	3.59
Singapore	46 and above	-1.69	5.00	0.06	1.56	4.94
	20-45	-2.06	5.65	0.06	2.06	5.71
	19 and under	-1.62	5.31	0.46	1.69	5.85
Singapore Total		-1.80	5.33	0.17	1.78	5.48
Grand Total		-1.60	4.68	0.26	0.98	4.33

Table 2: Mean Sum Scores by Country and Age Group

Across the four domains analyzed, the personal average sum score is the greatest, followed by country, then school and finally, friends and family. In terms of school, both countries still have positive values. The average of friend and family sum scores, however, has a negative value (**Table 2**).

In terms of age groups, individuals under 19 generally had the greatest total sum score, followed by respondents aged above 46, then by respondents aged 20–45. While this is true for respondents under 19 in both countries, in Indonesia, respondents from 20–45 have greater average sum scores compared to those above 46 (**Table 2**).

Comparing both countries, it can be seen that generally, Singapore has a higher average sum score in the categories of personal awareness and country (**Table 2**). On the other hand,

Indonesia has higher average sum scores in the categories for friends and family and also school (**Table 2**).

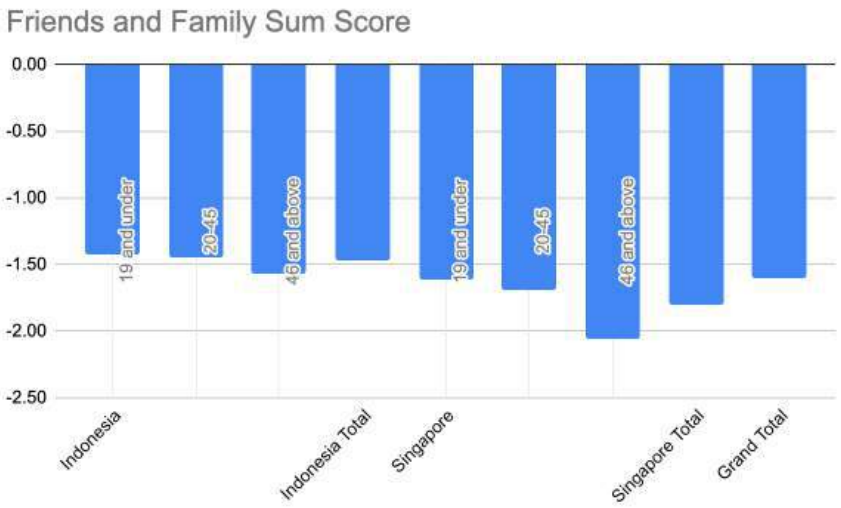

Fig 1: Mean Friends and Family Sum Scores by Country and Age Group

Respondents from both Indonesia and Singapore from all age groups had negative group averages of family and friend sum scores (**Fig 1**). Overall, Singapore had a greater average family and friend sum score compared to Indonesia. The age group with the highest average sum score in both countries was under 19, followed by 20-45 and 46+ with the least average sum score.

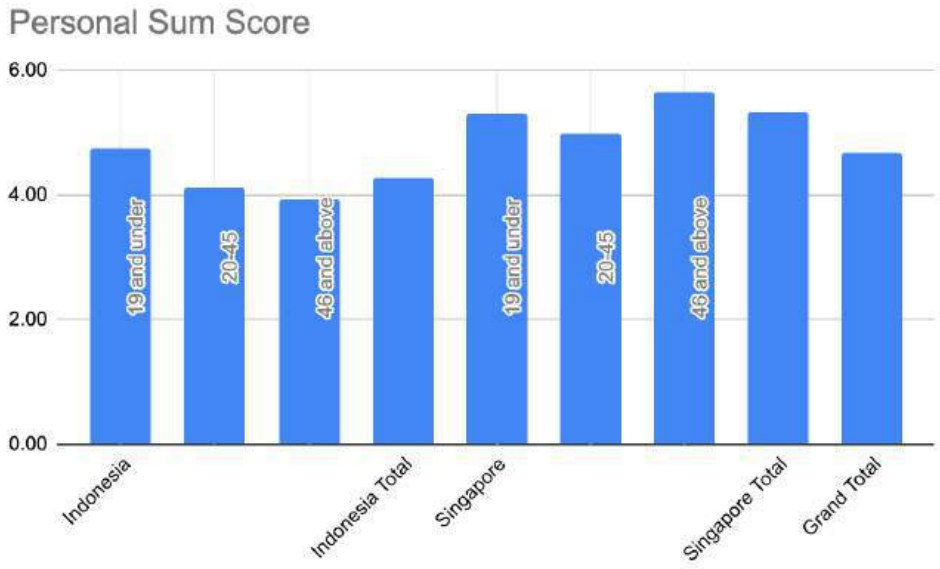

Fig 2: Mean Personal Knowledge Sum Scores by Country and Age Group

When comparing ages in the personal sum score bar chart, participants above the age of 46 had the greatest average in Singapore compared to other age groups (**Fig 2**). This was followed by participants under the age of 19 then between 20–45. In Indonesia however, participants aged under 19 had the greatest average, followed by participants aged between 20–45 then over the age of 46.

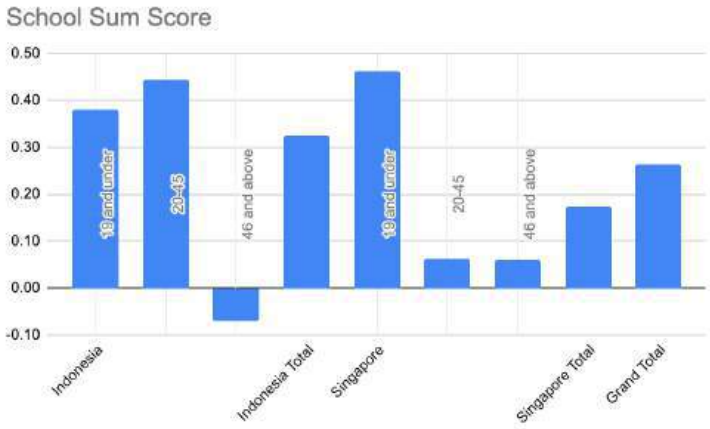

Fig 3: Mean School Sum Scores by Country and Age Group

When comparing school sum scores between the two countries, Singaporean respondents belonging to the age groups of under 19 and above 46 had higher average school sum scores compared to Indonesian respondents (**Fig 3**). However, Indonesians in the 20–45 age group had noticeably higher school sum scores compared to Singaporeans of that age group. This difference led Indonesia to have a higher overall school sum score compared to Singapore. The interactions between country and age in other domains, however, were more consistent in both countries. In other words, for domains other than school, if the average sum score was higher in one country than the other, the individual age groups also had higher sum scores for that country.

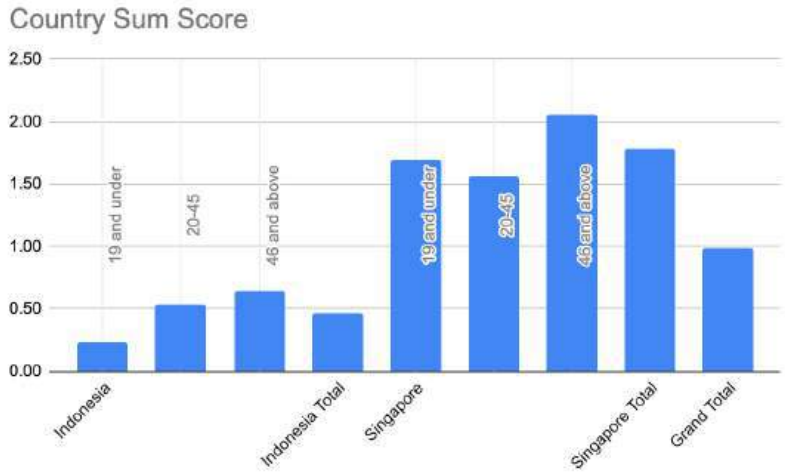

Fig 4: Mean Society and Country Sum Scores by Country and Age Group

Comparing society & country sum scores of both countries, Singapore had a noticeably higher average scores across all ages compared to Indonesia (**Fig 4**). In terms of age groups, respondents above the age of 45 had the greatest average sum score, followed by respondents under 19, and finally respondents between the ages of 20–45 in Singapore. In Indonesia, respondents above the age of 45 also had the greatest average score, followed by respondents between the ages of 20–45, and finally respondents under the age of 19.

	F-Statistic	P-Value	Degrees of Freedom
Know Someone With Learning Difficulties	3.545	0.062	1.0
Social Media	1.128	0.351	6.0
Country	0.001	0.981	1.0
Age Group	0.502	0.607	2.0
Country:Age Group	0.221	0.802	2.0

Table 3: ANCOVA Results for Friends and Family Sum Score

	F-Statistic	P-Value	Degrees of Freedom
Know Someone With Learning Difficulties	0.003	0.954	1.0
Social Media	1.323	0.252	6.0
Country	0.277	0.599	1.0
Age Group	3.060	0.051	2.0
Country:Age Group	0.778	0.462	2.0

Table 4: ANCOVA Results for School Sum Score

	F-Statistic	P-Value	Degrees of Freedom
Know Someone With Learning Difficulties	1.995	0.954	1.0
Social Media	0.550	0.769	6.0
Country	0.982	0.323	1.0
Age Group	0.086	0.918	2.0
Country:Age Group	0.249	0.780	2.0

Table 5: ANCOVA Results for Personal Sum Score

	F-Statistic	P-Value	Degrees of Freedom

Know Someone With Learning Difficulties	1.836	0.178	1.0
Social Media	1.878	0.091	6.0
Country	9.037	0.003	1.0
Age Group	0.315	0.731	2.0
Country:Age Group	0.145	0.864	2.0

Table 6: ANCOVA Results for Country Sum Score

For the domain of friends and family, knowing someone with a learning difficulty approached significance although it did not reach the threshold of $P<0.05$ (**Table 3**). Similarly, for the domain of school, being of a different age group also approached significance although it did not reach the threshold of 0.05 as well (**Table 4**). For the domain of personal knowledge, no values approached significance (**Table 5**). However, knowing someone with a learning difficulty has the highest F-value compared to other factors, which may indicate a potentially meaningful relationship. Finally, for the domain of Country, the country in which a respondent is from significantly affected the result with a $P=0.003$ and an $F=9.03$ (**Table 6**).

Open-ended questions

Topic	Indonesia (%)	Singapore (%)	Grand Total (%)
Advances in Science & Education	24.1	36.4	28.5
Changing Social Attitudes	12.7	18.2	14.6
Social Media & Public Visibility	40.5	29.5	36.6
Increased Awareness & Information Access	22.8	15.9	20.3

Table 7: Frequency and Percentage of Open-Ended Response Themes by Country

For the final two open-ended questions, we asked respondents to tell us about their thoughts on what caused changes in public and personal opinion on learning difficulties and if they would care to share any further insights. We then categorized these responses into the above four topics, with some responses being categorized into more than one topic: Social Media and Public Visibility, Advances in Science and Education, Increased Awareness and Information Access, and Changing Social Attitudes. Overall, Social Media and Public Visibility were the most common topics with 45 responses (**Table 7**). This was followed by Advances in Science and Education with 35 responses, Increased Awareness and Information Access with 25 responses, and Changing Social Attitudes with 18 total responses. For example, one Indonesian respondent stated that “Social media has helped society be more supportive towards people who

have learning disabilities." In Indonesia specifically, Social Media and Public Visibility was also the most common topic, while most Singaporean respondents attributed the change in perception to Advances in Science and Education instead, with 16 responses. One respondent from Indonesia also stated that the changes were "Because of increasingly advanced technology."

Discussion

In line with the hypothesis, respondents from Singapore had better perceptions and understanding about learning difficulties compared to respondents from Indonesia, although differences varied across the four domains. This is because it has a higher total average sum score compared to Indonesia. However, the results differ with the hypothesis in terms of the age groups. The hypothesis predicts that as the age decreases, the perceptions will improve. While respondents under 19 have the highest total average sum score, respondents aged 46 and above have greater total average sum score compared to respondents aged 20–45. Although respondents under 19 had the highest average scores, respondents aged 46 and above had greater total average sum scores than those aged 20–45. However, age group was not a statistically significant predictor of perception scores.

The findings from the open-ended questions attribute the changes in perceptions of learning difficulties over time to the increase in use of social media and public visibility. Over the past few decades, social media has grown into one of the world's most popular modes of communication and exchanging information (Zhao and Zhang, 2017). Unlike in the past, social media has allowed more information about learning difficulties to be more easily spread to more people (Zhao and Zhang, 2017). The development of science and technology in recent years has also been very relevant to improving perceptions about learning difficulties, as research in psychology, neuroscience, and education has improved understanding of conditions such as dyslexia, dyscalculia, and ADHD (Peterson and Pennington, 2015; Butterworth et al., 2011; Willcutt et al., 2005). People have improved their understanding through scientific fact and evidence, which has contributed to a decrease in social stigma and negative perceptions (Haft et al., 2023). This has also led education systems to change into more inclusive ones, considering the needs of people with learning difficulties (Ahmed et al., 2022; Thaver and Lim, 2014). One respondent responded in the open-ended section saying that "Public perception of learning difficulties has changed due to better scientific understanding, advocacy and awareness campaigns, stronger legal protections in education, and the influence of role models who openly share their experiences. These factors have shifted attitudes from stigma and deficit to acceptance and inclusion."

The under 19 age group had the highest average score in both countries. This means that they generally have better understanding and awareness about learning difficulties compared to other age groups. However, because age group was not a statistically significant predictor in ANCOVA models, this is a purely descriptive difference in this sample. This difference may be due to, just as discussed above, the increase in use of social media. One of the questions in the demographic section of the research paper was "How many hours per day do you use social media?" Respondents aged under 19 chose 3–4 hours per day as the most popular option, with

58.8% selecting this response. This is more than the age group of 20–45, where the most popular option was also 3–4 hours per day at 44.2%. However, this contrasts with respondents above 46, where the most popular option was 1–2 hours per day, with 41.9% selecting this response. This shows how an increase in exposure to information about learning difficulties through increasing use of social media may lead to respondents under 19 having the best perceptions about learning difficulties compared to other age groups (Zhao and Zhang, 2017). This is because social media may expose respondents to more information about learning difficulties, leading to more awareness and understanding about the issue (Zhao and Zhang, 2017). However, some respondents also mention that social media might cause learning difficulties instead of leading to awareness. One respondent said, "The rapid development of social media and gadgets has brought about changes in social interactions, particularly among children; many students now experience learning difficulties as a result of spending excessive time using these devices." Therefore, in this age of social media, it is crucial to maintain a balance between using it to receive its benefits while also avoiding its downsides.

Singapore has overall better perceptions of learning difficulties - particularly in areas of society and personal knowledge - than Indonesia as it has a higher average sum score. This may be due to increasing efforts in Singapore to promote a more inclusive education system for students with special educational needs (Thaver and Lim, 2014). Additionally, Singapore has a stronger economy, which may contribute to how learning difficulties are perceived and supported (Huff, 1995). More funds and resources may be available for medical and psychological consultations, allowing children with learning difficulties to receive more professional support and increasing formal discussions about this issue with both children and their parents. This is also reflected in the sample as both respondents with a learning difficulty that answered the survey were from Singapore. This could help explain why Singaporean respondents generally showed higher awareness and understanding of learning difficulties compared to Indonesian respondents.

Indonesia, however, had better perceptions in the domain of school and family and friends. This might be due to the educational system present and the amount of stigma present in both countries. Compared to Singapore's highly competitive education system, Indonesia may have a more relaxed one in some contexts. Singapore's high-performing and competitive education system may create greater academic pressure on students (Deng and Gopinathan, 2016; Chua et al., 2026). In high-pressure academic environments, stigma surrounding learning difficulties may also make students and families less willing to openly discuss these difficulties (Haft et al., 2023; Katz et al., 2022). This might be why Indonesia has higher perceptions towards learning difficulties in these two areas. However, this finding should be interpreted carefully, as Indonesia's education system still faces inequalities in resources, infrastructure, and access to support across different regions and school types (Wirandana and Khoirunurrofik, 2024; James et al., 1996; Bedi and Garg, 2000). It is also worth noting that all age groups and countries had averages of family and friend sum scores meaning that both countries and all age

groups don't have frequent conversations and discussions about learning difficulties between their friends and family.

From the results of the ANCOVA test, it is found that the country in which a respondent is from significantly affected the Country sum score. One possible explanation is differences in country's access to information, public awareness and discussion, and educational systems. For the domain of friends and family, knowing someone with a learning difficulty approaches significance. This may potentially be caused by someone in a respondent circle such as a family member or friend to have a learning difficulty, leading the respondent to have better understanding about learning difficulties. For the domain of school, being of a different age group also approached significance although it did not reach the threshold of 0.05 as well. This may potentially be caused by the generational difference in education where people of different age groups went to school at different eras.

There are several limitations in this paper. Firstly, the participants were recruited mainly through social circles. This may lead to bias in the paper as the sample may not represent the entire population of Indonesia and Singapore. Secondly, instead of collecting the exact ages of the participants, the survey categorized the participants into 7 age groups: 15 years and under, 16–19 years, 20–25 years, 26–35 years, 36–45 years, 46–55 years and 56 years and above. This may lead to the comparison between the age groups being less detailed and precise. Thirdly, some of the open-ended responses were translated from Bahasa Indonesia to English. Errors in the translation of these responses and misinterpretations of the open-ended questions may contribute to less accurate data. Fourth, some people have lived in their country for less than a year. This may cause their results to not be representative of the general population. Another limitation is that multiple ANCOVA tests were conducted without applying a correction for multiple comparisons and, therefore, results should be interpreted cautiously. Finally, there is an uneven number of subjects between countries and age groups. This means that the data must be interpreted carefully, as some groups may be overrepresented while others may not fully represent the wider population.

Future research on this topic should have more systematic ways of recruiting participants such as through online social media or survey platforms or through public survey distribution to make the sample more representative of the population. Future research can also collect the exact ages of participants instead of collecting participant's age range, which will lead to more accurate findings. Additionally, it might be worth taking into account which specific area of the country participants live in. Different areas of a country might have different perceptions due to a number of factors, including differences in resources, school systems, and access to support (Ahmed et al., 2022; Wirandana and Khoirunurrofik, 2024). Social media has been a big part of the findings from this research. Future research can go into which types of social media are harmful or beneficial to the topic of learning difficulties.

Conclusion

This survey study compares how perceptions about learning difficulties differ between Indonesia and Singapore across different age groups. Overall, the results showed that Singapore had better perceptions out of the two countries, although this differed across the four domains studied. It also showed that respondents aged under 19 had the best perceptions out of the three age groups compared, with respondents above 46 having the second highest score, followed by respondents aged 20–45. However, ANCOVA analysis found that the age group was not a significant factor while country significantly affected country sum scores. The study also found that the differences in perspective may be due to different factors such as social media, public visibility, advancements in science, better education systems, and economic factors. Previous research supports the idea that social media can influence how people access health and educational information (Zhao and Zhang, 2017), while scientific developments have improved understanding of learning difficulties as neurodevelopmental conditions (Peterson and Pennington, 2015; Butterworth et al., 2011; Willcutt et al., 2005). From the results of the study, respondents most frequently attributed the change in perceptions of learning difficulties to the increasing use of social media, suggesting that online platforms have played an important role in spreading awareness. Overall, the findings show that awareness of learning difficulties is increasing, but the level of understanding still differs between countries, age groups, and social settings.

Works Cited

Ahmed, Syeda Kashfee, et al. "Teacher Professional Development for Disability Inclusion in Low- and Middle-Income Asia-Pacific Countries: An Evidence and Gap Map." Campbell Systematic Reviews, vol. 18, no. 4, 2022, e1287. https://doi.org/10.1002/cl2.1287.

Andriana, Elga, and David Evans. "Voices of Students with Intellectual Disabilities: Experiences of Transition in 'Inclusive Schools' in Indonesia." British Journal of Learning Disabilities, vol. 49, no. 3, 2021, pp. 316–328. https://doi.org/10.1111/bld.12411.

Bedi, Arjun S., and Ashish Garg. "The Effectiveness of Private versus Public Schools: The Case of Indonesia." Journal of Development Economics, vol. 61, no. 2, 2000, pp. 463–494. https://doi.org/10.1016/S0304-3878(00)00065-1.

Butterworth, Brian, Sashank Varma, and Diana Laurillard. "Dyscalculia: From Brain to Education." Science, vol. 332, no. 6033, 2011, pp. 1049–1053. https://doi.org/10.1126/science.1201536.

Chua, Yi Chian, et al. "Academic Expectations Stress, Its Associations with Symptoms of Depression and Anxiety, and Preliminary Exploration of Suicidality: Results from a National Youth Mental Health Survey in Singapore." Frontiers in Psychology, vol. 17, 2026. https://doi.org/10.3389/fpsyg.2026.1785980.

Deng, Zongyi, and S. Gopinathan. "PISA and High-Performing Education Systems: Explaining Singapore's Education Success." Comparative Education, vol. 52, no. 4, 2016, pp. 449–472. https://doi.org/10.1080/03050068.2016.1219535.

Dobson Waters, Sharon, and Carole J. Torgerson. "Dyslexia in Higher Education: A Systematic Review of Interventions Used to Promote Learning." Journal of Further and Higher Education, vol. 45, no. 2, 2021, pp. 226–256. https://doi.org/10.1080/0309877X.2020.1744545.

Gelber, Scott M. "From Compensation to Accommodation: The History of Learning Disabilities in American Higher Education." History of Education Quarterly, vol. 65, no. 1, 2025, pp. 8–28. https://doi.org/10.1017/heq.2024.16.

Haft, Stephanie L., Caroline Greiner de Magalhães, and Fumiko Hoeft. "A Systematic Review of The Consequences of Stigma and Stereotype Threat for Individuals with Specific Learning Disabilities." Journal of Learning Disabilities, vol. 56, no. 3, 2023, pp. 193–209. https://doi.org/10.1177/00222194221087383.

Handayani, Nurvia, and Nurul Hanifa. "Pengaruh Ketimpangan Pendapatan, Tingkat Pendidikan dan Kemiskinan Terhadap Pertumbuhan Di Indonesia." Independent: Journal of Economics, vol. 4, no. 1, 2024, pp. 112–124. https://doi.org/10.26740/independent.v4i1.60729.

Huff, W. G. "What Is the Singapore Model of Economic Development?" Cambridge Journal of Economics, vol. 19, no. 6, 1995, pp. 735–759. https://doi.org/10.1093/oxfordjournals.cje.a035339.

Jamaludin, Azilawati, et al. "Low-Progress Math in a High-Performing System: The Role of

Math Anxiety in Singapore's Elementary Learners." The Asia-Pacific Education Researcher, vol. 33, no. 5, 2024, pp. 1071–1082. https://doi.org/10.1007/s40299-023-00773-7.

James, Estelle, Elizabeth M. King, and Ace Suryadi. "Finance, Management, and Costs of Public and Private Schools in Indonesia." Economics of Education Review, vol. 15, no. 4, 1996, pp. 387–398. https://doi.org/10.1016/S0272-7757(96)00035-0.

Katz, Idit, Marianna Alesi, and Angelica Moè. "Homework Stress and Learning Disability: The Role of Parental Shame, Guilt, and Need Frustration." Learning Disabilities Research & Practice, vol. 37, no. 4, 2022, pp. 231–241. https://doi.org/10.1111/ldrp.12294.

Landerl, Karin, Alana Bevan, and Brian Butterworth. "Developmental Dyscalculia and Basic Numerical Capacities: A Study of 8–9-Year-Old Students." Cognition, vol. 93, no. 2, 2004, pp. 99–125. https://doi.org/10.1016/j.cognition.2003.11.004.

Landerl, Karin, et al. "Dyslexia and Dyscalculia: Two Learning Disorders with Different Cognitive Profiles." Journal of Experimental Child Psychology, vol. 103, no. 3, 2009, pp. 309–324. https://doi.org/10.1016/j.jecp.2009.03.006.

Ma, Kang. "Revisiting the English Language Test of Singapore's Primary School Leaving Examination: A Validity Argument." Language Testing in Asia, vol. 15, 2025. https://doi.org/10.1186/s40468-025-00402-8.

McCurdy, Amy L., et al. "Measurement of Parental Autonomy Support: A Review of Theoretical Concerns and Developmental Considerations." Journal of Family Theory & Review, vol. 12, no. 3, 2020, pp. 382–397. https://doi.org/10.1111/jftr.12389.

Peterson, Robin L., and Bruce F. Pennington. "Developmental Dyslexia." Annual Review of Clinical Psychology, vol. 11, 2015, pp. 283–307. https://doi.org/10.1146/annurev-clinpsy-032814-112842.

Petretto, Donatella Rita, et al. "The Use of Distance Learning and E-Learning in Students with Learning Disabilities: A Review on the Effects and Some Hint of Analysis on the Use during COVID-19 Outbreak." Clinical Practice & Epidemiology in Mental Health, vol. 17, 2021, pp. 92–102. https://doi.org/10.2174/1745017902117010092.

Sharfi, Kineret, and Sara Rosenblum. "Activity and Participation Characteristics of Adults with Learning Disabilities: A Systematic Review." PLOS ONE, vol. 9, no. 9, 2014, e106657. https://doi.org/10.1371/journal.pone.0106657.

Smith-Spark, James H., et al. "Executive Functions in Adults with Developmental Dyslexia." Research in Developmental Disabilities, vols. 53–54, 2016, pp. 323–341. https://doi.org/10.1016/j.ridd.2016.03.001.

Swanson, H. Lee, and Maureen Hoskyn. "Experimental Intervention Research on Students with Learning Disabilities: A Meta-Analysis of Treatment Outcomes." Review of Educational Research, vol. 68, no. 3, 1998, pp. 277–321. https://doi.org/10.3102/00346543068003277.

Thaver, Thana, and Levan Lim. "Attitudes of Pre-Service Mainstream Teachers in Singapore

towards People with Disabilities and Inclusive Education." International Journal of Inclusive Education, vol. 18, no. 10, 2014, pp. 1038–1052. https://doi.org/10.1080/13603116.2012.693399.

Willcutt, Erik G., et al. "Validity of the Executive Function Theory of Attention-Deficit/Hyperactivity Disorder: A Meta-Analytic Review." Biological Psychiatry, vol. 57, no. 11, 2005, pp. 1336–1346. https://doi.org/10.1016/j.biopsych.2005.02.006.

Wirandana, Priaji Agung, and Khoirunurrofik. "Educational Inequality in Indonesia: Are Intergovernmental Fiscal Transfers Effective in Reducing the Gap?" Educational Studies, vol. 50, no. 6, 2024, pp. 1424–1443. https://doi.org/10.1080/03055698.2022.2103647.

Zhao, Yuehua, and Jin Zhang. "Consumer Health Information Seeking in Social Media: A Literature Review." Health Information & Libraries Journal, vol. 34, no. 4, 2017, pp. 268–283. https://doi.org/10.1111/hir.12192.

Personalisation, Judicial Capture, and the Politics of Gender: Democratic Backsliding under Trump and Erdoğan By Sophia Fox

Abstract

This essay examines the relationship between executive personalisation, judicial transformation, and women's-rights erosion under Donald Trump and Recep Tayyip Erdoğan. It brings together three strands of literature—democratic backsliding, executive personalisation and women's-rights erosion—to examine how they interact across two contrasting political systems. The paper argues that populist anti-pluralism can facilitate executive personalisation by presenting political leaders as the authentic representatives of the people and portraying institutional constraints as obstacles to popular sovereignty. Where judicial independence is weakened or judicial composition is politically reshaped, courts can become important sites through which political preferences are entrenched. The paper treats women's-rights erosion as both an important consequence of these political developments and a useful lens through which broader democratic erosion can be examined. Using a structured comparative case-study approach, the paper compares Turkey, where democratic backsliding has become extensive, with the United States, where institutional erosion has been more limited. Turkey demonstrates extensive institutional transformation through constitutional reform, judicial purges and the expansion of presidential authority, culminating in the withdrawal from the Istanbul Convention. The United States demonstrates a more limited form of institutional change, in which Trump's judicial appointments contributed to a transformation of the Supreme Court that became particularly consequential in Dobbs v. Jackson Women's Health Organization. The comparison demonstrates that similar political pressures can produce different institutional consequences depending on the political system in which they operate.

Keywords

Comparative Politics, Democratic Backsliding, Judicial Capture, Executive Personalisation, Gender Politics and Authoritarianism

Introduction

In March 2021, President Recep Tayyip Erdoğan withdrew Turkey from the Istanbul Convention (a treaty designed to prevent violence against women and protect women’s rights) by presidential decree, thereby unmaking a landmark treaty obligation that Turkey had been the first to ratify and was aimed at guaranteeing robust domestic rights protections (BBC). Fifteen months later, in the United States, Roe vs Wade, a groundbreaking Supreme Court decision to establish a constitutional right to pregnancy termination, was overturned (Totenberg and McCammon). These two occasions appear to share little beyond subject matter, being reported as very distinct cases: one a foreign policy story and the other a domestic legal change. This analysis will establish that their supposed disparity shares more than coincidental content

overlap; both instances represent symptoms of the same underlying process despite their different stages of progression.

The analysis will discuss how extensively the concentration of power in the executive, observed in both Turkey and the United States, has resulted in the weakening of institutions that uphold the democratic process; the impacts of this personalisation on democratic health will be examined through the concept of judicial capture as the core process through which backsliding occurs with the study of gender politics helping to recognise and diagnose the severity of decline. Thus, this paper asks:

How have executive personalisation and judicial transformation contributed to the erosion of women's rights in Turkey and the United States, and what does this reveal about democratic backsliding?

It argues that populist anti-pluralism can facilitate the personalisation of executive authority by weakening the legitimacy of institutional constraints,by relocating legislative authority into the judicial system. In all cases, the judicial system, originally intended to constrain executive power, acts as an insulated institution exempt from ordinary electoral accountability. Gender policy is one of the primary terrains through which personalist, populist rule is performed, and a measure because the severity of its rollback acts as a diagnosis in each case. However, the institutional consequences of personalisation depend on the political system in which it occurs. In Turkey, executive personalisation was accompanied by major constitutional changes and the restructuring of the judiciary. In the United States, Trump's influence over the judiciary operated primarily through existing constitutional appointment procedures. Women's-rights controversies therefore provide a useful lens through which to compare these different forms and degrees of democratic erosion.

Democratic backsliding has emerged in the twenty-first century as one of the most significant concepts across modern democracies, appearing in the post-Cold War era as a reversal of the global expansion of democratic governance (Bermeo 6). Turkey and the United States represent two regularly studied nations amongst scholars of comparative politics, but they follow distinctly different sequences of progression. The first, Turkey, has seen a transition from tutelary democracy to competitive authoritarianism spanning over two decades of Adalet ve Kalkınma Partisi (AKP) rule, whilst the second, the USA, has experienced a less severe degree of backsliding across one presidential term, within a consolidated constitutional democracy. Both cases, whilst systemically disconnected, demonstrate significant executive authority being contained in a single leader, that single leader engaging in sustained conflict with the judicial system, and a rollback of hugely impactful women’s rights developments.

Since democratic backsliding is commonly measured through features of democracy like media freedom, minority rights, and electoral integrity, the choice of women’s rights as an indicator requires justification. Each of these alternatives do carry weight as important measurements of democratic health and this paper does not claim that women’s rights equal

democracy, nor that they should replace conventional indicators. However, it does assert that the removal of established women's rights exposes weak institutional conditions and protection of democratic rights. A right that has been codified for an extended period of time should be upheld by democratic institutions; therefore, its removal tests whether these institutions can resist executives who threaten them. Gender also intersects with almost all other indicators. For example, legislative constraint is absent when treaties are revoked by decree and judicial independence is questioned when courts become the site of rights erosion. Moreover, civil liberties diminish as bodily autonomy and freedom of speech are curtailed, media freedom and electoral integrity are undermined when female journalists/candidates face harassment, and minority rights are threatened at the intersection of gender, class, ethnicity, and sexuality (Roggeband and Krizsán 9-10). For these reasons, women's rights should be treated as a core issue of comparative politics because the pattern of rollback reveals how and how far protections of democracy have weakened.

Existing scholarship (covering democratic backsliding, judicial capture, and gender politics) has developed as three channels of study (Roggeband and Krizsán 1, 9). This paper extends comparative politics by proposing an impact sequence – personalisation as the effect of populism, judicial capture as the institutional manifestation of this personalisation, and women's rights erosion as a key area of impact – across two visibly disparate systems. This comparison follows the Most Different Systems Design framework, developed by Przeworski and Teune (34-39) to test if the impact chain identified above holds in two cases with dissimilar institutional history and makeup, regional context, and economic development.

This analysis will consist of four parts: it will begin by establishing the theoretical vocabulary through which backsliding, judicial capture, and gender politics will be examined, drawing on existing literature. It will then apply this vocabulary to Turkey, tracing the shift from tutelary democracy to competitive authoritarianism under Erdoğan's personalisation, post-2016 judicial purges, and the instrumentalisation and subsequent erosion of women's rights over the past two decades. The penultimate section will map the trajectory of the United States throughout Donald Trump's tenure using this same vocabulary, arguing that his personalisation was achieved through a judicial appointment model and that Dobbs vs Jackson Women's Health was less severe than Turkey's 2021 Convention withdrawal but still operates as an example of norm-cascade reversal. Finally, the paper will conclude by drawing the two cases together to assert that the comparative severity of gender rights rollback offers an applicable measure of backsliding progression in both.

Literature review and theoretical framework

This review will establish the theoretical vocabulary through which a comparison between Donald Trump's America and Recep Erdoğan's Turkey will be made. Judicial capture will be used throughout as the connecting mechanism between the two cases; it will also include an examination of literature on populism and polarisation, which will help determine the political conditions that allowed for democratic backsliding. Finally, it will argue that women's rights

should not solely be studied as an example of democratic erosion, but instead as a diagnosis for the severity of this deterioration.

Backsliding and Personalisation of Power

The comparison of a highly established constitutional democracy like the United States against Turkey, a competitive authoritarian regime, requires the definition of what is actually being determined as "democratic backsliding" by visiting the established principles of modern democracies. Firstly, Dahl's (1-9) "Polyarchy" determines a "democracy" to be a state in which the government is consistently and effectively responsive to the desires of its inhabitants, providing the baseline indicators of a functioning liberal state: simultaneously high levels of public contestation and inclusivity (the right to participate). Instead of an idealistic vision of representative democracy, he coined the concept of a polyarchy to describe the realities of world order; this framework will remain the baseline to which democratic backsliding in Turkey and the United States will be measured. Bermeo (6) refines this as she distinguishes between contemporary backsliding and historical variations, such as executive coups and coups d'etat, prevalent in the Cold War Years.

Despite democratic backsliding being broadly understood as the weakening of institutions that sustain democracy, the distinction of "executive aggrandisement" provides the connection through which the broadly disparate political conditions in the United States and Turkey can be paralleled. Bermeo defines executive aggrandisement as occurring when incumbents weaken checks on their power incrementally (10-13). This tends to operate through a series of legally sanctioned changes, rather than a single decisive rupture from accountability and the democratic process. She theorises executive aggrandisement as being the dominant form of backsliding in the twenty-first century because its apparent legality and gradual nature prevent it from triggering public alarm (10-13). Unlike the seismic shifts created by overt coups or election fraud, each step can be defended as ordinary politics so that the concentration of power is only visible with hindsight. Esen and Gümüşçü apply this to Turkey by examining the use of constitutional referendums by the AKP (most notably the 2017 shift to executive presidency), as well as the utilisation of emergency powers following the 2016 coup attempt, thereby identifying legal attempts to dismantle prior checks on the executive (1586-601). In the American case, scholars like Levitsky and Ziblatt point to more incremental tools that often take years to fully implement: an expansive use of executive orders/emergency declarations and deliberate Supreme Court appointment strategies (106-12). Ultimately, the mechanism is shared between both case studies, despite the speed and scale differing considerably.

However, it is important to note that Bermeo hypothesised the concept of "executive aggrandisement" with reference to democratising and hybrid regimes, such as Turkey, in which accountability was only partially institutionalised in the first place (10-13). Applying this theory to the United States, which is a consolidated democracy with a robust, independent judiciary and free press, is therefore contestable. Nonetheless, this review will proceed on the basis that the

mechanism identified by Bermeo is a highly useful channel of comparison and its inclusion is a deliberate choice, rather than an uncontested application (10-13).

Finally, Levitsky and Ziblatt operationalise the formation of an illiberal democracy with their idea of "referee capture" (manipulation of court systems, electoral bodies, and regulators)(78-81), which is extended by the judicial capture literature discussed below.

Judicial Transformation and Institutional Capture

Whilst executive aggrandisement can be labelled as the general process through which democratic backsliding is occurring in the United States and Turkey, this research question specifically interrogates judicial capture as the mechanism. Starting with judicialisation, Hirschl's "Towards Juristocracy" states that legal systems become sites of "hegemonic preservation" in which increasingly threatened political elites transfer legislative and constitutional power into the judiciary (38-49). This acts as a method of interest preservation in the face of declining popular support. The chapter of Hirschl's book on the "Constitutionalisation and the Judicialisation of Mega-Politics" explains why courts more commonly adjudicate questions of constitutional legitimacy or executive authority (169-211): once elites can no longer source their preferred outcomes from the majoritarian process because of volatile public support, they have an incentive to relocate adjudication on contentious matters to courts as they are less accountable to voters and can be shaped through judicial appointments. This theory is applied to Turkey in Esen and Gümüşçü's (1586-601) account of the nation's transition away from tutelary to competitive authoritarian rule. It directly names judicial capture as one of the tools through which the AKP heavily limited the ability of the opposition to campaign and compete for power, making it highly improbable that the incumbent is defeated. In the United States, the Dobbs case provides a parallel. Dobbs vs Jackson Women's Health Organisation (2022) overturned both Roe vs Wade (1973) and Planned Parenthood vs Casey (1992), removing the federal constitutional right to abortion across the United States and returning its oversight to individual states. This outcome was made possible by the confirmation of three judges appointed during Trump's tenure from 2017 to 2020, which produced an ideological majority on the Supreme Court. Johnson argues that the overturning of Roe vs Wade exemplifies a departure from "substantive due process". The narrowing of the Due Process Clause represents judicial capture through the appointment process rather than through blatant court-packing, but it has nevertheless seen accusations of democratic erosion. Ultimately, both examples demonstrate judicial capture operating with the same logic (relocation of sovereignty to benefit political actors) despite being divided by direct court-packing in Turkey's case versus the gradual reshaping of the Supreme Court's composition in the United States. The question is therefore not whether the American judiciary was captured in the same sense as Turkey's, but whether political control over judicial appointments can produce some of the same consequences for rights protection.

Populism and Polarisation

In order to determine the popular attitudes that breed support for democratic backsliding, one must look to literature on populism. For example, Müller characterises the core logic of populism as being anti-pluralist. In other words, populist leaders claim to be the sole authentic representation of popular will, thereby delegitimising independent institutions (such as courts) as obstacles to the desires of state citizens rather than democratic safeguards (19-20). This provides a justification for both Trump and Erdoğan's dismissal of judicial intervention against them, but also for their manipulation of court composition being branded as necessary for effective rule. McCoy, Rahman, and Somer further validate the comparison between the United States and Turkey in their analysis of "pernicious polarisation" (16-20). They argue that polarisation causes positive-sum interests to transform into zero-sum interests in which mutual exclusion eliminates the possibility of collective action. This reinforces perceptions of society in terms of "Us versus Them" and also acts as a precursor to democratic erosion across seemingly disparate political landscapes with varying regime histories. Inglehart and Norris add a "cultural backlash" dimension, exploring the consequences of promoting post-materialist values over traditional values ("Trump" 2-4). Progressive social movements have gained momentum on a global scale over the last few decades, building on liberal principles of autonomy and self-expression. Ultimately, their characterisation of populist authoritarianism as a backlash response to social and cultural change explains why gender roles and family values lie at the heart of both Trump and Erdoğan's leadership. The contestation of women's rights in both cases, granted to a more extreme degree in Turkey, makes Norris and Inglehart's thesis a justification for this essay's focus on women's rights. Namely, gender policy being a terrain onto which broader anxieties are projected, demonstrates why the examination of gender rights disparity not only confirms the presence of democratic backsliding, but also acts as an instrument to measure its severity. Populism and polarisation therefore do more than just create conditions for backsliding. Instead, by casting the "people" as a single bloc opposed to corruption and the illegitimate "other", they establish demand for one individual who claims to embody that bloc; this popular inclination to accept personalisation makes it the precursor to the phenomenon.

The Personalisation of Politics

Despite this, populism only explains the popular conditions for backsliding, rather than identifying how anti-pluralist sentiment becomes concentrated in office. The literature on the personalisation of politics provides this missing link, examining how this concentration occurs.

Firstly, Poguntke and Webb offer the foundational framework in "The Presidentialisation of Politics", arguing that executives across democracies have become increasingly presidentialised (4-11). They essentially identify a pattern through which the power held by these executives is becoming highly concentrated within an individual leader, often at the expense of other decision-making bodies like cabinets and parties. They distinguish three separate "faces" of this process. The first face is the executive face, which describes the leader's growing autonomy as a result of cabinet cooperation. The second is the party face, or the leader's growing autonomy from the party itself. Finally, the electoral face encompasses the personalisation of campaigning

around the leader as an individual brand, thereby diminishing party politics. Ultimately, Poguntke and Webb (4-11) attribute this phenomenon to several causes: the internationalisation of politics (whereby inter-governmental negotiation and cooperation shifts power to individual leaders and key advisors), the growth of the state, meaning presidentialisation is required logistically, electronic communication which favours symbolism over substance, and finally, the loss of social anchoring in which more pluralist support bases require leaders with universally-perceived personality traits. Although their thesis was developed with parliamentary systems in mind, Poguntke and Webb (4-11) do not confine presidentialisation to one constitutional form, but rather confirm that it can occur within presidential systems. This, therefore, makes it transferable to both the United States (an already-presidential system which is seeing its executive and party faces intensify) and Turkey (which established presidentialisation constitutionally in 2017). Because the executive face concentrates decision-making authority in a single leader, the leader's personal autonomy is thereby exercised through the executive office, meaning personalisation becomes inseparable from presidentialisation. This is why the paper locates personalisation within presidentialisation in later arguments.

Samuels and Shugart extend this by defining the mechanism through which personalisation reshapes party organisations. Their principal-agent framework argues that the separation of powers "presidentialises" parties because executives are elected distinctly from the legislature and so parties are incentivised to nominate more personalistic candidates who will maximise votes (1-12). This therefore means that party structures become organised around the requirements of the executive rather than around policy goals, thereby weakening the ability of the party itself to act as an internal check on the leader.

Women's Rights as a Lens on Democratic Backsliding

Kandiyoti links personalisation with populism and gender politics, grounding the "cultural backlash" thesis in practice. She characterises gender politics in Turkey as being a "masculinist restoration" (104-106). In other words, she identifies the erosion of women's rights as a reassertion of patriarchal authority as a defensive response to perceived threats against existing gender hierarchies. This concept has since been applied to the United States, where scholars argue that a similar rebuild of masculine authority is occurring through social policy/appointments, rather than through blatant constitutional change. Connell's concept of hegemonic masculinity (76-81) further explains why populist leaders tend to utilise gender politics. Because a dominant masculinity is constructed to counteract both femininity and pluralism, it offers established anti-pluralist vocabulary that can be exploited by populist figureheads. Finally, Hudson et al. 's "Sex and World Peace" includes the central claim (that the physical and legal status of women acts as a leading indicator of political stability and quality of governance) (1-10), providing the grounds for treating gender policy as diagnostic. Finnemore and Sikkink's (895-905) norm life-cycle model allows Turkey's 2021 withdrawal from the Istanbul Convention and the overturning of Roe vs Wade in the United States to be used as norm-cascade reversals. By extending their framework to run their model backwards, one can

characterise these reversals as moments where previous institutional and constitutional norms are unmade by government officials. Arat demonstrates that the AKP's instrumentalisation of women's rights was integral to the establishment of competitive authoritarian rule in Turkey, with the 2010 referendum's gender equality language later being used to cover judicial appointment changes that entrenched party loyalists (13-14).

In summary, existing scholarship offers three developed but largely independent channels of analysis. The first discusses the personalisation of executive power, the second uses judicial capture as a mechanism for this personalisation, and the third documents women's rights erosion within the individual contexts of Turkey and the United States. No existing source blends all three into a single chain of cause and effect or connects the cases of both nations in a singular analysis. This paper has argued that doing so is justified, as demonstrated by Hirschl's "judicialisation" and Hudson et al.'s framing of gender as a government stability indicator.

Taken together, these literatures provide a framework for examining a relationship between executive personalisation, institutional transformation and women's-rights erosion. The contribution of this paper is therefore not to propose an entirely new theory of democratic backsliding, but to bring these existing perspectives together through a comparison of two contrasting cases.

It determines that the personalisation of politics should be considered a byproduct of populism rather than a parallel occurrence, that judicial capture is the mechanism through which the personalised executive entrenches popular mandate in constitutional power, and that the erosion of women's rights is not only a common policy feature, but instead the clearest example as to how far this capture has progressed.

Discussion: Backsliding in Turkey

This paper will identify the starting point for subsequent Turkey analysis as its prior status, a "Tutelary Democracy", rather than a consolidated democracy. Identified by Esen and Gümüşçü (1583-86), the term signifies a hybrid system in which competitive elections coincide with unelected guardians; in the case of Turkey, these "guardians" consisted primarily of the military and the judiciary, who retained the power to override outcomes of the democratic process if they deemed the secular, Kemalist order to be under threat. Dahl's "Polyarchy" framework (1-9) reinforces the inability to classify Turkey pre-2002 as a functioning democracy: despite possessing one marker of democratic status (public contestation) due to the presence of multi-party, competitive elections, it still fell short of the second indicator (inclusivity) since preferences of the electorate could be overridden by non-majoritarian forces. Therefore, Turkey's trajectory under the AKP cannot be classified as a direct transition from a healthy democracy to an authoritarian regime, but instead, as Esen and Gümüşçü recognise (1581-83), Turkey has shifted from one flawed democracy to another (competitive authoritarianism): contestation is visible in the nation's elections, yet the institutional structures systematically favour the incumbent, meaning party pluralism is superficial.

The AKP's actions in power further illustrate this shift. Firstly, on paper, the AKP's victory was democratically legitimate, and the 2010 constitutional referendum (which restructured the Constitutional Court and the Supreme Board of Judges and Prosecutors) was initially framed as democratisation, as it eliminated the guardianship of the judiciary. Despite this apparent advancement of democracy, it can also be read through Hirschl (38-49) and Esen and Gümüşçü's (1586-601) frameworks as the first major act of judicial capture. Instead of judicial liberalisation, the reforms transferred power to appoint the judiciary's senior bodies to the executive and away from the secular establishment, failing to eliminate tutelage altogether. Finally, the failed coup attempt of July 2016 provided the conditions for rupture: a state of emergency was declared in its aftermath, thereby allowing the AKP to carry out sweeping purges of the judiciary, military, civil service, and academia. This period of emergency rule both preceded and created the environment enabling the 2017 Constitutional Referendum, which formally abolished the office of Prime Minister and transferred legislative and appointment powers to an executive presidency. Ultimately, this sequence demonstrates the pattern predicted by Bermeo's definition of executive aggrandisement: whilst each step was legally defensible, Turkey was ultimately transformed into a competitive authoritarian regime through a series of incremental reforms (10-13).

Erdoğan's personalisation

When applying Poguntke and Webb's presidentialisation thesis to Erdoğan, all three "faces" of presidentialisation can be seen as becoming concentrated around a single figure (4-11). Firstly, the executive face is most evident following the 2017 referendum, in which cabinet government and the concept of collective ministerial responsibility were abolished; instead, this political order was replaced in favour of a presidency that appoints and dismisses ministers, whilst simultaneously possessing the ability to issue decrees with the force of law (lacking the countersignature of the prime minister). The second face is visible in Erdoğan's relationship with the AKP. After co-founding the party in 2001, he has remained its dominant figure across two decades and multiple offices (first holding the position of prime minister, then president). Under his leadership, the AKP's internal candidate selection and policy direction have become entirely subordinate to his preferences, aligning with Samuels and Shugart's account (1-12) of how a directly elected executive is able to presidentialise the party, producing their legitimacy. For example, before 2017, it was constitutionally mandated that Turkish presidents sever party ties upon taking office; however, the subsequent constitutional changes removed this requirement, and Erdoğan formally resumed chairmanship in May 2017, therefore ensuring continued intervention in party direction and policy. Finally, the electoral face is apparent in the personalisation of AKP campaigns themselves, most notably post-2014, which centre around Erdoğan as an embodiment of strong, decisive leadership, rather than around collective party intentions.

Drawing on Elgie and Passarelli's (2018) distinction between formal, constitutional presidentialisation, and informal personalisation, it is evident that Erdoğan's case demonstrates

both phenomena operating simultaneously. Before 2017, his dominance over the AKP organisation and cabinet decision-making was largely confined to informal personalisation (i.e., concentration of authority in a single leader without formal constitutional change in Turkey). Post-2017, however, the referendum converted this personalist practice into constitutional presidentialisation, giving Erdoğan's authority an institutional foundation that is capable of outlasting electoral cycles. This is relevant within the context of this analysis because it demonstrates that personalisation is not solely a reflection of popularity or style, but rather a process (as identified by Bermeo's executive aggrandisement (10-13)) with a goal of entrenchment.

The popular conditions enabling this personalisation are grounded in populism and polarisation literature. Müller's account of populism as anti-pluralist rhetoric (19-20) is reflected in Erdoğan's recurring claim to represent the authentic "national will" (milli irade) against corrupt and self-interested secular elites. This logic purposefully delegitimises secular opposition in addition to reframing institutional checks on executive overreach as barriers to the people's sovereignty. McCoy, Rahman and Somer's "pernicious polarisation" (16-20) is also identifiable within Erdoğan's language as he deliberately utilises an "us versus them" binary choice when addressing the tension between religion and secularism. This has proven capable of sustaining mobilisation across two decades by converting differences in policy into an existential, zero-sum contest between two versions of Turkish identity. Alongside this, Norris and Inglehart's "cultural backlash" thesis (Cultural Backlash 31-36) explains the AKP's traditionalist voter base as a defensive reaction to the liberalising nature of the Kemalist republic over the twentieth century; this thereby explains the AKP's intention to promote itself as the dominant electoral force through which the backlash community can mobilise. Nevertheless, the question of why gender specifically became the terrain through which the AKP's traditionalist authority was asserted still remains. Kandiyoti's concept of "masculinist restoration" (104-06) answers this in the case of Turkey by characterising AKP gender politics as a defence of the patriarchy against the considerable threat of Kemalist social liberalisation. Rather than being a re-establishment of an already secure patriarchal structure, Kandiyoti's analysis reframes the relationship between personalisation and gender policy: Erdoğan's personalist rule was not merely used against women's rights, but his attitudes towards gender politics were one of the primary mechanisms through which he claimed to embody traditionalist desires.

Judicial capture as an AKP mechanism

Hirschl's concept of "hegemonic preservation" (38-49) is directly relevant to the case of Turkey. Despite the AKP's popular support being consistently strong, it has experienced periods of fluctuation – most notably following the loss of Istanbul and Ankara's mayoralties in 2019. These incidents of uncertainty act as a clear incentive for officials facing electoral pressure to relocate legislative power away from mechanisms that are vulnerable to volatility, and into structures like the judiciary that are not subject to immediate popular accountability. Esen and Gümüşçü's study of competitive authoritarianism in Turkey (1586-601) directly references

judicial capture as a central tool through which this entrenchment has been achieved by the AKP. By packing the courts that select YSK (Supreme Electoral Council) members, Erdoğan has effectively captured the Council externally; since this takeover, the Supreme Electoral Council has rejected numerous formal complaints from opposition parties as well as acting alongside a loyalist-packed judiciary to selectively prosecute, fine and disqualify opposition politicians, most visibly the pro-Kurdish HDP (Esen and Gümüşçü, 1586-601). This produces the defining feature of competitive authoritarian regimes, as identified by Esen and Gümüşçü (1583-86): in form, elections stay relatively contested, yet the incumbent's defeat is almost impossible in practice.

The specific mechanism through which the high courts have been captured by the AKP is the 2016 purge of the judiciary. Following the failed coup in 2016, the declaration of national emergency facilitated a huge reshuffle of court officials; several thousand judges and prosecutors were dismissed with immediate effect, and in many cases, prosecuted (Esen). These individuals were then replaced through an appointment process brought under the effective control of President Erdoğan and the AKP-dominated Council of Judges and Prosecutors, as formalised by the reforms in 2010 and 2017. Here, we can see a clear example of court-packing that differs from the appointment process exploited in the US: rather than waiting for vacancies to arise through ordinary process, the Turkish state manufactured the vacancies directly and filled them within a short time period, achieving in months the court composition transformation that would take several years in established democracies. This comparison is worth noting because the disparity between the overt purge model in Turkey and the gradual appointment model in the US demonstrates how judicial capture exists across ranging speed and scale.

The Erosion of Women's Rights in Turkey

Finnemore and Sikkink's norm life-cycle model (895-905) provides a useful framework to interpret the 2021 Istanbul Convention withdrawal, when extended to run its stages in reverse. Their original model theorises three stages: norm emergence driven by "norm entrepreneurs" who call attention to prevalent social issues, socialisation cascading the norm past the tipping point (norm adoption by several states), and finally, norm internalisation whereby norms become so widely accepted that they are taken-for-granted. These same stages can be read backwards to identify a reversal of the norm life-cycle: when a previously internalised international and domestic norm is entirely unmade by government action. Instead of being underenforced by the AKP, Erdoğan's government deliberately and definitively ruptured gender norms by withdrawing from the Istanbul Convention (a treaty specifically designed to combat domestic violence and protect women). Despite being the first nation to ratify the Convention in 2011, Turkey's unilateral withdrawal by presidential decree a decade later fits the extension of this model and signifies a decisive norm of death, as identified by Panke and Petersohn (5-12). They determine that the prospects of norm change are contingent on the relative strength of the norm proponents in relation to the norm challengers. In the case of the Convention withdrawal, it is clear, given the judiciary's capture in 2010 and 2017 (as well as the AKP's parliamentary domination), that there is an absence of actors who are able to sufficiently defend the norm from the dominance of

these challengers. One must also consider, however, the mechanism of withdrawal. Since it was achieved by presidential decree rather than through parliamentary practice, it constitutes an exercise of personalised executive authority under the powers granted following the 2017 referendum. Ultimately, the Convention withdrawal also functions as a particularly visible illustration of the extent of institutional and normative change because it determines how far personalisation had progressed by 2021, permitting a symbolic norm of legal significance to be reversed without parliamentary involvement.

Arat's analysis determines backsliding under the backdrop of women's rights erosion to have begun after the 2010 referendum (2, 14). She determined the AKP to have instrumentalised women's rights language before the constitutional referendum to create the conditions required for judicial capture. Leading up to the vote, the AKP deliberately pledged to strengthen the gender equality clause, marketing the expansion of rights protections to the public. This rhetoric was specifically used to secure a broad coalition of support needed to pass the reforms, which ultimately allowed for the restructuring of the judicial appointment bodies and the subsequent entrenchment of AKP loyalists (13-14). Within this, the erosion of women's rights can be seen as a measure for judicial capture because gender equality did not merely function as an affected policy area once the court packing had been completed, but rather, the rhetoric acted as a mechanism for mobilisation.

Finally, when accounting for Hudson et al.'s assertion (1-10) that the physical and legal status of women can be read as an indicator of political stability and governing quality, the severity and pace of the trajectory of women's rights provides an additional indicator of the broader institutional and political transformation taking place This validates the argument that women's rights are not only an affected policy area within democratic erosion, but instead, the immediate withdrawal from binding international law constitutes a clear measurement tool of severe personalisation and democratic backsliding.

Discussion: Backsliding in the USA

In contrast to the Turkish case, American backsliding has occurred to a much lesser degree and has taken place from within an already-consolidated constitutional democracy, rather than the nation experiencing a decisive transition to authoritarianism. This starting point makes democratic backsliding harder to measure and explain. Whilst Turkey's descent into competitive authoritarianism can be seen through a series of definitive constitutional reforms, the American system retains a robust separation of powers, judicial independence, and a protracted amendment process. What this analysis will instead study is the agency of individual political actors operating within existing constitutional rules: strategic appointments, gatekeeping decisions, and the erosion of unwritten norms.

Personalisation under Trump

The application of Poguntke and Webb's presidentialisation thesis to Donald Trump demonstrates personalisation occurring without formal constitutional change (4-11). Despite this

absence, it is visible that all "three faces" of presidentialisation are apparent within the American presidency. The first, executive face, is evident within the increased preference for unilateral executive orders and emergency declarations rather than legislative coalitions. The party face can be viewed in the increasing centrality of Trump to Republican electoral politics and party organisation, most notably in the 2020 decision not to adopt a new party platform at the (10-13) National Convention. Finally, the electoral face, arguably the most blatant channel of presidentialisation, has seen Trump become an individual figure of loyalty, building a significant "Make America Great Again" support base with himself at the centre, rather than any real references to a policy agenda. In comparison to Turkey's case, one must affirm that the presidentialisation of the American presidency did not require any amendment to its formal, constitutional powers. Instead, it demonstrates personalisation operating through the erosion of established norms such as mutual toleration and restraint (Levitsky and Ziblatt 102-12), which had previously ensured a constrained executive.

Elgie and Passarelli's formal-informal distinction, as discussed above, is similarly useful here: Turkey's risk became clear in a single moment (2017 referendum) regardless of who holds office, yet the American presidency has long been a directly elected office before Trump's tenure, meaning that the constitutional starting point inherited by the individual in office already has a greater degree of formal power concentration. This poses a different threat to democratic health since power cannot be further concentrated through legal amendment, meaning this process must be achieved through informal, norm-based channels; whilst this does suggest the possibility of reversal by whoever subsequently takes office, the fact that no constitutional violation exists means it is much harder to detect and challenge.

The popular conditions that made this personalisation viable and electorally successful follow the same populism and polarisation literature as Turkey's case study. Müller's anti-pluralism concept (19-20) is directly reflected in Trump's consistent rhetoric of a "Real America" in which he represents the only authentic representation of popular will. His framing of Democratic opponents as unpatriotic and corrupt, his public attacks on media and press, and his disapproval of independent institutions all position these political features as obstacles between himself and the true will of the people, diminishing legitimate contestation. McCoy, Rahman, and Somer's account of pernicious polarisation (30-36) directly names an alarming intensification of affective polarisation within American politics following the 2008 financial crisis. According to their study, racial resentment, ethno-nationalism and racial prejudice all surged with the election of America's first president of colour, and this rhetoric was fuelled by Trump's mobilisation of an "us versus them" tension. His scapegoating of immigrants, minorities, and liberal elites as having caused "American Carnage" (Trump) further bred this affective polarisation, enhancing zero-sum perceptions and framing the "other" as an existential threat. In doing so, Trump has ensured that partisan identities increasingly come to function as a marker of social belonging rather than policy disparities, ultimately resulting in institutional checks on the executive becoming battlegrounds for these partisan affiliations. Norris and Inglehart's "cultural backlash" thesis (Cultural Backlash 31-36) adds another layer by explaining

how Trump's core support base operates as a reaction to decades of post-materialist and progressive social change. This, in turn, positions his presidency, like Erdoğan's, as a manifestation of this backlash community, despite occurring against the backdrop of institutional secularism and rights protection.

Judicial Appointments and Institutional Transformation in the United States

Hirschl's "hegemonic preservation" (38-49), as in the case of Turkey, can be applied to the United States, yet the methodology through which it is achieved differs between the two nations. Instead of purging sitting judges, as has been the Turkish case since 2016 (discussed above), judicialisation has operated in the US through the regular appointment process. This method has been deliberately and increasingly utilised over the past few decades to ensure a durable ideological base in the Supreme Court as a buffer in the event of electoral volatility or unstable legislative dynamics. The Republican Party's control over the Senate confirmation process can be viewed through two examples and their cumulative effect: the first is the refusal to hold hearings on Merrick Garland's 2016 nomination (demonstrating the political importance of appointment timing), and the second is the expedited confirmation of Amy Coney Barrett in the final weeks of the 2020 election (which exemplifies strategic use of an available appointment). Together with Trump's other conventional appointments, these events completely shifted the ideological composition of the Supreme Court. This reflects the logic of Hirschl's argument in that elites have had growing incentives to relocate the security of preferences into less electorally accountable institutions in order to entrench these policy preferences in the face of unreliable electoral outcomes.

Nevertheless, this mechanism is different from Turkey in significant ways. Formal judicial independence did remain as sitting judges were not removed. As well as this, more broadly, judicialisation and judicial capture in the United States is contingent upon vacancies rather than being purge-driven, it depends on the timing of deaths, retirements, and electoral cycles. Instead of a single removal event, it is slower and less directly controllable by any one President. Despite this lesser speed and scale, outcomes are much more durable, and so once ideologically beneficial appointments have been achieved, life tenure ensures that any subsequent office holder is unable to reverse it, unlike in the case of Turkey, whereby the loyalist judiciary could be completely reshuffled if the incumbent were to fall from power.

Dobbs vs Jackson Women's Health Organisation (2022) provides the central study through which judicial capture and women's rights erosion converge. Johnson extends the point further by identifying the mechanism through which the appointment-driven capture was achieved as having effects that will reach beyond abortion rights as an immediate question; with substantive due process being the doctrine under which the Supreme Court had entrenched "fundamental" rights under the Fourteenth Amendment, including reproductive autonomy, marital privacy, and others, he predicts that the language provided in the Dobbs case will have detrimental effects on future rights rulings. The Supreme Court ruling in the Dobbs case determines that further claims of rights protection must be "deeply rooted in the Nation's history

and tradition" (Dobbs 231, 237-38), thereby significantly narrowing the tool through which future rights claims will be read. Furthermore, this shall disproportionately impact rights, notably women's rights, which have been institutionally realised recently, as a result of social and normative change.

Women's Rights in the United States

To a similar end, when applying the same extension of Finnemore and Sikkink's norm life-cycle model (895-905)developed above in relation to Turkey, one can view the Dobbs case as a norm-cascade reversal. By taking the stages within the original model in reverse, Dobbs can be read in parallel to Turkey's 2021 Istanbul Convention withdrawal: in both instances, a right that had been internalised and institutionalised for decades is unmade by government action, operating through two alternative routes. In the American case, the constitutional guarantee to reproductive autonomy was stripped as a result of the Supreme Court's ideological majority and subsequent ruling, rather than through presidential decree, yet despite this apparent absence of executive action, the condition identified by Panke and Petersohn (5-12) as necessary for a successful norm challenge is relevant, as in the Turkish analysis. The absence of an actor able to impose meaningful cost for the violation is directly applicable to the US, as the captured judiciary was itself the agent who implemented the reversal, but also the highest court in the land, therefore meaning no equivalent check is available above it. Ultimately, whether this agent is a presidential decree, as in Turkey, or a judicial ruling, is less relevant in relation to the shared features that make the reversal possible; in each case, personalisation had already relocated authority over norm-regulation into institutions that are not vulnerable to regular electoral and legislative checks, making the reversal impossible to challenge due to a lack of proponents possessing the power to impose a cost for its violation.

The analysis included in the judicial capture section above is directly relevant to the status of women's rights in the United States as a whole because the Dobbs case's narrowing of substantive due process constitutes a norm-cascade reversal that reaches beyond the singular ruling. The same history-tradition requirement used to remove reproductive rights legislation threatens other rights that are recognised under the same doctrine. The reversal should, therefore, not be looked at as an isolated rights loss, but rather as damage to the foundations behind twentieth-century rights protections. In other words, the American case carries a risk that is not present in the Turkish withdrawal: the altering of the entire methodology through which rights protections are interpreted across domestic law will have significant effects on future domestic rights claims.

Finally, Hudson et al.'s assertion (1-10) that the physical and legal status of women functions as an indicator of democratic health (political stability and governance quality) allows backsliding severity in the United States to be measured through a gender-rights lens. Turkey's trajectory, spanning from the 2010 referendum's instrumentalisation of rights expansion and gender-focused language, up until the judicial purges still occurring today, exemplifies a degree of backsliding that is extensive, both internationally and domestically. The American trajectory,

whilst severe, is much more contained. Returning authority of reproductive rights regulation to individual states did not eliminate protection by amendment, but instead narrowed rights interpretations for future claims – a wholly domestic shift. Overall, the Hudson et al. diagnostic framing supports the underlying argument of this analysis, demonstrating that the erosion of women's rights is not merely a symptom of democratic backsliding but an easily comparable measure of harshness across seemingly disparate cases.

Conclusion

This review set out to determine the extent to which the personalisation of executive power under Trump and Erdoğan has driven democratic backsliding through judicial capture, using the erosion of women's rights as a mechanism of this process. The comparison suggests that executive personalisation can contribute to democratic erosion by weakening the institutional constraints that limit political leaders, but that the mechanisms and consequences of this process vary considerably across political systems.

The Turkish case constitutes the chain at its most complete with the 2017 constitutional referendum transforming Erdoğan and the AKP's informal dominance into a formal presidential system; the post-2016 purges of the judiciary representing immediate instances of judicial capture rather than gradual reshaping; and the instrumentalisation of women's rights language in the 2010 referendum before the culmination of rights erosion with the 2021 Convention withdrawal, showing personalisation operating through gender politics. The United States demonstrates the same underlying political dynamics but to a less far-reaching extent: personalisation without constitutional amendment through erosion of restraint, appointment-process manipulation and gradual judicial capture, and the narrower norm-cascade reversal of the Dobbs case standing parallel to Turkey's withdrawal.

This gender politics comparison, measured through Hudson et al.'s diagnostic framing (1-10), does not merely show both cases exhibiting some level of rights protection rollback and democratic backsliding, but instead it orders their severity. In the case of gender rights protection being treated as one affected policy area among others, one is only able to see the presence of democratic decline. Rather than an incidental example, this analysis offers the erosion of women's rights as a measurement tool for democratic health, as validated by the 2012 study.

The ordering of severity also encompasses this paper's intended contribution to comparative politics as a whole. The pairing of Turkey and the United States demonstrates convergence on the same outcome despite divergence on historical trajectory. Ultimately, where a most-similar-systems design would isolate a single point of difference, this analysis treats the sheer dissimilarity of the two cases as its own test: a shared impact chain producing similar outcomes (granted, to lesser severity in the case of the US), regardless of the absence of overlapping background conditions, is stronger evidence of mechanism presence. Furthermore, this allows the analysis to extend specific concepts rather than merely apply them in the referenced context. For example, Bermeo's (10-13) executive aggrandisement and Poguntke and Webb's presidentialisation thesis, developed for hybrid or parliamentary regimes, demonstrate

applicability in a consolidated democracy, provided they operate through more informal registers (4-11).

One limitation of this study must be acknowledged in that this analysis has, at times, depicted the judiciary as a passive site into which power is relocated. This is not entirely accurate. Granted, in the Turkish case, judicial autonomy is comparatively more subordinate to the preferences of the AKP and Erdoğan since demonstration of loyalty is a required condition for these positions. However, in the American case, judges are selected through the Federalist Society's vetting pipeline and are entirely willing to exercise the authority offered to them, as well as being permitted to maintain ideological positions and personal interpretations of justice. A further limitation is that this paper focuses primarily on institutional and elite-level processes. It therefore pays less attention to civil society, social movements, opposition parties and public opinion, all of which can constrain or facilitate democratic erosion.

Overall, the stakes of this analysis' framing extend beyond this paper: as it deems gender policy to be valid as a diagnosis for backsliding severity, the erosion of women's rights functions as an early warning signal available to the press, ordinary citizens, and democracy-monitoring organisations before more visible indicators of backsliding (e.g., election manipulation or press suppression) become visible.

.

Works Cited

Arat, Yeşim. 'Democratic Backsliding and the Instrumentalization of Women's Rights in Turkey'. Politics & Gender, vol. 18, no. 4, Aug. 2021, pp. 1–31, https://doi.org/10.1017/s1743923x21000192.

Bermeo, Nancy. 'On Democratic Backsliding'. Journal of Democracy, vol. 27, no. 1, 2016, pp. 5–19, https://doi.org/10.1353/jod.2016.0012.

Connell, R. W. Masculinities. Polity, 2005.

Dahl, Robert A. Polyarchy: Participation and Opposition. Yale UP, 1971.

Elgie, Robert, and Gianluca Passarelli. 'Presidentialisation: One Term, Two Uses – Between Deductive Exercise and Grand Historical Narrative'. Political Studies Review, vol. 17, no. 2, 2019, https://doi.org/10.1177/1478929918754500. Accessed 6 Oct. 2026.

Esen, Berk. 'Judicial Transformation in a Competitive Authoritarian Regime: Evidence From the Turkish Case'. Law & Policy, vol. 47, no. 1, 2024, https://doi.org/10.1111/lapo.12250. Accessed 6 Oct. 2026.

Finnemore, Martha, and Kathryn Sikkink. 'International Norm Dynamics and Political Change'. International Organization, vol. 52, no. 4, 1998, pp. 887–917.

Gümüşçü, Berk Esen. 'Rising Competitive Authoritarianism in Turkey'. Third World Quarterly, vol. 37, no. 9, 2016, pp. 1581–606.

Hirschl, Ran. Towards Juristocracy: The Origins and Consequences of the New Constitutionalism. Harvard UP, 2004.

Hudson, Valerie M. Sex and World Peace. Columbia UP, 2012.

Inglehart, Ronald, and Pippa Norris. 'Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash'. HKS Faculty Research Working Paper Series, vols RWP16–026, 2016, https://doi.org/10.2139/ssrn.2818659. Accessed 6 Oct. 2026.

Johnson, Richard. 'Dobbs V. Jackson and the Revival of the States" Rights Constitution'. The Political Quarterly, vol. 93, no. 4, 2022, https://doi.org/10.1111/1467-923x.13193. Accessed 6 Oct. 2026.

Kandiyoti, Deniz. 'Locating the Politics of Gender: Patriarchy, Neo-Liberal Governance and Violence in Turkey'. Research and Policy on Turkey, vol. 1, no. 2, 2016, pp. 103–18, https://doi.org/10.1080/23760818.2016.1201242. Accessed 6 Oct. 2026.

Levitsky, Steven, and Daniel Ziblatt. How Democracies Die. Broadway Books, 2018.

McCoy, Jennifer. 'Polarization and the Global Crisis of Democracy: Common Patterns, Dynamics, and Pernicious Consequences for Democratic Polities'. American Behavioral Scientist, vol. 62, no. 1, 2018, pp. 16–42, journals.sagepub.com/doi/10.1177/0002764218759576. Accessed 6 Oct. 2026.

Müller, Jan-Werner. What Is Populism? Penguin Books, 2016.

Norris, Pippa, and Ronald Inglehart. Cultural Backlash: Trump, Brexit, and Authoritarian Populism. Cambridge UP, 2019.

Panke, Diana, and Ulrich Petersohn. ‘Norm Challenges and Norm Death: The Inexplicable?’. Cooperation and Conflict, vol. 51, no. 1, 2016, pp. 3–19, https://doi.org/10.1177/0010836715597948. Accessed 6 Oct. 2026.

Poguntke, Thomas, and Paul Webb. The Presidentialization of Politics: A Comparative Study of Modern Democracies. Oxford UP, 2005.

Przeworski, Adam, and Henry Teune. The Logic of Comparative Social Inquiry. Wiley-Interscience, 1971.

Roggeband, Conny, and Andrea Krizsán. ‘Democratic Backsliding and the Backlash Against Women’s Rights: Understanding the Current Challenges for Feminist Politics’. UN Women, 2020, www.unwomen.org/en/digital-library/publications/2020/06/discussion-paper-democratic-backsliding-and-the-backlash-against-womens-rights. Accessed 6 Oct. 2026.

Samuels, David J., and Matthew S. Shugart. Presidents, Parties, and Prime Ministers: How the Separation of Powers Affects Party Organization and Behavior. Cambridge UP, 2010.

Totenberg, Nina, and Sarah McCammon. ‘Supreme Court Overturns Roe V. Wade, Ending Right to Abortion Upheld for Decades’. NPR, 24 June 2022, www.npr.org/2022/06/24/1102305878/supreme-court-abortion-roe-v-wade-decision-overturn. Accessed 6 Oct. 2026.

Trump, Donald J. ‘Inaugural Address’.

‘Turkey Erdogan: Women Rise up Over Withdrawal From Istanbul Convention’. BBC News, 26 Mar. 2021, www.bbc.co.uk/news/world-europe-56516462. Accessed 6 Oct. 2026.

Methods in Ground-Effect Aerodynamics: A Chronological Review By Terumasa Iriyama

Abstract

Ground effect is a phenomenon in which lift slightly increases and drag decreases when an object moves close to the ground, making its prediction important during takeoff and landing. This review examines the methods used to analyze ground effect, beginning with hand calculation in the 1920s and continuing to emerging LLM-driven Computational Fluid Dynamics (CFD). Four historical periods and an emerging direction are compared, each evaluated by its dominant method, cost, range of testable conditions, and the technology that enabled the transition. Methodological transitions occurred when a new technology overcame the principal limitation of the preceding period. No method completely replaced its predecessor: modern CFD still relies on wind-tunnel testing for validation. The expected next transition is LLM-driven CFD case automation. Existing reviews of ground-effect aerodynamics address specific application domains, such as race cars and wing-in-ground-effect craft, or questions of simulation fidelity. To the author's knowledge, no prior review has been identified that provides a century-scale, primary-source chronology of how fixed-wing ground-effect analysis progressed across analytical theory, ground-simulation experiments, flight validation, CFD, and emerging LLM-driven workflows.

Nomenclature

$h = wing\ height\ above\ the\ ground$

$b = wingspan$

$S = wing\ area$

$AR = aspect\ ratio$

$C_L = lift\ coefficient$

$C_D = drag\ coefficient$

$Re = Reynolds\ number$

Source units are retained where needed:
1 ft = 0.3048 m; 1 ft² = 0.09290304 m²; 1 ft/s = 0.3048 m/s; 1 kn = 0.514444… m/s.

Introduction

This review presents an overview of the technological transformation in investigating ground effects over the past century. The 1920–1960 period relied on analytical solutions and wind-tunnel testing with small-scale models and mechanical measurements. Then, in the 1960–1990 period, thanks to the development of electrical instruments and a refined wind-tunnel setup, the moving-belt ground plane, flight testing and wind tunnel experiments validated each other, while expanded analytical solutions still remained in use alongside them. Subsequently, in the 1990–2010 period, improvements in hardware and software enabled CFD to be conducted affordably and with acceptable accuracy. This did not replace real-life experiments (Slotnick et al.). Finally, in the 2010–2023 period, CFD was used in design optimization and automation

frameworks.

In the 1920–1960 period, beginning in the 1920s of research on ground effect aerodynamics, it was common to rely on analytical solutions. From the 1910s, it was already known to some researchers that induced drag decreases and lift increases slightly when aircraft reach near the ground, but it was not until 1922 that Wieselsberger was able to calculate it. His method was the Method of Reflection, which replaces the ground with the reflected wing and applies Prandtl's biplane theory (Wieselsberger). This novel method was used as the foundation for an analytical method for the following 100 years. After that, research on the wind tunnel using a scaled model was introduced. The wind tunnel experiment in the 1920–1960 period was used to support expanded analytical solutions, such as applying the chord and thickness (Tani et al.).

An important parameter when comparing model-scale tests with full-scale flight is the Reynolds number, $Re = \frac{UL}{\nu}$, where U is the flow speed, L is a characteristic length such as wing chord, and ν is the kinematic viscosity. The Reynolds number measures the relative importance of inertial and viscous effects (NASA Glenn Research Center). Matching it helps preserve viscous similarity, but does not by itself match geometry, ground clearance, angle of attack, or boundary conditions. In the studies reviewed here, the model-scale ground-effect tests were run at Reynolds numbers of the order of 10^5–10^6 (7×10^5 in the descending-model wind-tunnel tests and 3.2×10^5 in the NACA 4412 airfoil-in-ground-effect CFD validation case), (Chang and Muirhead; Hsiun and Chen) whereas the flight tests were flown at full-scale landing conditions; (Baker et al.; Corda et al.; Curry) the Reynolds numbers compared in this review therefore range from model scale to full scale rather than taking a single value.

However, wind-tunnel testing was still limited by the mechanical measurement equipment of the era. Another problem was the boundary layer that develops on a fixed tunnel floor, which is a limitation Klemin first identified in 1934, proposing a moving-belt floor as the remedy (Klemin). Purpose-built aeronautical rigs followed in the early 1960s (Butler et al.) and Turner (1966, 1967) subsequently validated and documented the technique at NASA Langley for powered-lift configurations (Turner, “Endless-Belt”; Turner, “Moving-Belt”).

The 1970s, the start of the 1960–1990 period, saw a shift in the research method. Thanks to the introduction of electrical measurement equipment like the strain gauge, and the widespread use of the moving-belt floor, which became common in the late 1960s, wind tunnel testing became increasingly reliable. Simultaneously, flight testing was introduced into ground effect aerodynamics investigation. While these technological improvements made these methods quantitative and shifted the method of research, neither became a source beyond any doubt. While moving-model experiments, which move a model through still air, exposed the fixed-floor wind tunnel boundary-layer bias that Klemin previously presented and verified the reliability of the moving-belt floor wind tunnel, (Turner, “Endless-Belt”; Turner, “Moving-Belt”) instrumented flight tests provided a full-scale check on whether wind-tunnel and analytical predictions reproduced the magnitude of the ground effect (Baker et al.).
It was already common in the 1950s for airplanes to measure their flight characteristics using

onboard instruments (Sjoberg). However, it was not until this period that ground-based optical/radar tracking, synchronized to onboard recordings by a radio tone, made precise low-altitude trajectory measurement possible (Baker et al.).

While new methods emerged, traditional analytical solutions kept their place as the most reliable source, and some researchers were able to delve into an expanded analytical solution in three dimensions (Widnall and Barrows). The research center at Los Alamos introduced a completely new method in 1965: the 1960–1990 version of Computational Fluid Dynamics(CFD). This method solved the Navier-Stokes equations(NSE), which are the governing equations of fluid dynamics (Harlow and Welch). However, CFD in the 1960–1990 period was prohibitively slow for industrial use.

In front-line aeronautics, the practical method was the Panel Method in 1966 by Hess and Smith (Hess and Smith). Instead of directly solving the NSE, which took too much time on the computer specifications of that era, they distributed sources over the body surface and solved a boundary integral equation, so that potential flow about arbitrary bodies could be computed; ground effect can be represented in such methods by an image system of panels.

Although full-scale flight testing remained expensive, the end of the Cold War in 1989 weakened the basis of national security, which supported the broader aeronautic research in the US. Between 1994 and 2000, NASA's aeronautics budget declined by 45% (Ferguson). This austerity did not end fighter-based ground effect testing, with testing on the F-15 in 1994 and with the F-16XL later in the decade (Corda et al.; Curry). Indeed, flight testing became more selective, and wind tunnel testing became increasingly dominant during the transition to the next decade.

It was in the late 1990s, during the 1990–2010 period, that modern CFD was introduced into ground-effect aerodynamics research. It was a fittingly groundbreaking shift that CFD promised design-relevant answers at a fraction of the cost of flight testing, as Chapman had anticipated as early as 1979 (Chapman). This shift was enabled by advances in these two factors: hardware and software (Slotnick et al.). Chapman also framed the rise of CFD as a race between the demand for complex three-dimensional viscous flow simulation and the rapidly improving speed and memory of digital computers (Chapman).

On the hardware side, Moore reported in 1965 that exponential growth in realistic integrated-circuit complexity, observing that the number of components per circuit doubled every year, approximately (Moore). Denser and cheaper processors and memories integrated with vector supercomputers made large structured-grid flow solvers practical, while later massively parallel systems expanded CFD from isolated calculations to a progressively larger and faster process (Graham et al.).

On the software side, turbulence models helped make engineering calculations of turbulent flow affordable. In Reynolds-averaged Navier–Stokes (RANS) simulations, these models approximate the effects of turbulent fluctuations on the mean flow. The k–ε and k–ω models each solve two additional transport equations: one for the turbulent kinetic energy, k, and another for its dissipation rate, ε, or specific dissipation rate, ω. Menter's shear stress transport

(SST) model blends a k–ω formulation near walls with a transformed k–ε formulation away from walls (Menter; Launder and Spalding). Separately, object-oriented software frameworks such as FOAM supported modular implementation of physical models and numerical methods (Weller et al.).

Retrospective assessments such as NASA's CFD Vision 2030 study attribute modern capability to both enormous hardware growth and algorithmic gains of comparable or greater effect, while retaining ground and flight tests as complementary sources of validation and risk reduction (Slotnick et al.).

While the introduction of modern CFD was revolutionary, its realistic use was to calculate a single case due to time and computational resource limitations. CFD didn't completely replace flight testing and wind tunnel testing, either. In validated regimes, CFD can now provide design-relevant information comparable to wind-tunnel or flight testing and often offers much richer spatial detail; however, experiments remain indispensable for establishing credibility in separated, transitional, off-design, and otherwise weakly validated flows (Slotnick et al.).

Finally, in the late 2010s, during the 2010–2023 period, computational and framework improvements enabled CFD to be utilized in design optimization. For example, in 2023, Abu Salem et al, presented a framework for Aircraft Take-Off Considering Ground Effect Aerodynamics in Conceptual Design (Abu Salem et al.). Another example is Jesudasan's 2023 research, which provided a wing planform optimization, incorporating aerodynamics into an extensive framework including stability (Jesudasan et al.).

However, CFD is still considered to be part of a complicated engineering process, despite commercial CFD vendors having invested in usability (Pandey et al.). Engineering projects requiring ground-effect aerodynamics still typically require specialist CFD expertise. Recent work has shown that large language model (LLM) agents can construct and iteratively repair OpenFOAM cases using retrieval-augmented generation (RAG) and execution feedback (Pandey et al.; Chen et al.). MetaOpenFOAM demonstrated a multi-agent natural-language-to-CFD workflow, (Chen et al.) while the peer-reviewed OpenFOAMGPT extended this approach across case setup, boundary-condition and turbulence-model modification, and diverse RANS/LES applications (Pandey et al.).

While fields outside ground-effect aerodynamics are developing LLM-driven CFD, no published study appears to have applied LLM-driven CFD to ground effect with experimental validation. Following the previous century's technological transformation, it is highly likely that the LLM-driven methods still have to be verified with the previous methods presented in this paper.

Existing reviews of ground-effect aerodynamics are organized primarily by application domain: race cars (Zhang et al.), wing-in-ground-effect craft (Rozhdestvensky; Cui and Zhang), powered-lift and rotorcraft vehicles (Matus-Vargas et al.), low-Reynolds-number experimental wings (Lee and Lin), compressible flow (Doig), or by questions of simulation fidelity (Barber et al.; Barber). To the best of our knowledge, none provides a century-scale, primary-source

chronology of how fixed-wing ground-effect analysis progressed across analytical theory, ground-simulation experiments, flight validation, CFD, and emerging data-driven or LLM-assisted workflows. Early surveys exist (Pistolesi; Foshag) but predate the methods reviewed here. This paper concludes that every time a new method is introduced, it does not replace the existing method but interdependently verifies and improves accountability.

Methodology

This narrative review compares selected publications by resource requirements, testable conditions, enabling technology, and validation evidence. It is not a systematic census of the literature: no reproducible search-and-screening record or adoption-frequency dataset is available. Numerical comparisons retain each source's configuration, height definition, and normalization; percentages and coefficient increments are not pooled. Values read approximately from published graphs are identified as such, and no new regression, hypothesis test, or statistical meta-analysis is performed. Cost categories are qualitative judgments based on the equipment and computational resources described, not measured expenditure.

A) Chronology of selected methods

Table 1 below compares the methods used across the review's historical periods. These historical periods are chosen based on the appearance of new methods; around 1960, wind-tunnel testing gained a practical moving-belt ground plane, and flight testing gained popularity with the development of measurement equipment; around 1990, CFD moved from a research tool to an affordable, industry-ready solver; around 2010, hardware and software developments made it realistic to utilize CFD in optimization cases; and since 2024, LLM coding agents have become publicly available and capable of generating CFD cases; this review treats that development as an emerging direction rather than a fifth period.

Table 1. Representative studies within the review periods.

Method	1920–1960	1960–1990	1990–2010	2010–2023	2024– (outlook)
Analytical solutions	(Wieselsberger; Tani et al.)	(Widnall and Barrows)	—	—	—
Wind tunnel experiment	(Wieselsberger; Tani et al.)	(Chang and Muirhead; Butler et al.; Turner, "Endless-Belt"; Turner, "Moving-Belt")	Validation. (Hsiun and Chen)	—	—
Flight testing	General flight. (Sjoberg)	Ground effect. (Baker et al.)	Ground effect. (Corda et al.; Curry)	—	—
CFD	—	MAC/panels. (Harlow and Welch; Hess and Smith)	RANS; FOAM. (Hsiun and Chen; Weller et al.; Menter; Launder and Spalding)	Design uses. (Abu Salem et al.; Jesudasan et al.)	—
Computational optimization	—	Adjoint. (Jameson)	—	(Abu Salem et al.; Jesudasan et al.)	ML surrogate. (Narayanan et al.)
LLM-driven CFD	—	—	—	—	(Pandey et al.; Chen et al.)

Entries identify examples discussed in this review, not prevalence or dates of invention. A dash means no example is assigned to that cell, not that the method was unavailable. CFD = computational fluid dynamics; LLM = large language model. Some general numerical-method studies provide background rather than ground-effect validation.

A-1 Analytical methods and verification method proposals, 1920–1960

The equation below estimates the change in induced drag for an idealized finite wing, with ΔC_D as the induced-drag coefficient increment and σ as the influence coefficient (Wieselsberger).

$$\Delta C_D = \frac{-\sigma \cdot C_L^2}{\pi \cdot AR}(1) \text{ (Wieselsberger).}$$

$$\sigma = \frac{1-1.32\cdot(h/b)}{1.05+7.4\cdot(h/b)} \quad \text{(Applicable range: } \frac{1}{30} \leq \frac{h}{b} \leq \frac{1}{4}\text{) (Wieselsberger).}$$

Wieselsberger experimented using his model (σ=0.432), which gives ΔC_D = -0.015 C_L^2, matching the measured polar up to $C_L \approx 1$. This correction applies to the idealized induced-drag term only; it assumes lift is unchanged and excludes profile drag, separation, and the moving-ground boundary layer.

For h defined here as height above the ground, the stated range is 1/30 ≤ h/b ≤ 1/4 (Wieselsberger).

Wind tunnels measure aerodynamic forces on a scaled model in a controlled airstream (Turner, “Endless-Belt”). Wieselsberger compared his prediction with model measurements (Wieselsberger). A fixed floor introduces a boundary layer, but the direction and magnitude of its effect depend on configuration and ground clearance; Turner's large lift-loss bias concerned a powered-lift model (Turner, “Endless-Belt”).

The correction requires little computation, but represents an idealized induced-drag effect (Wieselsberger).

A-2 Experimental and numerical developments, 1960–1990

Circa 1970–1990, research in the 1960–1990 period shifted toward flight testing and wind tunnel testing with ground-effect-dedicated machinery.

In this period, the dominant methods incurred substantially higher costs than the analytical era: flight testing required instrumented full-scale aircraft and ground-based tracking, while moving-belt tunnels required purpose-built rigs (Butler et al.; Turner, “Endless-Belt”).

A-3 Flight validation and CFD development, 1990–2010

During the 1990–2010 period, research saw a transformation in which CFD increased its dominance while flight testing declined in use as budgets contracted, (Ferguson) and wind tunnel testing kept its role.

A-4 Computational design applications, 2010–2023

During the 2010–2023 period, research continued to evolve from advances made in the 1990–2010 period. The limitation previously preventing this technology from spreading was that single-case CFD could not search a design space; adjoint methods return gradients at a cost of roughly two flow solutions per iteration, independent of the number of design variables (Jameson).

A-5 LLM-driven CFD, 2024–2026

From 2024, researchers started to focus on using LLMs to build and fix CFD cases (Pandey et al.; Chen et al.).

B) Ground simulation in wind tunnels

B-1 Ground Representation in Wind Tunnels

Table 2 features four types of ground representation and their performance against ground-effect testing.

Table 2. Wind-tunnel ground representations used for ground-effect research.

Wind-tunnel testing type	Ground boundary layer	Testable range for ground effect	Introduction and validation
Fixed ground board (Turner, "Endless-Belt")	Board develops its own boundary layer, which separates ahead of high-lift models and shifts the stagnation streamline.	Conventional wings at moderate heights; overstates lift loss close to the ground for high-lift models.	Moving-belt remedy proposed by Klemin, 1934. (Klemin)
Moving-model carriage (Turner, "Endless-Belt")	Model moves through initially still air; avoids a tunnel-floor boundary layer.	Boundary-layer-free reference used to validate the other representations.	Used by Turner as the reference technique, 1966.
Moving-belt ground plane (Turner, "Endless-Belt")	Belt motion and upstream suction reduce the floor boundary layer.	Closer to the moving-model reference for the tested high-lift configuration.	Proposed 1934; (Klemin) RAE rig、1963; (Butler et al.) validated at Langley、1966.
Dynamic (constant sink-rate) testing (Chang and Muirhead)	Fixed ground board, with the model lowered toward it at constant sink rate.	Reproduces the descent itself; captures the sink-rate dependence that static tests miss.	Chang and Muirhead, 1987; dynamic XB-70 data matched flight.

RAE = Royal Aircraft Establishment (Butler et al.).

General wind-tunnel classifications (circuit type, test-section type, speed regime) are independent of the ground boundary-layer problem, so this review classifies tunnels by their

ground representation and model motion. These representations accumulated rather than replaced one another: the moving belt has remained the standard ground representation since its validation in the 1960s, (Butler et al.; Turner, "Endless-Belt") and the change after 1990 was in the tunnel's role, from primary measurement tool to the validation reference for CFD, not in its hardware.

B-2 Fixed-floor bias

The table below shows a difference in characteristics between the fixed-floor(mainly in the 1920–1960 period) wind tunnels and the moving-belt-floor(introduced in the 1960–1990 period) wind tunnels.

Table 3 summarizes the lift-loss comparison for the four ground representations.

Table 3. Fractional lift loss for Turner's full-span blown-flap model under four ground representations.

Ground representation	$\frac{\Delta L_h}{L_{h=\infty}}$ **at h/b ≈ 0.08**	$\frac{\Delta L_h}{L_{h=\infty}}$ **at h/b ≈ 0.25**	**Height where plotted loss approaches zero**
Moving model (reference)	≈ −0.45	≈ 0	h/b ≈ 0.25
Fixed ground board	≈ −0.72	≈ −0.33	h/b ≈ 0.6
Belt, $V_{belt} = V_{air}$	≈ −0.36	≈ −0.10	h/b ≈ 0.35
Belt, $V_{belt} = 0$	≈ −0.57	≈ −0.33	h/b ≈ 0.5

Above in Table 3, h is ground clearance, b is wingspan, belt speed is either zero or equal to the airspeed, and aspect ratio is 6, out-of-ground-effect lift coefficient is 9.5. Values are read from Turner's Fig 6 (Turner, "Endless-Belt").

The limitation of ground-effect wind-tunnel testing in the 1920–1960 period was not the balance or the model but the floor. A stationary ground board develops its own boundary layer, and for configurations that push high-energy flow toward the ground, this low-energy layer separates upstream of the model and relocates the stagnation streamline, altering the measured forces (Turner, "Endless-Belt"). Klemin identified the problem and proposed a moving belt in 1934, (Klemin) but the first quantitative demonstration of its magnitude came from Turner's comparison at NASA Langley. The same full-span blowing-flap model (aspect ratio 6, C_L = 9.5 out of ground effect) was tested by two techniques: driven through still air on a carriage, which produces no floor boundary layer and therefore represents the true ground effect, and mounted over a conventional fixed ground board in the tunnel. Read from Turner's Fig 6, the fractional lift loss $\frac{\Delta L_h}{L_{h=\infty}}$ at h/b ≈ 0.08 was approximately −0.45 for the moving model but approximately −0.72 over the fixed board; at h/b ≈ 0.25 the moving model showed essentially zero loss, whereas the fixed board still indicated a loss of roughly 0.33. The fixed floor thus overstated the lift loss by about 60% close to the ground and reported a ground effect at heights where none existed (Turner, "Endless-Belt").

B-3 Moving-belt ground plane [1960–1990]

The endless-belt ground plane reduced this bias. Turner's rig consisted of a 10 ft × 12 ft belt of plastic-impregnated woven wool driven at 0–100 ft/s, with a 1-inch suction slot at the leading edge to remove the upstream boundary layer, and a model carried on an internal strain-gauge balance on a remotely driven sting. When the same model was retested over the belt (Turner's Fig 6), the belt running at free-stream velocity reproduced the moving-model curve, while the stopped belt reproduced the fixed-board curve. Belt speed proved forgiving: the lift loss varied almost linearly with the belt-to-air velocity ratio, so extreme precision in belt speed was not required (Turner, "Endless-Belt"). Two qualifications apply. These magnitudes are specific to a high-circulation powered-lift configuration; Turner reported that the belt is required primarily for such configurations, or for any model at sufficiently small h/b, and not for conventional wings at moderate heights (Turner, "Endless-Belt"). Moreover, a similar installation had already been built at the RAE; (Butler et al.) Langley's contribution was the validation against a boundary-layer-free reference, not the invention.

In terms of the evaluation criteria used throughout this review, the moving belt raised the cost of ground-effect testing (a dedicated rig, suction system, and belt drive) but extended the testable range to high-lift, low-clearance conditions where the fixed floor was demonstrably wrong; the enabling technologies were the belt mechanism itself and the strain-gauge balance that allowed remote height sweeps (Turner, "Endless-Belt").

B-4 Dynamic (sink-rate) testing [1980s-2010]

Static ground representations, including the moving belt, hold the model at a fixed height, so they cannot capture the sink-rate dependence that flight testing had revealed (Section D). Chang and Muirhead reproduced the descent itself inside the tunnel: 60, 70, and 75 degree delta wings, an XB-70 wing, and an F-104A wing were lowered toward the ground board at constant sink rate in the University of Kansas tunnel at $Re = 7 \times 10^5$, with static runs alongside for comparison. The dynamic XB-70 data agreed closely with flight, while the static data overpredicted the lift increment (Chang and Muirhead).

In terms of the evaluation criteria, dynamic testing added little hardware cost, a drive to lower the model, but extended the testable range from steady heights to the landing maneuver itself; the enabling technology was the constant-sink-rate rig, and its accuracy was established against flight data (Section D, Results R-3).

C) Computational methods and their validation

C-1 Summary of CFD technological features

Below is the table describing each solver for CFD and its performance in the context of ground effect research, focusing on treatment of the governing equations, turbulence closure, assumed hardware, and representative ground-effect application (Harlow and Welch; Hess and

Smith; Moore; Graham et al.; Weller et al.; Menter; Launder and Spalding; Pandey et al.; OpenFOAM Foundation).

Table 4 compares the solver classes discussed below.

Table 4. Computational fluid dynamics (CFD) solver classes and their characteristics.

Solver class	**Governing equations**	**Turbulence closure**	**Hardware assumed**	**Representative ground-effect application**
Marker-and-cell. (Harlow and Welch)	Time-dependent incompressible NSE, finite-difference scheme.	None; Resolved viscous flow.	1960s computing; no comparable cost series.	None directly; foundation for later NSE-based solvers
Panel method. (Hess and Smith)	Incompressible potential flow using surface singularities.	None; inviscid flow.	Numerical surface discretization; no common cost comparison.	No ground-effect validation established by this cited example.
RANS with two-equation closure. (Menter; Launder and Spalding)	Reynolds-averaged NSE.	k-ε; k-ω SST.	Supercomputers, later commodity clusters. (Moore; Graham et al.)	NACA 4412 ground-effect comparison in R-4. (Hsiun and Chen)
Open-source object-oriented implementation. (Weller et al.; OpenFOAM Foundation)	Finite-volume NSE in a modifiable solver framework.	User-selectable closure libraries.	Commodity parallel hardware. (Graham et al.)	Modular framework; validation of a GE case remains separate.
LLM-driven CFD (Pandey et al.)	Not a new equation set; inherits the governing equations of the selected OpenFOAM solver (tested across laminar/turbulent, single-/multiphase, and heat-transfer cases).	No fixed closure; the agent selected or changed OpenFOAM RANS/LES models, including k-ε, RNG k-ε, k-ω SST, k–kl–ω, LRR, and Smagorinsky.	GPT-4o or o1-preview via the OpenAI API, a RAG database, and a machine running OpenFOAM; no dedicated CFD hardware was specified.	General tutorial and code-translation tests; no GE experimental validation.

Description for characters and abbreviations above: NSE, Navier–Stokes equations; RANS, Reynolds-averaged Navier–Stokes; SST, shear-stress transport; LLM, large language model; k, turbulent kinetic energy; ε, dissipation rate; ω, specific dissipation rate.

C-2 Marker-and-cell Method

Unlike compressible flow, which has a continuity equation for density, in incompressible flow there is no evolution equation for pressure.

An innovative feature of the MAC method was setting the main variables as pressure and two-dimensional velocity components instead of vorticity and the stream function, by staggering the pressure in the middle of the cell and velocity on the boundary of a cell (Harlow and Welch).

Harlow and Welch compared their free-surface calculation with liquid-column-collapse measurements by Martin and Moyce (Harlow and Welch; Martin and Moyce). That comparison is not a validation of an aircraft ground-effect calculation.

C-3 Panel Method

For eight Polaris-missile hydrodynamic coefficients, the panel method gives an average absolute error of 4.9% with a maximum individual error of 7.3%, against 11.7% and 23.5% for the conventional calculation (Hess and Smith).

C-4 RANS with two-equation closure

Launder and Spalding presented a two-equation turbulence-model formulation in 1974, (Launder and Spalding) and Menter introduced BSL and SST eddy-viscosity models in 1994 (Menter). These papers document particular developments rather than the first existence of all two-equation closures.

Instead of relying on a single-equation model, Launder and Spalding's method solved the transport equations for both turbulent kinetic energy(k) and dissipation rate(ε). These validations should not be treated as one study or as direct ground-effect tests.

C-5 OpenFOAM

OpenFOAM 1.0 was released in 2004 (OpenFOAM Foundation).

Another problem was that solvers were a commercial product, and approximately only 70% of the code was available to the public (Weller et al.).

The novelty lies in the code design rather than the equations themselves.

One test case was incompressible flow around a square cylinder, using Reynolds-averaged and large-eddy simulations (LES).

C-6 OpenFOAMGPT and LLM-driven CFD

OpenFOAMGPT is a 2025 research project aiming to automate CFD cases utilizing GPT-4o and o1-preview (Pandey et al.).

General-purpose LLMs lacked the solver-specific knowledge required to construct reliable OpenFOAM cases. OpenFOAMGPT addressed this gap by retrieving existing simulation setups and supplying them as domain-specific context, enabling the agent to generate more accurate and specialized configurations without retraining the foundation model (Pandey et al.).

A self-correcting loop operates through the coordination of three roles, which are Interpreter, Builder, and Runner, which monitor solver error logs and append details to the query for subsequent iterations upon failure. Domain knowledge is injected using RAG (LangChain, text-embedding-3-small, 1536 dimensions) (Pandey et al.).

The six case-generation demonstrations and four code-translation tests should be distinguished. In Table V, PitzDaily translation failed with GPT-4o but succeeded with o1-preview, whereas Motorbike failed with both (Pandey et al.). These were workflow failures and should not all be described as failures of numerical convergence.

D) Unique Case Studies on Ground Effect

This part focuses on the same aerodynamic tasks being evaluated by multiple methods that appeared in the A-C part. Fig 1 describes each aircraft used in the experiment.

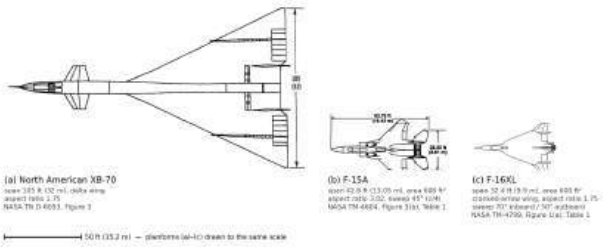

Fig 1: Wing planforms of the XB-70, F-15A, and F-16XL at a common scale; Dimensions and aspect ratios are shown for comparison. Scale bar is 50 ft, or 15.2 m. Source planforms are from Baker et al., Corda et al., and Curry (Baker et al.; Corda et al.; Curry). (a) North American XB-70: delta wing, span 105 ft (32 m), aspect ratio 1.75 (Baker et al.). (b) F-15A: span 42.8 ft (13.05 m), wing area 608 ft², aspect ratio 3.02, quarter-chord sweep 45° (Corda et al.). (c) F-16XL: cranked-arrow wing, span 32.4 ft (9.9 m), reference area 600 ft², aspect ratio 1.75, leading-edge sweep 70° inboard and 50° outboard (Curry).

Drawings reproduced from NASA TN D-6053 (Fig 3), NASA TM-4604 (Fig 3(a)) and NASA TM-4799 (Fig 1(a)). Dimension annotations in (b) are those of the original drawing (overall length and stabilator span); the F-15 wingspan is 42.83 ft (Corda *et al.*, Table 1) (Corda et al.).

D-1 Flight testing method transformation

There are two generations of methods conducted in the flight testing for ground effect evaluation.

This introduction significantly decreased the number of flight attempts in flight testing; for example, constant-angle-of-attack data were obtained during 5 runs, whereas the constant-altitude-flyby data required 44 runs (Baker et al.).

Another transformation in 1997 addressed the following problem: Optical tracking involved a heavy burden of setup and post-processing.

In addition, altitude data cannot be obtained from the aerodynamic measurement system because

ground effect itself distorts the barometric altitude readings from the nose boom (F-15 ≈6.5 ft, (Corda et al.) F-16XL ≈25 ft).

The 1970 method, constant-angle-of-attack approach, was validated by comparison with the previous constant-altitude-flyby approach. The precision was within $\pm$1.5 ft (Baker et al.). Curry checked DGPS height against known ground-clearance conditions; this was not a direct comparison with optical tracking. The precision for DGPS was within 1 ft (Curry).

D-2 XB-70: gap between static experiment and dynamic experiment

Baker *et al.* found agreement in trends but discrepancies in the magnitudes of lift and pitching-moment changes between flight, theory, and static wind-tunnel results (Baker et al.). Chang and Muirhead later measured the effects of controlled descent in a wind tunnel (Chang and Muirhead). Their result identifies a relevant variable; it does not establish descent as the sole cause of every earlier discrepancy.

D-3 F-15 and F-16XL: sink rate and aspect ratio influence on dynamic ground effect

The F-15 (aspect ratio 3.02, conventional fighter wing) and the F-16XL (aspect ratio 1.75, cranked-arrow delta wing) are both fighter jet aircraft tested at NASA facilities. In this section, the ground effect is quantified by the increment in lift coefficient between flight near the ground and flight out of ground effect, ΔC_L; a larger increment means a stronger ground effect. The sink rate, $\dot{h}$, is the vertical speed of descent, reported in the original studies in ft/s (1 ft/s = 0.3048 m/s).

To solve this issue, Corda, Stephenson, Burcham, and Curry from the NASA team investigated using an F-15 aircraft and flew 24 landings on 7 test flights at sink rates from 0.7 to 6.5 ft/s. The correlations fitted to the landing data (Corda *et al.*, Fig 11) are given below, with sink rate inserted numerically in ft/s; the fitted constants therefore apply to that unit convention. Here h' represents the sink rate, $\Delta C_{L,\,GE}$ represents the difference in lift coefficient, and $\Delta C_{D,\,GE}$ represents the difference in drag coefficient under ground effect, respectively: (Corda et al.)

$$\Delta C_{L,\,GE} = \frac{0.2}{1+h'} + 0.02 \text{ (Corda et al.).}$$

$$\Delta C_{D,\,GE} = 0.035 - 0.005h' \text{ (Corda et al.).}$$

Curry built another engineering model in 1997 using 24 F-16XL maneuvers over seven flights. He then applied the model to previously published F-15 data as well (Curry). In the equation below, $\gamma[rad]$ represents the flightpath angle (negative in descent) and $C_{L,\,OGE}$ represents the lift coefficient without the ground effect,

$$\frac{\Delta C_{L\,ge,\,dynamic}}{\Delta C_{L\,ge,\,steady-state}} = 1 + \frac{2\pi b^2 \gamma}{C_{L,\,OGE} S} \text{ (Curry).}$$

In the same F-15 study, Corda et al. found that the obtained F-15 data align with the correlation curve between the AR and the ground effect derived from the previous data that contain numerous aircraft that range from AR of 1.7 to 4.00, including the F-106B, the F-104A,

the F5D-1 and its ogee-wing modification, the XB-70 from the earlier NASA flight tests, and the forward-swept-wing X-29A (Corda et al., Table 4). For the aircraft correlation at the reported h/b = 0.3, the equation below gives the increment normalized by out-of-ground-effect lift; multiplying the ratio by 100 gives percent (Corda et al.).

$$\frac{\Delta C_{L,\,GE}}{C_{L,OGE}} \approx \frac{0.2}{AR}$$ (Corda et al.).

Results and Discussion

This section focuses on the evaluation between the methods and provides this review's own assessment about how the methods introduced earlier in this review improved ground-effect aerodynamics research.

R-1 Evaluation of methods against the review criteria

Table 5 summarizes the methods on four criteria: cost, testable range, enabling technology, and what validated the method.

Table 5. Evaluation of ground-effect methods against the review criteria.

Method	Cost	Testable range	Enabling technology	Validated by
Analytical (Wieselsberger)	Minimal	Idealized, then expanded to lift and pressure distribution	Prandtl biplane theory.	Wind tunnel. (Tani et al.)
Fixed-floor wind tunnel	Moderate	Static; moderate h/b Model-sized	Mechanical, then strain-gauge balances	Moving-model tests. (Turner, "Endless-Belt")
Moving-belt wind tunnel (Turner, "Endless-Belt")	High	High-lift, low h/b Model-sized	Endless belt, suction slot. (Butler et al.; Turner, "Moving-Belt")	Moving-model tests.
Dynamic wind tunnel (Chang and Muirhead)	Moderate~High	descending at fixed sink rate, Model-sized	sting on linear bearings, spring-loaded door.	XB-70 flight comparison.
Flight testing (Baker et al.)	Highest	Full-scale, dynamic	Cinetheodolite → DGPS; tone synchronization. (Curry)	Full-scale check on the other methods.
CFD (RANS)	Case-dependent	Validated regimes	Hardware、turbulence models, OpenFOAM. (Moore; Graham et al.; Weller et al.; Menter; Launder and Spalding; OpenFOAM Foundation)	Wind tunnel and flight. (Slotnick et al.)
CFD-based optimization (Abu Salem et al.; Jesudasan et al.)	Moderate	Design-space sweeps	VLM-based frameworks; optimization algorithms.	CFD and experiment.
LLM-driven CFD (Pandey et al.; Chen et al.)	Not established	Not established for ground effect	LLM agents.	Reference-CFD checks; no GE experimental validation.

R-2 Fidelity of ground simulation in wind tunnels

Fig 2 compares the four ground representations discussed in this section.

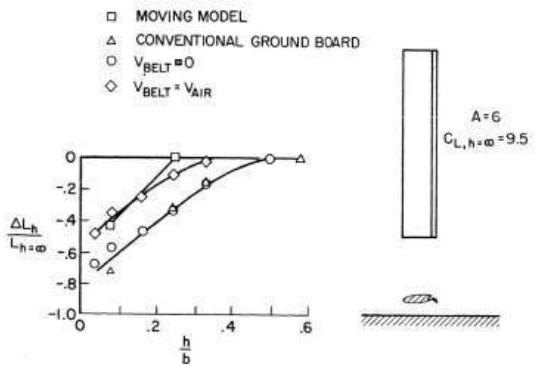

Fig 2: Turner's fractional lift loss for four ground representations.

Here, h is ground clearance, b is wingspan, ΔLh/L∞ is the lift change relative to out-of-ground-effect lift, Vbelt and Vair are belt and air speeds, aspect ratio 6; and the out-of-ground-effect lift coefficient 9.5; directly cited from Turner, Fig 6 (Turner, "Endless-Belt"). This section evaluates the influence of the moving belt on the wind tunnel testing, referring to Table 3 in section B-2.

Ground effect can reduce lift in powered-lift configurations (Turner, "Endless-Belt").

Among the tested tunnel representations, the belt moving at airspeed was closest to the moving-model reference (Turner, "Endless-Belt"). Residual differences and the high-lift configuration limit the scope of that conclusion.

R-3 Dynamic ground effect revealed by flight testing

Table 6 summarizes the study comparisons, and Fig 3 presents the flight-data comparisons.

Table 6. Comparisons revealing dynamic ground effect and informing its prediction.

Study	**Comparison**	**Main finding**
Baker *et al.* (1970). (Baker et al.)	Flight measurements compared with theoretical and static wind-tunnel predictions.	Trends agreed, but the magnitudes of changes in lift and pitching moment differed.
Chang and Muirhead (1987). (Chang and Muirhead)	Static and descending-model wind-tunnel measurements, with comparison to published flight data.	Ground-effect increments depended on sink rate. Dynamic XB-70-wing results agreed closely with flight data at an angle of attack of 9.3°.
Corda *et al.* (1994). (Corda et al.)	F-15 flight measurements at different sink rates; a separate comparison with static DATCOM predictions.	Lift increments increased as sink rate decreased. The static DATCOM prediction exceeded the flight-measured increment in the reported comparison. These results informed revision of the flight-simulator model.
Curry (1997). (Curry)	F-16XL descending-flight measurements compared with steady-state wind-tunnel predictions; engineering-model estimates compared with flight data.	Dynamic lift and drag increments were below steady-state predictions. A flightpath-angle correction provided an engineering estimate of the static–dynamic relationship for the F-16XL and F-15 data.

DATCOM = U.S. Air Force Stability and Control Data Compendium (McDonnell Douglas Corporation).

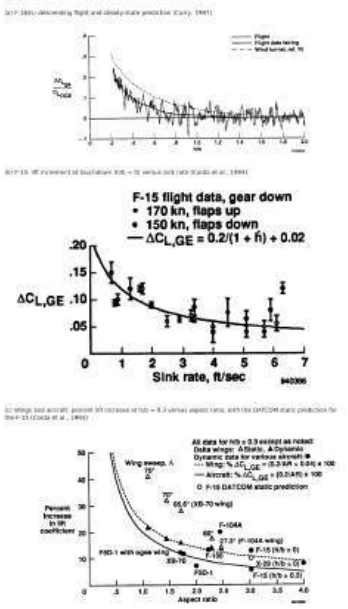

Fig 3: Dynamic ground effect in flight.

(a) normalized lift increment versus height-to-span ratio for the F-16XL, with flight data, its fairing, and a steady-state wind-tunnel prediction; (Curry)

(b) lift-coefficient increment versus sink rate for the F-15 at 170 kn with flaps up and 150 kn with flaps down, with error bars representing curve-fit accuracy and the original correlation curve (Corda et al.).

(c) percent increase in lift coefficient at h/b = 0.3 versus aspect ratio for the wings and aircraft compiled by Corda *et al.*, with their wing and aircraft correlation curves and the DATCOM static prediction for the F-15 (Corda et al.; McDonnell Douglas Corporation).

In (a), $\Delta C_{L,GE} / C_{L,OGE}$ is the lift increment divided by the out-of-ground-effect lift coefficient and h/b is ground clearance divided by wingspan.

In (b), $\Delta C_{L,GE}$ is the coefficient increment, and sink rate is positive downward (ft/s). The source touchdown label h/b = 0 is relative to touchdown; the F-15 wing remained approximately 9 ft above the ground (Corda et al.).

In (c), open triangles are static and filled triangles are dynamic wing data, filled circles are dynamic aircraft data, and the open circle is the DATCOM static prediction for the F-15; panels are reproduced from Curry, Fig 7(a), Corda *et al.*, Fig 11(a), and Corda *et al.*, Fig 12(a), respectively. The “ref. 15” label inside panel (a) belongs to Curry’s original report (Corda et al.; Curry).

From Fig 3 above, it can be read that at h/b = 0.3, the flight value is 5.6%, compared with approximately 10% from the static DATCOM prediction reported by Corda *et al.* The static and dynamic delta-wing points in panel (c) also differ, but do not establish a universal factor of two.

Simultaneously, $\Delta C_{L,\,GE}$, which represents the difference in lift coefficient, increased from approximately 0.05 to 0.10 as sink rate decreased from 6.5 ft/s to 0.7 ft/s in the reported flight results (Corda et al.).

Flight and descending-model experiments both show that sink rate affects ground-effect increments under the studied conditions (Chang and Muirhead; Corda et al.). These comparisons do not test steady RANS directly, so no RANS error direction is inferred from them.

R-4 Predictive capability from analytical theory to CFD

Fig 4 below plots the analytical correction coefficient discussed in this section.

Tani *et al.* found that maximum lift decreased at moderate ground clearances but increased when the wing was brought unusually close to the ground (Tani et al.). This non-monotonic behavior shows a limitation of a correction that changes induced drag while keeping lift unchanged (Wieselsberger; Tani et al.).

Viscous CFD extended the analysis to the effects of Reynolds number, ground clearance, and angle of attack on airfoil performance. Hsiun and Chen used a steady, two-dimensional calculation with the standard k–ε turbulence model for a NACA 4412 airfoil and compared the predicted lift with measurements at Re = 3.2×10^5 and H/c = 0.4, where H is the height of the airfoil's midchord above the ground. Their viscous predictions agreed more closely with the measurements than the inviscid predictions as angle of attack increased; this comparison supports the method for the tested conditions (Hsiun and Chen).

Computational methods also made it possible to use aerodynamic predictions in shape optimization. Jameson formulated an adjoint approach in which the computational cost of each design iteration was of the same order as two flow solutions, because the adjoint equation had a complexity comparable to that of the flow equation (Jameson). This result concerns the efficiency of the optimization procedure and does not establish the accuracy of a ground-effect prediction.

Narayanan *et al.* applied artificial neural networks (ANNs) and Kolmogorov–Arnold networks (KANs) to predict optimized NACA 4412 airfoil shapes and aerodynamic coefficients in ground effect. Within the study's parameter range, the authors reported computational costs reduced by factors of approximately six for ANN and eight for KAN relative to CFD. At a validation condition of $\alpha = 3°$ and h/c = 0.2, ANN deviations from CFD were approximately 1.3% in lift and below 1.5% in drag and pitching moment; KAN deviations were below 1% in lift and pitching moment and below 4% in drag (Narayanan et al.). These values describe agreement with CFD for the reported validation condition, rather than errors relative to wind-tunnel measurements.

Taken together, these studies show an expansion from analytical drag corrections to viscous-flow prediction, shape optimization, and surrogate modeling. However, their validation cases and comparison measures differ, so the evidence examined in this review does not support a single error range between CFD predictions and wind-tunnel measurements that applies across ground-effect configurations.

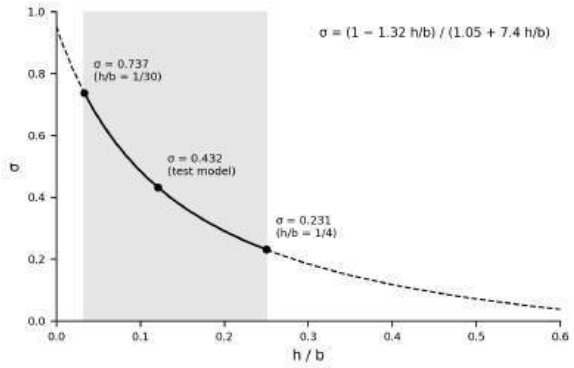

Fig 4: Wieselsberger's correction coefficient.

The influence coefficient $\sigma = (1 - 1.32\ h/b)/(1.05 + 7.4\ h/b)$, the form of Eq. (1) used in A-1, is plotted against h/b, where h is the height of the wing above the ground and b is the wingspan (Wieselsberger). The shaded band is the stated range of validity, $1/30 \leq h/b \leq 1/4$; the dashed segments extrapolate the formula outside it. Marked points: $\sigma = 0.737$ at $h/b = 1/30$; $\sigma = 0.432$ at $h/b = 0.121$, the value for Wieselsberger's test model (height 15 cm, span 124 cm); $\sigma = 0.231$ at $h/b = 1/4$. The curve was calculated from Eq. (1) using Python and Matplotlib 3.8.3 (https://matplotlib.org/); no data points are reproduced from the source figure.

Discussion: validation and limitations

Reading through R-1 to R-4, it can be seen that moving model experiments and flight testing were the most common validation methods, since these methods most accurately reflect reality. However, from the example of analytical solutions and a fixed-floor wind tunnel, it can also be assumed that if two methods omit the same aspect of a phenomenon, their results might align even though both deviate from reality.

Limits of applicability

Each method has its own limitation; that is the reason why a variety of methods remain to date.

Turner's fixed-floor test overstated lift loss by approximately 60% relative to the moving-model loss at one low-clearance condition (Turner, "Endless-Belt"). This is a configuration-specific result, not a general error for fixed-floor tunnels. LLMs can assist case preparation, but the reported failures and need for human oversight remain relevant (Pandey et al.).

Flight testing, which is a precise method that can collect data on real aircraft under real conditions, has its own limitations as well: namely, the accuracy of the measurements heavily depends on the pilot's control and weather conditions. Moreover, it is mostly impossible to obtain data on independently fixed sink rate and angle of attack.

Outlook for LLM-driven CFD

The research paper that presented OpenFOAMGPT does not include cases implemented with conditions under ground effect; moreover, the research only focused on tutorials, which are the most basic cases provided by the publisher of OpenFOAM (Pandey et al.). The author of the research also addresses that "human oversight remains crucial for ensuring accuracy and adapting to shifting contexts."

Conclusion

This review presented a chronological overview that spans the last 100 years of methodology conducted on ground effect. The chronology is divided into four historical periods, with LLM-driven CFD discussed as an emerging direction. Each method was evaluated through three criteria: cost, the applicable range, and the technology that enabled it. This review then focused on each method's transformation and validation to a specific case through different methods.

This discussion above leads to three conclusions. First, agreement between methods that share the same assumptions does not constitute verification. Second, the methods were, as shown in DGPS and moving-belt validated through multiple other methods, as a larger amount of validation makes the method more reliable, in this review. Finally, each method has its own limitations, but the diverse methods complement each other and provide a reliable outcome. However, the LLM-driven CFD is not yet validated by any other methods, leaving room for future development.

The review is selective rather than systematic and does not quantify adoption rates or provide a common cost or error scale. Its focus is wing and aircraft aerodynamics, including selected ground-effect-aircraft studies; it is not a comprehensive survey of race-car or wing-in-ground-effect vehicle research.

Acknowledgements

The author thanks Feiyuan Long, for guidance on the scope, structure, and sourcing of this review.

Author

Terumasa Iriyama is a senior at Kaisei High School, Tokyo, Japan. He conducted the literature search, the analysis, and the writing of this review, and prepared Fig 4 from the reviewed equation.

AI Acknowledgement

OpenAI Codex and Claude Code were used to assist with source checking, manuscript organization, language revision, reference formatting, and figure preparation. The author is responsible for reviewing and verifying all assisted content before submission.

Works Cited

Abbott, I. H., and A. E. von Doenhoff. Theory of Wing Sections: Including a Summary of Airfoil Data. Dover Publications, 1959. ISBN 978-0-486-60586-9.

Abu Salem, K., et al. “A Simulation Framework for Aircraft Take-Off Considering Ground Effect Aerodynamics in Conceptual Design.” Aerospace, vol. 10, no. 5, 2023, article 459. https://doi.org/10.3390/aerospace10050459.

Baker, P. A., et al. Flight Evaluation of Ground Effect on Several Low-Aspect-Ratio Airplanes. National Aeronautics and Space Administration, 1970. NASA Technical Note D-6053. https://ntrs.nasa.gov/citations/19700033422. Accessed 13 Aug. 2026.

Barber, T. J. “Aerodynamic Ground Effect: A Case Study of the Integration of CFD and Experiments.” International Journal of Vehicle Design, vol. 40, no. 4, 2006, pp. 299–316. https://doi.org/10.1504/IJVD.2006.009068.

Barber, T. J., et al. “Causes for Discrepancies in Ground Effect Analyses.” The Aeronautical Journal, vol. 106, no. 1066, 2002, pp. 653–668. https://doi.org/10.1017/S0001924000011726.

Butler, S. F. J., et al. A Moving-Belt Rig for Ground Simulation in Low-Speed Wind Tunnels. Aeronautical Research Council, 1963. Reports and Memoranda No. 3451, Council No. 25,874. Bibliographic record: Aeronautical Research Council, Reports and Memoranda No. 3550, p. 1. https://reports.aerade.cranfield.ac.uk/bitstream/handle/1826.2/2820/arc-rm-3550.pdf. Accessed 7 Sept. 2026.

Chang, R. C., and V. U. Muirhead. “Effect of Sink Rate on Ground Effect of Low-Aspect-Ratio Wings.” J. Aircr., vol. 24, no. 3, 1987, pp. 176–180. https://doi.org/10.2514/3.45413.

Chapman, D. R. “Computational Aerodynamics Development and Outlook.” AIAA J., vol. 17, no. 12, 1979, pp. 1293–1313. https://doi.org/10.2514/3.61311.

Chen, Y., et al. “MetaOpenFOAM: An LLM-Based Multi-Agent Framework for CFD.” arXiv, 2024, arXiv:2407.21320. Preprint. https://arxiv.org/abs/2407.21320.

Corda, S., et al. Dynamic Ground Effects Flight Test of an F-15 Aircraft. National Aeronautics and Space Administration, 1994. NASA Technical Memorandum 4604, H-1999. https://ntrs.nasa.gov/citations/19950005778. Accessed 13 Aug. 2026.

Cui, Erjie, and Xin Zhang. “Ground Effect Aerodynamics.” Encyclopedia of Aerospace Engineering, edited by R. Blockley and W. Shyy, Wiley, 2010. https://doi.org/10.1002/9780470686652.eae022. [要確認]

Curry, R. E. Dynamic Ground Effect for a Cranked Arrow Wing Airplane. National Aeronautics and Space Administration, 1997. NASA Technical Memorandum 4799, H-2177. https://ntrs.nasa.gov/citations/19970026087. Accessed 7 Sept. 2026.

Doig, G. “Transonic and Supersonic Ground Effect Aerodynamics.” Progress in Aerospace Sciences, vol. 69, 2014, pp. 1–28. https://doi.org/10.1016/j.paerosci.2014.02.002.

Ferguson, R. G. NASA's First A: Aeronautics from 1958 to 2008. National Aeronautics and Space Administration, 2013. NASA SP-2012-4412, ch. 6, p. 175. https://www.nasa.gov/wp-content/uploads/2015/04/NASAsFirstA-508-ebook.pdf. Accessed 13 Aug. 2026.

Foshag, William F. Literature Search and Comprehensive Bibliography of Wings in Ground Effect and Related Phenomena. David Taylor Model Basin, 1966. DTMB Report 2179, DTMB-AERO-1098, AD-633139. https://ntrs.nasa.gov/api/citations/19680023297/downloads/19680023297.pdf. [要確認]

Graham, S. L., et al., editors. Getting Up to Speed: The Future of Supercomputing. National Academies Press, 2005, pp. 35–42. https://doi.org/10.17226/11148.

Harlow, F. H., and J. E. Welch. “Numerical Calculation of Time-Dependent Viscous Incompressible Flow of Fluid with Free Surface.” Phys. Fluids, vol. 8, no. 12, 1965, pp. 2182–2189. https://doi.org/10.1063/1.1761178.

Hess, J. L., and A. M. O. Smith. “Calculation of Potential Flow about Arbitrary Bodies.” Prog. Aerosp. Sci., vol. 8, 1967, pp. 1–138. https://doi.org/10.1016/0376-0421(67)90003-6.

Hsiun, C.-M., and C.-K. Chen. “Aerodynamic Characteristics of a Two-Dimensional Airfoil with Ground Effect.” J. Aircr., vol. 33, no. 2, 1996, pp. 386–392. https://doi.org/10.2514/3.46949.

Jameson, A. “Aerodynamic Design via Control Theory.” J. Sci. Comput., vol. 3, no. 3, 1988, pp. 233–260. https://doi.org/10.1007/BF01061285.

Jesudasan, R., et al. “Investigating Planar and Nonplanar Wing Planform Optimisation for Ground Effect Aircraft.” Aerospace, vol. 10, no. 11, 2023, article 969. https://doi.org/10.3390/aerospace10110969.

Klemin, A. “A Belt Method of Representing the Ground.” J. Aeronaut. Sci., vol. 1, no. 4, 1934, pp. 198–199. https://doi.org/10.2514/8.51.

Launder, B. E., and D. B. Spalding. “The Numerical Computation of Turbulent Flows.” Comput. Methods Appl. Mech. Eng., vol. 3, no. 2, 1974, pp. 269–289. https://doi.org/10.1016/0045-7825(74)90029-2.

Lee, T., and G. Lin. “Review of Experimental Investigations of Wings in Ground Effect at Low Reynolds Numbers.” Frontiers in Aerospace Engineering, vol. 1, 2022, article 975158. https://doi.org/10.3389/fpace.2022.975158.

Martin, J. C., and W. J. Moyce. “Part IV. An Experimental Study of the Collapse of Liquid Columns on a Rigid Horizontal Plane.” Philos. Trans. R. Soc. London, Ser. A, vol. 244, no. 882, 1952, pp. 312–324. https://doi.org/10.1098/rsta.1952.0006.

Matus-Vargas, A., et al. “Ground Effect on Rotorcraft Unmanned Aerial Vehicles: A Review.” Intelligent Service Robotics, vol. 14, no. 1, 2021, pp. 99–118. https://doi.org/10.1007/s11370-020-00344-5.

McDonnell Douglas Corporation. USAF Stability and Control DATCOM. U.S. Air Force Flight Dynamics Laboratory, 1960. Revised 1976.

Menter, F. R. “Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications.” AIAA J., vol. 32, no. 8, 1994, pp. 1598–1605. https://doi.org/10.2514/3.12149.

Moore, G. E. “Cramming More Components onto Integrated Circuits.” Electronics, vol. 38, no. 8, 1965, pp. 114–117.

Narayanan, J. P., et al. “Leveraging Machine Learning for Airfoil Shape Optimization and Estimation of Aerodynamic Coefficients in Ground Effect.” Aerosp. Sci. Technol., vol. 177, part B, 2026, article 112355. https://doi.org/10.1016/j.ast.2026.112355.

NASA Glenn Research Center. “Reynolds Number.” https://www.grc.nasa.gov/WWW/k-12/airplane/reynolds.html. Accessed 7 Sept. 2026.
OpenFOAM Foundation. “OpenFOAM 1.0 Released.” https://openfoam.org/release/1-0/. Accessed 13 Aug. 2026.
Pandey, S., et al. “OpenFOAMGPT: A Retrieval-Augmented Large Language Model (LLM) Agent for OpenFOAM-Based Computational Fluid Dynamics.” Phys. Fluids, vol. 37, no. 3, 2025, article 035120. https://doi.org/10.1063/5.0257555.
Pistolesi, E. Ground Effect—Theory and Practice. National Advisory Committee for Aeronautics, 1937. NACA Technical Memorandum 828. https://ntrs.nasa.gov/citations/19930094588. [要確認]
Rozhdestvensky, K. V. “Wing-in-Ground Effect Vehicles.” Progress in Aerospace Sciences, vol. 42, no. 3, 2006, pp. 211–283. https://doi.org/10.1016/j.paerosci.2006.10.001.
Sjoberg, S. A. Flight Measurements with the Douglas D-558-II (BuAero No. 37974) Research Airplane: Static Lateral and Directional Stability Characteristics as Measured in Sideslips at Mach Numbers up to 0.87. National Advisory Committee for Aeronautics, 1950. NACA Research Memorandum L50C14. https://ntrs.nasa.gov/citations/19930086174. Accessed 13 Aug. 2026.
Slotnick, J. P., et al. CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences. National Aeronautics and Space Administration, 2014. NASA/CR-2014-218178. https://ntrs.nasa.gov/citations/20140003093. Accessed 13 Aug. 2026.
Tani, I., et al. “The Effect of Ground on the Aerodynamic Characteristics of a Monoplane Wing.” Report of the Aeronautical Research Institute, Tokyo Imperial University, vol. 13, no. 156, 1937, pp. 22–76. https://jaxa.repo.nii.ac.jp/records/35365. Accessed 13 Aug. 2026.
Turner, T. R. A Moving-Belt Ground Plane for Wind-Tunnel Ground Simulation and Results for Two Jet-Flap Configurations. National Aeronautics and Space Administration, 1967. NASA Technical Note D-4228.
Turner, T. R. “Endless-Belt Technique for Ground Simulation.” Conference on V/STOL and STOL Aircraft, National Aeronautics and Space Administration, 1966, pp. 435–446. NASA SP-116. https://ntrs.nasa.gov/citations/19660015317. Accessed 13 Aug. 2026.
Weller, H. G., et al. “A Tensorial Approach to Computational Continuum Mechanics Using Object-Oriented Techniques.” Comput. Phys., vol. 12, no. 6, 1998, pp. 620–631. https://doi.org/10.1063/1.168744.
Widnall, S. E., and T. M. Barrows. “An Analytic Solution for Two- and Three-Dimensional Wings in Ground Effect.” J. Fluid Mech., vol. 41, no. 4, 1970, pp. 769–792. https://doi.org/10.1017/S0022112070000915.
Wieselsberger, C. Wing Resistance Near the Ground. National Advisory Committee for Aeronautics, 1922. NACA Technical Memorandum 77. https://digital.library.unt.edu/ark:/67531/metadc54179/. Accessed 13 Aug. 2026.
Zhang, X., et al. “Ground Effect Aerodynamics of Race Cars.” Applied Mechanics Reviews, vol. 59, no. 1, 2006, pp. 33–49. https://doi.org/10.1115/1.2110263.

A Directional-Statistics Study of the Normalized Zero Spacings of the Riemann Zeta Function By Carlo Marville Harefa

Abstract

The consecutive spacings between non-trivial zeros of the Riemann zeta function are examined through the lens of directional statistics. Each spacing is unfolded according to the local density of zeros, mapped onto the circle using a fixed cutoff of $C = 4$, and the resulting angles are fitted with a von Mises distribution. For the zeta zeros, the estimated concentration parameter κ declines steadily with sample size, from 3.780 at $N = 10^2$ to 3.093 at $N = 10^6$. Across 100 independent samples from the Gaussian Unitary Ensemble, the median κ instead remains close to 2.91. The mean direction is approximately 1.55 for both the zeta zeros and the GUE controls. A sensitivity analysis over 21 cutoff values identifies those choices of C that avoid angular wraparound. These results warrant careful interpretation. Since the zeros are deterministic, and since the mapping imposes a circular structure on bounded linear spacings, the reported quantities characterize the behavior of the analysis pipeline rather than the zeros themselves. They offer no evidence that the zeros are random, nor that circular concentration is an intrinsic property of the zeros.

Keywords Mathematics; Number Theory; Riemann zeta function; zero spacings; directional statistics.

Introduction

Over 150 years ago, German mathematician Bernhard Riemann proposed the famous Riemann Hypothesis (RH) in his landmark paper, "Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse" (On the number of primes less than a given quantity) (Riemann). In short, the RH states that all "non-trivial zeros" of the Riemann zeta function lie on a single vertical line in the complex plane. Yet, despite extensive efforts by mathematicians over the past century, the conjecture has remained unresolved. Although it may appear to be an abstract problem, mathematicians have long admired its apparent elegance, particularly for its unobvious connection to number theory.

In 1972, number theorist Hugh Montgomery, who was studying the gaps between zeta zeros, and physicist Freeman Dyson, who was studying energy level spacings in quantum systems, compared their work at a tea party at Princeton—and found the statistics were identical. Later in 1987, A. M. Odlyzko, in his paper "On the Distribution of Spacings Between Zeros of the Zeta Function," (Odlyzko) studied the gaps between the zeros of the Riemann zeta function. In his observation, Odlyzko found that the distribution of these gaps agrees well with predictions in quantum physics, specifically the Montgomery pair correlation conjecture and the behavior of energy-level spacings in quantum systems, known as the Gaussian Unitary Ensemble (GUE) (Odlyzko). Even so, Odlyzko himself admitted that several features of the data on zeros are not predicted by the GUE theory, suggesting that this relation is still not fully understood.

Later studies have investigated how the spacings of zeta zeros at finite height diverge from the predictions of random matrix theory, as well as finite-size corrections to GUE spacing statistics (Bogomolny et al.; Forrester and Mays). This work shows that comparisons based on finite samples can diverge noticeably from asymptotic results.

In an attempt to understand the zeros of the Riemann zeta function, most studies treat these gaps or spacings as linear data, as in the previously discussed study by Odlyzko and another by Rudnick and Sarnak, which treated the zeros as points on a number line (Odlyzko; Rudnick and Sarnak). This paper, on the other hand, takes a different approach by treating the gaps as circular data instead, using directional statistics. Specifically, the von Mises distribution will be applied to the spacings of zeros, and the results will be reported. It is important to note that this paper is not an attempt at resolving the Riemann Hypothesis. Instead, using directional analysis in lieu of the traditional linear approach, it will offer an alternative perspective on the distribution of zeros that may reveal patterns or structures not immediately apparent.

Background Information

The Riemann Zeta Function and Its Zeros

The Riemann zeta function, denoted as $\zeta(s)$, is a complex-valued function initially defined as the infinite sum

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s},$$

known as a Dirichlet series, which converges for $R(s) > 1$ (Karatsuba and Voronin). Through a process called analytical continuation, mathematicians were able to extend the definition of this function so that $\zeta(s)$ is defined for nearly all complex numbers s. However, there is another definition for $\zeta(s)$ as a product instead of a sum, that is, the Euler product. For $s = \sigma + it, \sigma > 1$, the following identity holds:

$$\zeta(s) = \prod_{p} \frac{1}{1-p^{-s}},$$

where the product is taken over all prime numbers p (Karatsuba and Voronin). Because $\zeta(s)$ can be purely written in terms of primes, then the behavior of $\zeta(s)$—its zeros, how fast it grows, etc.—must somehow encode information about how primes are distributed.

Riemann himself showed that the non-trivial zeros of $\zeta(s)$ actually control the error term in the prime counting function, particularly how far the actual count of primes below some number x deviates from the smooth approximation given by the prime number theorem ($Li(x)$, the logarithmic integral) (Karatsuba and Voronin). In simpler terms, if RH is true, the error term stays as small as possible, meaning primes are distributed as evenly as the theory allows; if false, primes could clump or scatter more irregularly than expected. This is why mathematicians are eager to prove the RH—that every non-trivial zero has real part $\frac{1}{2}$—even going to the extent of putting a monetary prize offered for its resolution; because of the aforementioned link between

the zeros and prime-counting error terms, progress on RH would greatly improve our understanding of the distribution of primes.

Zero Spacings and Unfolding

Beyond their role in counting primes, the non-trivial zeros have also been studied through their spacing, that is, the gaps between consecutive zeros along the critical line. The Riemann-von Mangoldt formula for the number of non-trivial zeros up to height T has leading terms

$$N(T) \approx \frac{T}{2\pi} log\left(\frac{T}{2\pi}\right) - \frac{T}{2\pi},$$

where log is the natural logarithm (Mangoldt). Differentiating the leading term gives the local density of zeros,

$$\rho(T) = \frac{1}{2\pi} log\left(\frac{T}{2\pi}\right),$$

and the corresponding average spacing is approximately

$$\frac{1}{\rho(T)} = \frac{2\pi}{log\left(\frac{T}{2\pi}\right)}.$$

As T increases, this density grows and the average spacing shrinks. This motivates unfolding: each raw spacing is rescaled by the local density at its height,

$$unfolded\ spacing = (raw\ spacing) \cdot \rho(T) = (raw\ spacing) \cdot \frac{log\left(\frac{T}{2\pi}\right)}{2\pi}.$$

The unfolded spacings are nonnegative real numbers with local mean approximately 1, allowing comparison of a gap near the tenth zero with one near the hundred-thousandth zero.

Circular Statistics and the von Mises Distribution

As previously mentioned, these spacings are often treated as linear data. However, instead of doing exactly that, this paper maps the spacings onto a circle, converting each into an angle, and turns them into directional/circular data. Unlike linear data, circular data cannot be simply averaged using ordinary arithmetic methods. For example, arithmetically averaging the two angles 1° and 359° would yield 180°. However, on a circle, their true midpoint is 0° because they lie beside each other near the $0^{\circ}/360^{\circ}$ boundary. This occurs because linear statistics measure distance along a straight number line, while points on a circle "wrap around," making the two ends of that line essentially neighbors.

Because of this, modeling circular data requires tools different than those used for linear data, namely the von Mises distribution. The von Mises distribution is the circular analog of the normal distribution on a line, that is, a probability density function for circular data. The von Mises density function is given by

$$f(\theta|\mu, \kappa) = \frac{e^{\kappa cos(\theta - \mu)}}{2\pi I_0(\kappa)},$$

for $\theta \in [0, 2\pi)$, $\mu \in [0, 2\pi)$, $\kappa > 0$, and where

$$I_0(z) = \frac{1}{2\pi}\int_0^{2\pi} e^{z\cos\theta}\, d\theta, \qquad z \in C,$$

the modified Bessel function of the first kind of order zero (Mardia and Jupp). The two parameters of the von Mises distribution, μ and κ, represent the mean direction and the concentration of the distribution, respectively. To visualize, for small values of κ, the distribution is very close to uniform, and $\kappa = 0$ corresponds exactly to a uniform distribution over the interval; on the other hand, for large values of κ, the distribution would peak—sharper as κ increases—at the direction μ (Mardia and Jupp). This behavior is shown in Figure 1.

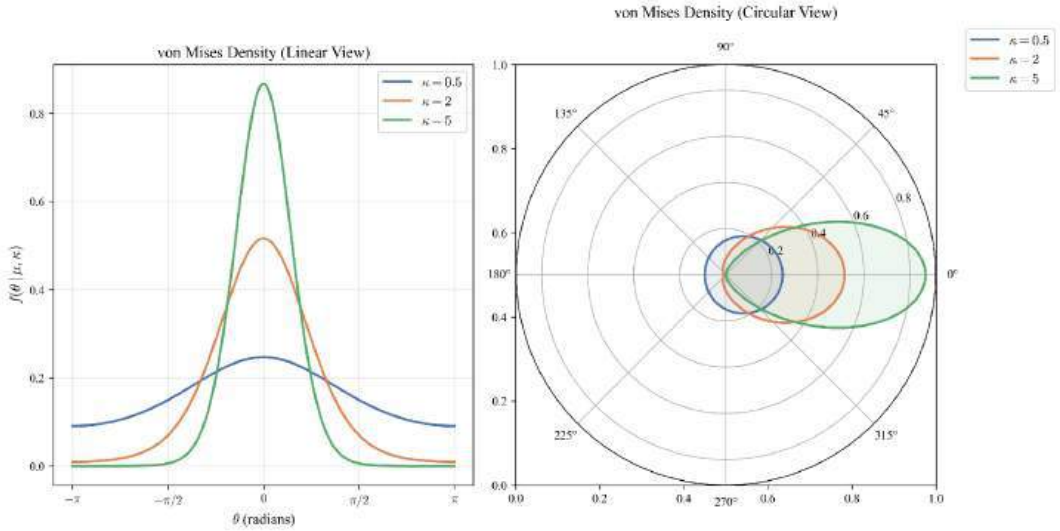

Fig 1: The von Mises density for $\kappa = 0.5, 2, 5$, shown linearly (left) and on the circle (right).

Data and Methodology

Data Source

To perform statistical analysis on the zeros, one must have a hold of the zeros themselves. The zeros in this paper are sourced from Odlyzko's high-precision tables of non-trivial zeta zero heights, hosted via the LMFDB database (The LMFDB Collaboration). However, the web interface of LMFDB has a retrieval limit of (exactly) the first 100000 zeros. As such, to retrieve samples beyond this limit, the raw binary .dat files are directly downloaded from the public directory and parsed. The .dat files store the zero heights in a compressed binary format, which are then converted into a flat array and saved as a .npy file for compatibility with Python.

The data used in this paper are the heights of the first N non-trivial zeros of $\zeta(s)$,

$$\rho_n = \frac{1}{2} + i\gamma_n,$$

where γ_n denotes the height of the n-th zero. The values of N considered are

$$N \in \{10^2, 10^3, 10^4, 10^5, 10^6\}.$$

The motivation for testing multiple values of N is to observe whether the estimated concentration κ changes systematically as sample size increases.

For each sample, the first N zero heights produce $N - 1$ consecutive spacings.

Note that these are nested subsets rather than independent samples—the first N zeros for smaller N are contained within the datasets for larger N. As such, the observed trend in κ reflects the effect of extending the same sequence of zeros, not independent draws from the zero distribution.

Unfolding

For each zero, the spacings are computed as the difference between two consecutive zeros,

$$\Delta_n = \gamma_n - \gamma_{n-1}.$$

As previously mentioned, each raw spacing is rescaled using the unfolding formula, evaluated at the midpoint height between consecutive zeros. Thus, the unfolded spacings are defined as

$$X_n = \Delta_n \cdot \frac{log\left(\frac{(\gamma_n+\gamma_{n-1})/2}{2\pi}\right)}{2\pi} = \Delta_n \cdot \frac{log\left(\frac{\gamma_n+\gamma_{n-1}}{4\pi}\right)}{2\pi}.$$

Mapping to Circular Data

Unfolded spacings are not naturally circular. As such, to perform circular statistics on the unfolded spacings, they must be mapped to angles $\theta \in [0, 2\pi)$. One approach is to use min-max normalization, which rescales each unfolded spacing X_n so that all values fall within a specified range—in the case of circular data, $[0, 2\pi)$ (Zheng and Casari). For this interval, each X_n is transformed according to

$$\theta_n^{mm} = \frac{X_n - X_{min}}{X_{max} - X_{min}} \cdot 2\pi,$$

where X_{min} and X_{max} denote the minimum and maximum unfolded spacings in the sample, respectively, and θ_n^{mm} is the resulting min-max angle. The smallest spacing is thus mapped to 0, the largest to 2π, and all remaining spacings are scaled proportionally within the interval.

One issue with this approach, however, is that the positions 0 and 2π always correspond to the smallest and largest values in that particular sample. This means the 'ruler' used to place unfolded spacing X_n on the circle changes from one sample to another; an unfolded spacing X_n of, say, 1.5 might land at a different angle in a sample of 1,000 zeros than in a sample of 1,000,000 zeros, simply because the largest spacing observed differs between samples.

To avoid this, instead of letting each sample of spacings define its own minimum and maximum, a fixed reference value C can be chosen ahead of time, independent of any particular sample. Every unfolded spacing X_n is then divided by this same C and scaled onto the circle. As

long as C is chosen to comfortably exceed the largest unfolded spacing expected to occur, a spacing of 1.5 will always land at the same angle, regardless of which sample it came from.

$$\theta_n = \frac{X_n}{C} \cdot 2\pi.$$

To avoid any wraparound in the circle, it must remain below 2π, which is equivalent to requiring $X_n < C$ for all spacings. In this paper, $C = 4$ is used.

The analysis verifies that no spacing reaches C, so no angle wraps around.

The earlier, sample-dependent min-max mapping was ultimately discarded because it induces a spurious sample-size trend in κ; this is documented in the Appendix.

Remark 1. *Unlike genuinely periodic circular data (e.g., wind direction), the mapped spacings here are a bounded linear quantity artificially embedded on a circle. The von Mises distribution is used as a convenient unimodal density for bounded data, not because the underlying phenomenon is directional in any physical sense.*

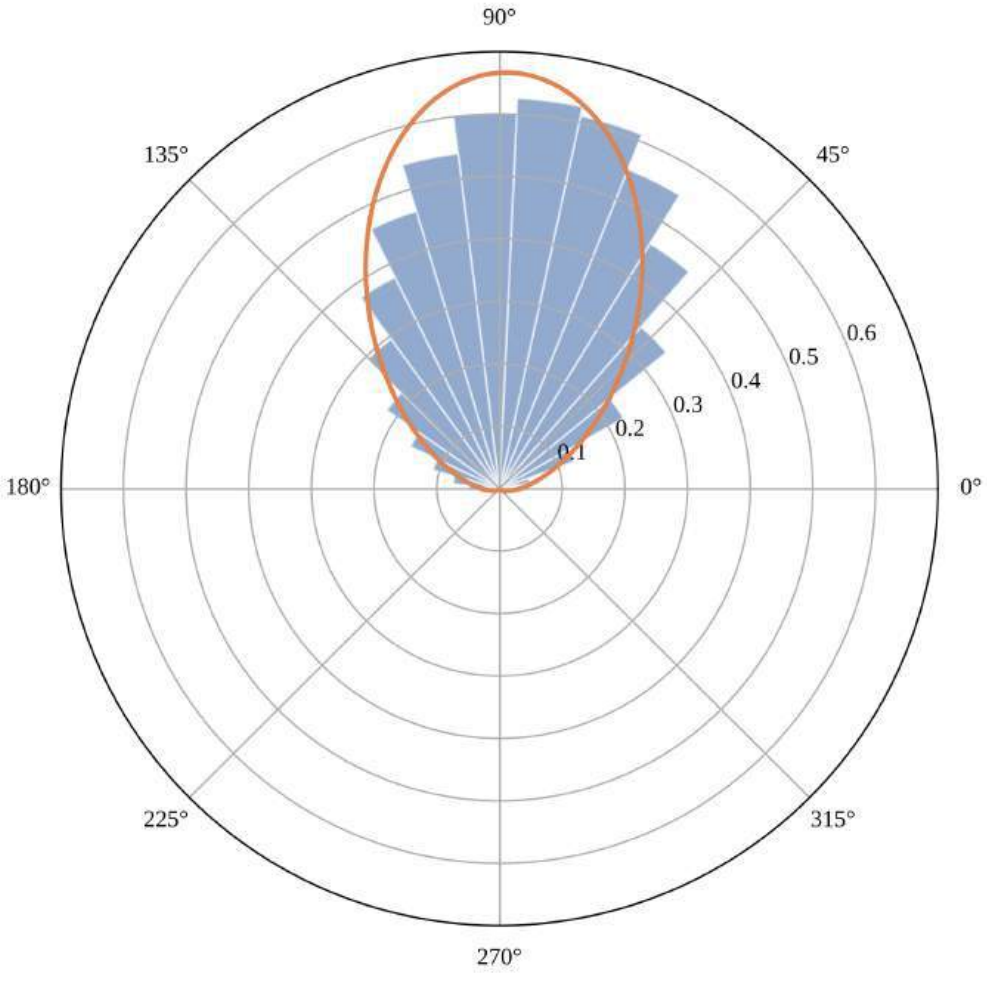

Fig 2: Unfolded zeta zero spacings for $N = 10^6$ mapped to angles with $C = 4$, with the fitted von Mises density.

von Mises Fitting

After obtaining the circular data θ_n, the von Mises distribution is fitted to the angles. The estimated parameters are the mean direction μ and concentration parameter κ. For each sample size N, the same fitting procedure is applied, producing one pair (μ, κ) for the zeta zero spacings.

For angles θ_n, the estimated mean direction is the circular mean, $\hat{\mu} = atan2(s^-, c^-)$ reduced to $[0, 2\pi)$, where $c^- = n^{-1}\sum_n cos\theta_n$ and $s^- = n^{-1}\sum_n sin\theta_n$. Let $R^- = \sqrt{c^{-2} + s^{-2}}$. The concentration estimate $\hat{\kappa}$ is obtained by solving $I_1(\hat{\kappa})/I_0(\hat{\kappa}) = R^-$, where I_1 is the modified Bessel function of the first kind of order one (Mardia and Jupp).

GUE Control

To provide a comparison baseline, the study generates 100 independent control samples for each sample size N, each containing $N - 1$ spacings. A standard chi variate with three degrees of freedom, $Y \sim \chi_3$, is scaled as $s = Y\sqrt{\pi/8}$. This sampling procedure has the Wigner-surmise density for the Gaussian Unitary Ensemble (GUE),

$$P(s) = \frac{32}{\pi^2}s^2 e^{-4s^2/\pi},$$

with mean spacing 1 (Mehta). The Wigner surmise is a close approximation to, but not exactly, the GUE nearest-neighbor spacing distribution. For every control sample, the spacings are mapped using the same fixed cutoff and fitted using the same procedure as the zeta data. For each N, the control summary reports the median fitted parameters and the 95% range of fitted κ across the 100 independent samples; this between-sample spread is the control uncertainty.

The analysis code is written in Python. The principal scripts are `main_results.py` (Table 1 and Figures 2, 3, and 7) and `analyze_unfolded_spacings.py` (Figure 5).

Cutoff Sensitivity

The analysis was repeated for 21 cutoffs from $C = 2.0$ to 6.0 in steps of 0.2. Cutoffs at or below the largest spacing in either the zeta data or the control were excluded; at $N = 10^6$ this leaves $C = 4.0$ to 6.0, while smaller N allow lower cutoffs. The largest spacing is 3.21 for the zeta zeros and 3.91 across all control samples.

Results

This section presents the estimated von Mises parameters (μ, κ) obtained by applying the methodology described above to the first N non-trivial zeta zeros, for $N = 10^2, ..., 10^6$, alongside the corresponding results from the GUE control.

Table 1: Zeta von Mises estimates with bootstrap standard errors and GUE-control medians and 95% ranges across 100 samples, for $C = 4$.

N	Zeta $\hat{\mu}$ (SE)	Zeta $\hat{\kappa}$ (SE)	Control $\hat{\mu}$ (median)	Control $\hat{\kappa}$ (median) [95% range]
10^2	1.56683 (0.05864)	3.78018 (0.40259)	1.55799	2.98962 [2.51136, 3.66716]
10^3	1.55346 (0.01897)	3.41510 (0.12200)	1.54955	2.91419 [2.74707, 3.13939]
10^4	1.55106 (0.00591)	3.25777 (0.03473)	1.54706	2.91357 [2.85708, 2.97740]
10^5	1.55054 (0.00202)	3.15707 (0.01104)	1.54650	2.91433 [2.89185, 2.93438]
10^6	1.55019 (0.00065)	3.09305 (0.00355)	1.54695	2.91504 [2.91047, 2.92072]

Standard errors for the zeta estimates were estimated via bootstrap resampling ($B = 1,000$ resamples). At each iteration, the angle data θ_n were resampled with replacement and the von Mises parameters were re-estimated; the standard deviation of the resulting estimates is reported as the standard error. Since the zeta zeros are deterministic, these standard errors describe the resampling variability of the estimates. For the control, uncertainty is represented by the spread across the 100 independent samples.

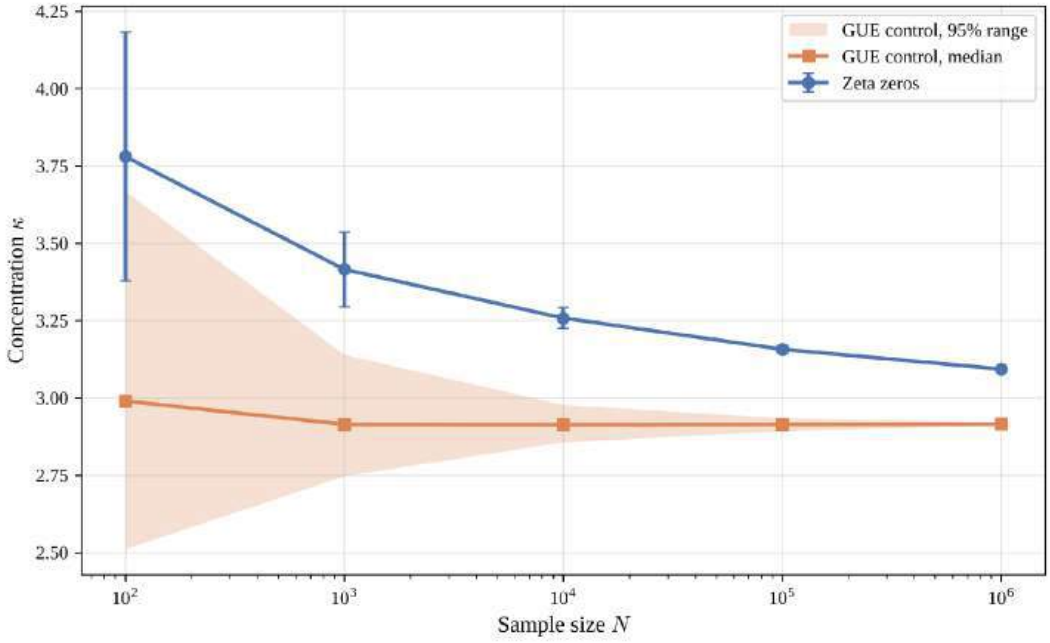

Fig 3: Estimated κ versus sample size N for the zeta zero spacings (error bars: bootstrap standard errors) and the GUE control (median and 95% range across 100 samples), with $C = 4$.

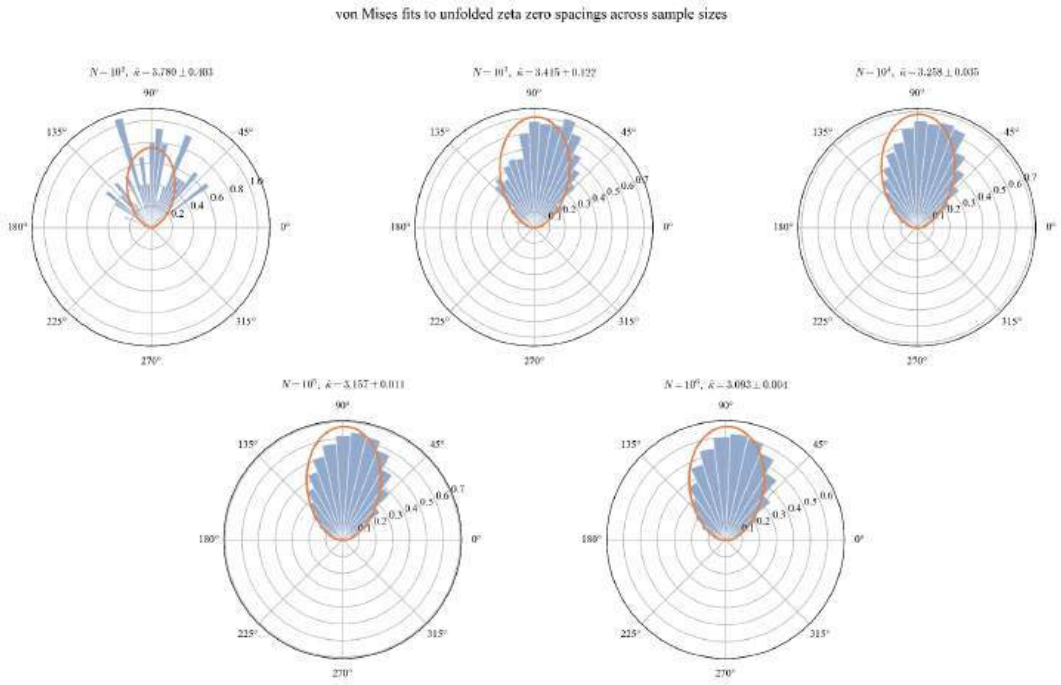

Fig 4: von Mises fits to the unfolded zeta zero spacings for $N = 10^2$ to 10^6 ($C = 4$); each panel shows $\hat{\kappa} \pm$ bootstrap standard error.

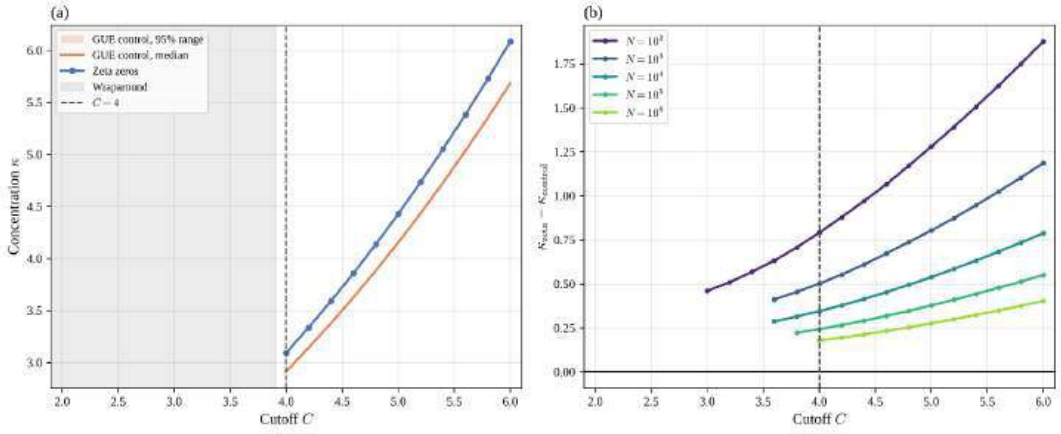

Fig 5: Sensitivity to the cutoff C. (a) κ versus C at $N = 10^6$ for the zeta zeros and the GUE control (median and 95% range across 100 samples); the shaded region marks cutoffs at or below the largest spacing (3.91), which are excluded because the angles would wrap around. (b) Difference between the zeta κ and the control median for each N.

Discussion

As seen in Figure 3 and Table 1, for $C = 4$, the zeta estimate of κ decreases from 3.780 at $N = 10^2$ to 3.093 at $N = 10^6$, while the median GUE-control estimate stays near 2.91 (Table 1, Figure 3). The gap between the two narrows as N grows but stays positive at every sample

size tested. Both datasets give a mean direction of about 1.55. This is to be expected, as unfolding sets the mean spacing to 1, and with $C = 4$ a spacing of 1 maps to an angle of $\pi/2 \approx 1.571$. The mean direction therefore mainly reflects the normalization and says little about how the two datasets differ. Figure 5 shows that the absolute values of κ change with the cutoff, so the more informative quantity is the zeta-control difference evaluated across the valid cutoffs.

It is important to recall that the obtained results are a reflection of how the zeros behave under one specific unfolding and circular mapping, and should not be taken as a fundamental property of the zeta zeros themselves. The difference in concentration between the zeta zeros and the GUE samples is also consistent with known effects: earlier studies have shown that at finite heights, zeta-zero spacings slightly deviate from random-matrix predictions, but these deviations diminish gradually as the height increases (Bogomolny et al.; Forrester and Mays).

The methodology itself has several limitations. Firstly, the value of κ depends on the choie of the cutoff C, so it describes the spacings only as they appear after the mapping, not as an absolute property of the original data. The sensitivity analysis covers only the tested range of cutoffs. Secondly, the control samples come from the Wigner surmise, which, albeit close, is not an exact approximation of the true GUE spacing distribution. Part of the observed difference may therefore come from the control model. Finally, the zeta samples are nested, with each larger sample containing the smaller ones, implying that the estimates at different values of N are not independent of one another. Despite these limitations, circular statistics may still serve as a useful exploratory tool for describing spacing data.

These findings have no bearing on the truth of the Riemann Hypothesis. They do point to two areas for further work: principled choices of transformation, and the role of finite-sample effects in statistical studies of zeta-zero spacings.

Conclusion

This study unfolded consecutive spacings of the Riemann zeta zeros, mapped them to angles using fixed cutoffs, and modeled the resulting angles with the von Mises distribution. At $C = 4$, the zeta estimate of κ fell from 3.780 at $N = 10^2$ to 3.093 at $N = 10^6$. Across 100 GUE-control samples, the median κ stayed close to 2.91. Both datasets had a mean direction of about 1.55, consistent with the unit mean spacing imposed by unfolding. The difference in concentration narrowed as N increased, and the cutoff-sensitivity analysis showed that the comparison is meaningful only for cutoffs large enough to prevent wraparound. These results characterize this exploratory circular-statistics pipeline and do not bear on the truth of the Riemann Hypothesis. Future work could compare alternative mappings and test whether finite-N GUE models, rather than the Wigner surmise, change the results.

AI Acknowledgement

During the preparation of this work, the author used GitHub Copilot and ChatGPT in order to troubleshoot computational implementations and assist with document formatting (i.e.,

equation rendering, reference formatting, and document conversion), respectively. After using these tool, the author reviewed and edited the content as needed and took full responsibility for the content of the publication.

Acknowledgements

The author gratefully acknowledges Professor George Aivaliotis for his mentorship and guidance throughout this research project. The author also thanks Indigo Research for providing the research mentorship program through which this work was conducted. Finally, the author is grateful to the LMFDB Collaboration and Andrew M. Odlyzko for making high-precision Riemann zeta zero datasets publicly available, which made this study possible.

Appendix

The Discarded Min-Max Mapping and Its Sample-Size Artifact

Before opting for the fixed-cutoff C method, the min-max normalization was used to map zero spacings to angles. This appendix documents the earlier pipeline—the approach—supporting the claim in section 3.3 that the latter becomes the dominant artifact in the resulting concentration parameter κ.

The earlier pipeline

The discarded pipeline differed from the final method in two respects. First, it operated on the *raw* spacings $\Delta_n = \gamma_n - \gamma_{n-1}$ standardized to zero mean and unit variance,

$$Z_n = \frac{\Delta_n - \bar{\Delta}}{\sigma_\Delta},$$

rather than on the Riemann-von Mangoldt unfolded spacings X_n. Second, it mapped these standardized values to angles using the sample-dependent min-max transform, with X_n replaced by Z_n:

$$\theta_n^{mm} = \frac{Z_n - Z_{min}}{Z_{max} - Z_{min}} \cdot 2\pi.$$

The von Mises fit was then performed on θ_n^{mm} using the fitting procedure described above.

Results under the discarded pipeline

Table 2 reports the fitted parameters under this earlier pipeline, for the same sample sizes N and the same bootstrap procedure ($B = 1,000$ resamples) used in the main text.

Table 2: Estimated von Mises parameters (μ, κ) for the discarded min-max pipeline, with bootstrap standard errors based on 1,000 resamples.

N	$\hat{\mu}$	$SE(\hat{\mu})$	$\hat{\kappa}$	$SE(\hat{\kappa})$
10^2	1.36771	0.09598	1.74343	0.22378
10^3	1.11839	0.01706	3.73669	0.16634
10^4	0.85861	0.00370	7.17072	0.12611
10^5	0.66776	0.00089	12.34367	0.06746
10^6	0.54142	0.00023	19.24018	0.03170

Why the trend is an artifact

Under the discarded pipeline, κ grows without bound as N increases, that is, from 1.74 at $N = 10^2$ to 19.24 at $N = 10^6$, an order of magnitude, which is a direct consequence of combining z-scoring with a sample-dependent min-max map.

The mechanism is as follows. The single largest raw spacing in the entire dataset is the very first gap, $\gamma_2 - \gamma_1 \approx 6.887$, which is fixed and does not change with N. As N grows, however, the local density of zeros increases and the standard deviation σ_Δ of the raw spacings shrinks. Consequently, the standardized value Z_n of that fixed largest gap grows steadily, where its z-score rises from ≈ 4.4 at $N = 10^2$ to ≈ 24.8 at $N = 10^6$. Because min-max stretches the angular ruler to span from the minimum to this single, ever-more-extreme maximum, the bulk of the spacings is compressed into an increasingly narrow arc. The von Mises fit interprets this artificial compression as rising concentration, as reflected by the inflation of κ as N increases. Figure 6 illustrates the latter for $N = 10^5$, with the mapped angles confined to a narrow arc and the von Mises fit sharply concentrated around it.

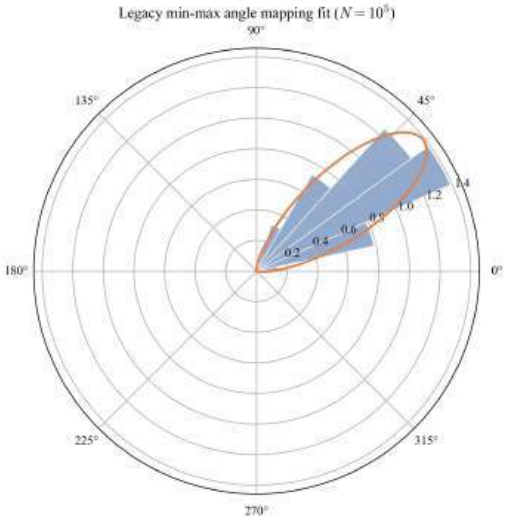

Fig 6: Legacy min-max angle mapping and von Mises fit for $N = 10^5$.

Two properties of the final method remove this artifact. Unfolding rescales each spacing by the local density, so the unfolded spacings have mean 1 and a stable range at every height rather than a shrinking variance. The fixed cutoff C then maps spacings onto the circle using a constant 'ruler', as previously mentioned, so a spacing of a given size always lands at the same angle regardless of N. Under this method, κ declines from 3.78 to 3.09 across the same range of N (Table 1), while remaining above the GUE-control median.

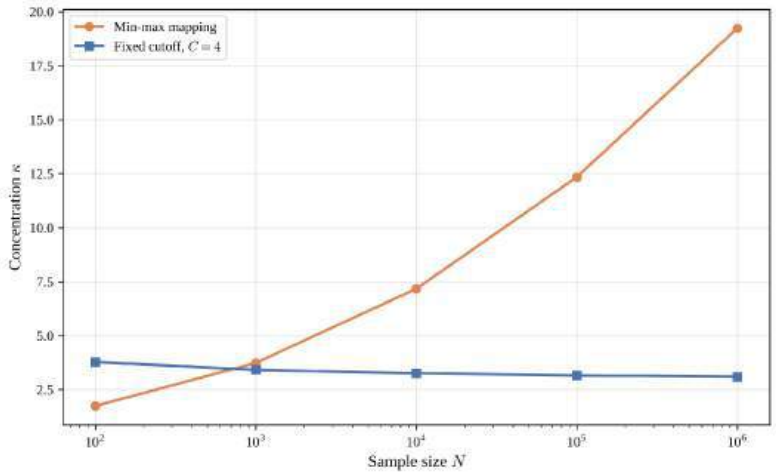

Fig 7: Estimated κ versus sample size N under the discarded min-max mapping of z-scored raw spacings and the final fixed-cutoff mapping of unfolded spacings ($C = 4$).

Works Cited

Bogomolny, Eugene, et al. “On the Spacing Distribution of the Riemann Zeros: Corrections to the Asymptotic Result.” Journal of Physics A: Mathematical and General, vol. 39, no. 34, 2006, pp. 10743–10754. https://doi.org/10.1088/0305-4470/39/34/010.

Forrester, Peter J., and Anthony Mays. “Finite size corrections and Odlyzko’s data set for the Riemann zeros.” Proceedings of the Royal Society A, vol. 471, no. 2182, 2015, article 20150436. https://doi.org/10.1098/rspa.2015.0436.

Karatsuba, Anatoly A., and S. M. Voronin. The Riemann Zeta-Function. Translated by N. Koblitz, De Gruyter, 1992. https://doi.org/10.1515/9783110886146; https://www.degruyterbrill.com/document/doi/10.1515/9783110886146/html. Accessed 26 July 2026.

The LMFDB Collaboration. “Zeros of the Riemann Zeta Function.” The L-Functions and Modular Forms Database, 2026, https://www.lmfdb.org/zeros/zeta/. Accessed 26 July 2026.

Mangoldt, Hans von. “Zu Riemanns Abhandlung „Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse“.” Journal für die reine und angewandte Mathematik, vol. 114, 1895, pp. 255–305. https://doi.org/10.1515/crll.1895.114.255; https://www.degruyter.com/document/doi/10.1515/crll.1895.114.255/html. Accessed 26 July 2026.

Mardia, Kanti V., and Peter E. Jupp. Directional Statistics. Wiley, 1999. https://doi.org/10.1002/9780470316979; https://onlinelibrary.wiley.com/doi/book/10.1002/9780470316979. Accessed 26 July 2026.

Mehta, Madan Lal. Random Matrices. 3rd ed., Elsevier/Academic Press, 2004.

Odlyzko, A. M. “On the Distribution of Spacings between Zeros of the Zeta Function.” Mathematics of Computation, vol. 48, no. 177, 1987, pp. 273–308. https://doi.org/10.1090/S0025-5718-1987-0866115-0; https://www.ams.org/mcom/1987-48-177/S0025-5718-1987-0866115-0/. Accessed 26 July 2026.

Riemann, Bernhard. “On the Number of Prime Numbers Less than a Given Quantity.” Translated by David R. Wilkins, Monatsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin, 1859, pp. 671–680. https://www.claymath.org/wp-content/uploads/2023/04/Wilkins-translation.pdf. Accessed 26 July 2026.

Rudnick, Zeév, and Peter Sarnak. “Zeros of Principal L-Functions and Random Matrix Theory.” Duke Mathematical Journal, vol. 81, no. 2, 1996, pp. 269–322. https://doi.org/10.1215/S0012-7094-96-08115-6; https://projecteuclid.org/journals/duke-mathematical-journal/volume-81/issue-2/Zeros-of-principal-L-functions-and-random-matrix-theory/10.1215/S0012-7094-96-08115-6.full. Accessed 26 July 2026.

Zheng, Alice, and Amanda Casari. Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists. O’Reilly, 2018.

Engineering Challenges Faced in the Implementation of Atomic Clocks for Autonomous Deep Space Navigation

By Nomeda Matonyte[1] *

Abstract

One of the most exciting applications of atomic clocks is deployment onboard spacecraft to enable a more responsive one-way tracking method instead of two-way tracking method. However, one-way tracking requires an onboard clock stable enough to remain within the timing error budget, a demand that current onboard ultra stable oscillators cannot achieve long-term but atomic clocks can. Atomic clocks can reach sufficient stability, requiring complex systems that create engineering challenges, such as maintaining long-term stability; operability at low size, weight and power; survivability and transition to autonomous navigation. These challenges are identified through research including simulations, and measured by a NASA Technology Demonstration Mission. The study discusses engineering challenges facing the use of onboard atomic clocks in deep space navigation, as well as optical clocks, a frontier technology and a potential alternative to standard microwave atomic clocks in deep space navigation systems. A conclusion is reached that none of the space-qualified atomic clocks are qualified for all deep space missions, since all three criteria are never simultaneously satisfied: long-term stability, low size, weight and power, flight survivability.

Keywords

Engineering, aerospace engineering, deep space navigation, atomic clocks, one-way radiometric tracking

Introduction

As deep space probes complete missions further from Earth, such as NASA's New Horizons Pluto flyby, signal transmission time to Deep Space Network (DSN) ground stations grows, requiring autonomous navigation.[2] However, development of autonomous space navigation systems can aid the issue by changing the tracking method from two-way to one-way radiometric tracking.

Two-way radiometric tracking (Figure 1a) is the standard method for deep space navigation; it measures the time taken for a radio (hence radiometric) signal to travel from a DSN ground station on Earth to a spacecraft and back and the change in signal frequency (Doppler effect). The time taken to go from ground to the spacecraft is multiplied by the speed of electromagnetic waves to determine the distance, while the shift in frequency gives the spacecraft's velocity along the line of sight.[3] Two way radiometric tracking allows the calculations to be carried out using two timestamps from the same atomic clock in a DSN ground station, eliminating timing error.

However, two-way tracking requires the radio signal to travel to the spacecraft and back, resulting in twice the one-way light time and producing a delay before the position is calculated.

1. St. Clare's, Oxford, 139 Banbury Road, Oxford, OX2 7AL, United Kingdom
nomeda.matonyte@gmail.com

During transmission time back to Earth, the spacecraft moves further along its trajectory, making its exact real-time position difficult to calculate. In time-critical maneuvers, such as a change in trajectory, a position calculation delay can not be tolerated. The calculations and data processing time in DSN ground stations further increase delay. Autonomous navigation using one-way tracking enables responsive navigation, decreasing delay to only that of onboard position calculations.[4] Additionally, one-way tracking achieves autonomy by removing the need to stay in continuous contact with Earth, as it can perform position calculations locally using a stable atomic clock.[5]

One-way (Figure 1b) radiometric tracking employs a single radio signal transmitted from Earth to spacecraft to measure light time, the time taken for radio waves to travel between Earth and spacecraft, which is used to determine the position of a spacecraft. One-way tracking's working principle is as follows. The spacecraft forms the range observable by referencing the received signal against its own onboard atomic frequency standard (AFS), an atom oscillator oscillating between two energy levels, rather than the atomic clock in a DSN ground station.[6] High long-term stability is essential, to maintain a similar time as the atomic clocks in DSN ground stations. Unlike two-way radiometric tracking, which cancels out timing error, the timing error on the AFS is included in light time, directly translating to a position error.[1]

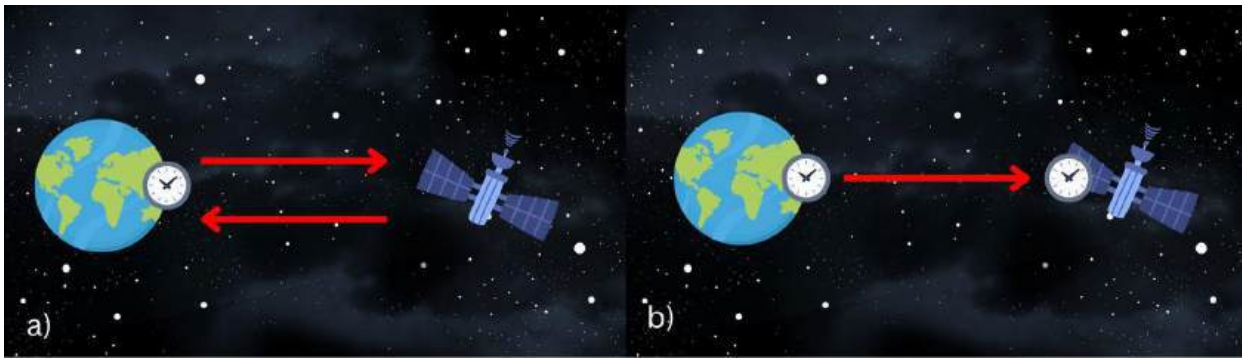

Fig 1. Tracking systems in deep space navigation a) Two-way radiometric tracking b) One-way radiometric tracking

This review paper reviews engineering challenges that must be resolved to use atomic clocks for one-way radiometric tracking in autonomous deep space navigation. In the first section, space-qualified frequency standards and optical clocks in development for potential deployment to deep space are overviewed. The main engineering challenges discussed in the following sections include atomic clock holdover stability; reducing size, weight and power (SWaP); deep space survivability; and onboard autonomy. The penultimate section analyses optical clock development and obstacles to deployability and the final section concludes the study.

Background on atomic clocks

This section establishes AFS taxonomy, including the rubidium atomic frequency standard (RAFS), the passive hydrogen maser (PHM), the deep space atomic clock (DSAC), the chip-scale atomic clock (CSAC), and optical clocks as well as how AFS stability is measured.

The working principle of an atomic clock is as follows: an atom oscillation between two fixed energy levels corresponds to a specific frequency, making it incredibly stable. In contrast, a mechanical oscillator's, for instance a quartz oscillator's, stability depends on aging physical dimensions of a manufactured device.[7] This article discusses passive AFS as their architecture is more applicable in autonomous deep space navigation due to lower power consumption than active AFS. Passive AFS require a driving device - an ultra stable oscillator, usually quartz — for atom oscillation. The USO probes the atom continuously and corrects the crystal's frequency to resonate with the frequency of the atom.[8] Atomic clock categories that are relevant to this study are optical and microwave atomic clocks. The difference between optical and microwave clocks is their oscillation frequency — microwave clocks have frequencies in the gigahertz range while optical clocks have a frequency range close to petahertz. Although a higher frequency requires a more complex system, it enables finer division of a second, a measurement that increases the potential stability of a clock as well.[7]

The stability of an atomic clock is measured using Allan deviation (AD), which quantifies fractional frequency fluctuation at averaging time τ. Stability is a function of τ, therefore it varies with different values of τ. Along with stochastic error quantified by AD, deterministic error such as drift contributes to creating timing error.[1]

Table 1 compares the space-qualified AFS candidate and optical clock long-term stability and SWaP. A trade-off between long-term stability and SWaP is evident, as lowest SWaP AFS have the lowest long-term stability, and the highest SWaP AFS have the highest stability.

Table 1. Comparison between long-term stability and SWaP in atomic clocks RAFS, PHM, DSAC, CSAC and optical clocks

	Long-term stability	SWaP
RAFS[9]	4×10^{-14} at 10,000s	2.4L, 3.3kg
PHM [9]	1×10^{-14} over 10,000s	28L, 18kg
DSAC[10]	$3\text{-}5\times10^{-15}$ at one day	19L, 19kg, 56W
CSAC[11]	10^{-11} at one day	15 cm^3, 125 mW
Optical clocks (only the laser cooling system)[12]	-	306 L, 163.8 kg, 615W

Rubidium atomic frequency standard

Rubidium atomic frequency standard (RAFS) is a microwave frequency standard. The RAFS physics package includes an ^{87}Rb-filled glass cell placed in a microwave cavity tuned near the 6.834 GHz Rb ground-state hyperfine transition frequency. It has a low SWaP of 3.3kg and 2.4L (power not reported), compared to other AFS and a long-term stability of 4×10^{-14} at 10,000s.[9,13] Flight heritage consists of mid earth orbit deployment on satellites such as GALILEO.[8]

Passive hydrogen maser

Passive hydrogen maser (PHM) is one of the most common atomic clocks onboard Global Navigation Satellite Systems, for instance GALILEO. Its physics package utilizes nested vacuum systems. The inner vacuum contains the oscillating hydrogen atoms while the cavity around it provides thermal insulation. Operates at a high SWaP of 28L and 18kg (power not reported), and a moderate stability of 1×10^{-14} over 10,000s.[9] The flight heritage consists of mid earth orbit deployment on GALILEO.[8]

Deep space atomic clock

Deep space atomic clock (DSAC) is a mercury trapped-ion clock designed in NASA's JPL to test the technology. Since SWaP was not the priority for the 2019 in-orbit demonstration, it has high SWaP at 19L, 19kg, 56W. The results from in orbit testing produced a high stability of $3\text{-}5\times10^{-15}$ at one day. Flight heritage consists of 2 year low earth orbit flight.[14]

Chip-scale atomic clock

Chip-scale atomic clock (CSAC) is a convergence of atomic spectroscopy and semiconductor technology made in an attempt to miniaturize atomic clocks. This AFS has the lowest SWaP, running on 125 mW and being chip sized, around 15 cm^3, its stability around 10^{-11} at one day. Flight heritage consists of low Earth orbit flights.[11]

Optical clocks

While optical clocks have the potential to have higher stability than microwave clocks due to operation at higher frequencies, the technology is in early development stages and confined to laboratories.[7] Discussed in detail in section 4.

As seen in table 1, none of the AFS satisfies all criteria at once, since high stability is a result of a physics package with increased vulnerability and SWaP. A low long-term stability increases position error because the observable is referenced to an onboard AFS, accumulating stability error over holdover intervals, which can extend to multiple days in one-way tracking. A durable AFS is crucial for longer missions, as the risk of system degradation mid-mission is smaller. A low SWaP is preferred to minimise the mass and power resource consumption of the navigation system on a spacecraft, that would otherwise be designated to payload. Therefore, the development of the three features of onboard atomic clocks is essential to establish a viable one-way radiometric tracking system.

Engineering challenges

Atomic clock implementation to utilize one-way tracking faces the following challenges: drift over a holdover interval; SWaP under deep space constraints; space survivability; and onboard autonomy. It is important to note that AFS flight heritage consists entirely of low to mid Earth orbit missions and demonstrations, therefore additional challenges may arise during deployment to deep space.

Stability and drift over holdover interval

To navigate independently, a spacecraft must accurately calculate its position, while without contact with DSN ground stations.[5] This period of time is the holdover interval. The reliability of a spacecraft's navigation during the holdover interval depends on the stability of the onboard clock, for instance a spacecraft with a higher stability AFS able to navigate autonomously with lower error. However, maintaining stability with little error onboard a deep space vehicle is challenging as the filter cannot remove the stochastic error.[1]

Over the holdover interval timing error is accumulated, which translates to position error, but it can be mitigated by applying filters to the calculations. In a 2021 simulation testing one-way radiometric tracking using DSAC when approaching Mars, T. Ely et al. use filters to remove deterministic error.[5] Deterministic error consists of drift, the AFS frequency systematically increasing or decreasing over time, and bias, a constant frequency offset.[1] Despite the removed error, filtering does not remove stochastic error occurring in the DSAC. Higher stability, represented by low AD, of an AFS shows low stochastic error.[1]

For this reason, high stability AFS maintain low position error better during a long holdover interval. For instance, a high stability AFS, such as the DSAC with a stability of $3\text{-}5\times10^{-15}$ at one day, would accumulate a lower timing error after a 3-5 hour holdover interval, than a low stability AFS, such as CSAC with a stability of 10^{-11} at one day.[13,15] Error accumulated over the DSN scheduled time of 3-5 hours is not measured for either AFS in the demonstration missions cited, so for the purpose of calculations a one day holdover interval is used. The position error at one day is 0.08-0.13 meters for the DSAC, while the position error for a CSAC is 260 meters.

Increasing stability to the system poses a challenge, as the onboard system is limited by resource constraints. A high stability AFS, such as a DSAC or a PHM, requires a complex system that increases SWaP.[10] Therefore, the maximum stability is limited by hardware constraints.

However, as AFS transfer from ground to low Earth orbit and require design adjustments, their long-term stability falls. In pre-flight testing, the DSAC successfully reached a low AD of 2×10^{-15} at 1 day, yet DSAC's performance in low Earth orbit is poorer with a higher AD of $3\text{-}5\times10^{-15}$ at 1 day.[1,10] Adjustments in the ion-trap configurations, due to failure in orbit (neon depletion) contributed to the lower stability in space as well.[10] Although the ground testing conditions aimed to mimic the space environment, it cannot be inferred that the stability will remain high in more adverse environments, such as deep space.

SWaP under deep space constraints

High stability AFS that reduce position error are a result of complex physics packages (laser-cooling systems, vacuums, cold-ion traps) that increase SWaP: higher volume, heavier weight and high power consumption.[13] Deep space probes are powered by radioisotope thermoelectric generators (RTG), which can last for longer periods due to slow power budget

decay, but have low power output and a finite power supply. Small power supply limits the use of equipment in physics packages such as traps, vacuums and lasers, that increase AFS stability. The least stable AFS is a CSAC at 10^{-11} at one day stability, which consumes 125 mW power.[11,13] In contrast, the most stable space qualified atomic clock DSAC at stability 3×10^{-15} at 23 days consumes the most power at 56W.[10] In the middleground stand PHM with a moderate stability yet a high SWaP and RAFS with low stability and a low SWaP relative to DSAC and PHM.[9]

However, deep space conditions such as a vacuum provide conditions that allow components for extreme condition creation in ground application to be discarded, reducing the mass. For example, PHM needs two vacuums in the ground version — one where physics of free hydrogen atoms takes place and one for insulation, to keep the thermal conditions stable. A space based PHM can utilize the surrounding space vacuum in place of the second vacuum.[16]

Optical clocks are not yet space qualified, but even the supporting laser-cooling system has higher SWaP than full systems of AFS, yet its long term stability in a laboratory setting exceeds microwave atomic clocks. A high stability offers a potential alternative to microwave AFS, although its physics package must undergo miniaturization to become space qualified.[7]
Space survivability and durability

Deep space has some of the harshest conditions including radiation from the sun, as there is no atmosphere to protect the hardware, extreme temperature fluctuations, and vibrations during launch. These hazards deplete the predicted lifetime of an AFS, that is already under constraint. Currently, the survivability challenge is estimated using ground-simulation of launching conditions and low to medium Earth orbit mission data.[5,17] New challenges may arise as AFS are deployed to deep space.

Radiation

AFS must be radiation hard with respect to ionizing radiation that will be encountered in deep space due to the lack of protection from an atmosphere. The hardware must be resistant enough to be reliable for the whole mission's duration, which is a challenge since deep space missions can last multiple years.[13] AFS are less sensitive to space radiations compared to quartz, therefore improvement of the driving quartz USO durability is imperative.[18] Additionally, electronics onboard deep space missions are at risk of reduced stability due to long term exposure (more than 3 years) to ionizing radiation. Multi-layered shielding is under development to prevent electronic system degradation.[19] Shielding can be applied to the computer hardware designated to position calculations and avoid error outside of the AFS.

Thermal fluctuations

Extreme temperatures are encountered by the spacecraft in space vacuum so systems that maintain thermal stability must be integrated. However, some thermal fluctuation inside the clock is inevitable despite the control system, so low thermal sensitivity has to be ensured. DSAC has shown no anomalous behaviors in between −12 ◦C and +30 ◦C.[20] Though thermal sensitivity is defined slightly differently across sources, the DSAC has a temperature sensitivity

of 1.3×10^{-15}/°C.[10] The PHM and RAFS have temperature sensitivities of 3×10^{-14}/°C and < 5×10^{-14}/°C respectively.[9,10] The CSAC temperature sensitivity is in the order of 2×10^{-11}/°C, as this is the case with most vapor-cell atomic clocks.[21] The AFS follow the same ranking based on thermal sensitivity as with stability — DSAC has the highest stability as well as least thermal sensitivity, followed by PHM, RAFS and CSAC. The physics package that provides high stability isolates the AFS from thermal fluctuations as well.

Vibrations during launch

Atomic clocks that are launched into space experience intense vibrations that can damage the clock hardware and design coherence. Therefore high durability is crucial. The DSAC clock has been tested against and survived a 14 g vibration for 2 minutes, which exceeds the launch loads.[20] PHM and RAFS have successfully passed vibration tests in preparation for deployment on GALILEO.[9] A CSAC was tested in 3 g RMS vibration, lower than experienced during launch yet enough to demonstrate robust design, and experienced some short term performance degradation.[22]

Hardware lifetime

Hardware lifetime is limited by the design of the AFS and is further depleted due to deep space hazards. PHM aging is caused by the limited hydrogen container capacity, bulk getters, ion pumps, and the total dose of ionising radiation. Despite the constraints the PHM on Galileo is anticipated to have a 12-year on-orbit lifetime.[9] DSAC lifetime is around 3-5 years, although a performance-degrading mechanism was discovered during testing after years of development.[10] RAFS lifetime is longer since Rb resources are not depleted as H resources are in a PHM.[13] CSAC employs a laser that has the longest lifetime in cold temperatures, however, a room temperature has to be maintained to ensure clock stability, offering a trade-off.[11]

All AFS passed testing in simulation and low to medium Earth orbit, with DSAC being the most recently tested for autonomous deep space navigation in low Earth orbit. However, the challenge is to deploy the AFS on a deep space demonstration to gather field data.

Onboard autonomy

For space navigation to move from two-way tracking to one-way tracking, the performance of AFS onboard must at least be on par with the performance with atomic clocks in DSN ground stations. Currently, only the DSAC stability is close to that of ground clocks.[14] The PHM and RAFS stability is sufficient for low and mid Earth orbit missions, as they are deployed on GALILEO, however below the ground clock standard.[16] Similarly, CSAC stability falls short of the ground clock standard, as miniaturization is prioritised in the design.[11] At present, most AFS error is too large for one-way tracking to reach two-way tracking level of accuracy.

Despite its high stability, two-way tracking demands continuous signal transmission from a DSN antenna during the communication pass, limiting the number of spacecraft that can be tracked simultaneously. In contrast, one-way tracking does not require a dedicated antenna for

the mission and can cruise for longer without contact with DSN ground stations. Using one-way tracking, a spacecraft navigates independently from several days to a week when cruising, which is followed by a communication pass of 8 hours or more.[1] Thus, a single DSN antenna can support a higher quantity of simultaneous missions.[23]

However, one-way tracking is limited due to switching between DSN ground stations and onboard processing constraints. As a consequence of one-way measurement, the onboard filter experiences hand-off uncertainty, additional error due to switching between DSN ground stations as the Earth rotates when a new ground station acquires the signal.[14] The hand-off uncertainty adds to the total timing error. Additionally, the filtering process uses Linearized Kalman Filter (LKF) and must take place onboard. It requires a lot of computational power, which is a limited resource, so in a 2021 simulation Ely et al. propose to filter without smoothing (iterating the process to increase accuracy). Additionally, only forward mapping is used to cut the computational burden.[5]

Optical navigation is already used in autonomous deep space navigation. When one-way radiometric tracking is fused with optical navigation, the the position estimate is improved by almost an order of magnitude.[5]

Optical clocks

Optical clock stability averages down to 6.6×10^{-19} at 1,800 s in laboratory setting, a substantially lower AD than that of a microwave atomic clock.[24] High stability optical clocks could decrease the timing and position error in one-way tracking facilitating the transition to autonomous deep space navigation, yet SWaP requirements for space-qualification are not yet fulfilled.

Despite rapid development, system complexity prevents optical clocks from being deployed to space for autonomous navigation purposes. Optical clocks require optical frequency combs (OFC) to bridge the optical frequency of several hundred terahertz to frequencies of tens of gigahertz that are accessible to frequency tracking electronics. Another complex system that optical clocks require is a multiple laser system, including a laser-cooling system to a few microkelvin. The bulkiness, high power consumption and cost of both the OFC and laser system technologies confine optical clocks to laboratories.[25] However, rapid developments, such as miniaturization of optical clocks to create chip-scale optical clocks, are happening to make the technology applicable in field settings.[8]

The OFC generators were successfully miniaturized and integrated on-chip, however the supporting electronic and optical system remains bulky, subverting any advantages gained by miniaturization. Thus OFC technologies are not yet deployable to space.[25] In contrast, laser cooling alkaline-earth atoms for optical clocks was achieved in the Chinese Space Station in 2025. The laser system is one of the two technology bottlenecks preventing optical atomic clocks from being space-qualified. The demonstration showed that solutions towards the elimination of the bottleneck are being developed. For instance, system survivability was demonstrated, the

system withstanding vibrations up to 3.4g at 17-40Hz and temperature fluctuation from 10-40°C.[12]

Although the experiment was successful, there are remaining challenges to be resolved. Firstly, the laser-cooling system, which is just one part of an optical clock, had a higher SWaP than any of AFS. Secondly, the demonstration was done in a structured environment of a space lab, whereas the clock must withstand extreme conditions to become space-qualified. Currently, spacecraft deployed on deep space missions are not equipped with an appropriate power supply and cannot support such a large system. The next bottleneck is miniaturizing the laser system and OFC to facilitate space deployment.

Conclusion

No single atomic clock is universally deployable to all deep space missions, since each atomic clock fails at least one of the following requirements: high long-term stability, low SWaP and flight survivability. Therefore, an atomic clock is selected depending on the prioritized requirement for the mission.[13]

RAFS and CSAC design prioritize low SWaP rather than long term stability, while the PHM and DSAC have a higher long term stability as well as a substantially higher SWaP. None of the AFS dominate.[9,10] Flight survivability has been tested in simulation and low to mid Earth orbit, accounting for challenges such as radiation, thermal fluctuation and vibrations, however, no AFS have been deployed to deep space, leaving deep space survivability unproven for all AFS.[20] Additionally, onboard autonomy is limited by hand-off uncertainty and simplification of filtering, two factors that add error to the position calculations in one-way radiometric tracking and must be accounted for.[5,14] Therefore, to facilitate autonomous navigation, the onboard atomic clock is mission-dependent.

Optical clocks are an ultra stable potential alternative to space-qualified microwave atomic clocks. High stability ensures minimal error accumulation over the holdover interval, however optical clocks have high SWaP. Future work could be done on the two bottlenecks currently preventing optical clocks from becoming space qualified: miniaturizing laser equipment and supporting equipment of OFC as well as reducing their power consumption.[25] Alternatively, more technology demonstration missions could be deployed to a deep space environment to verify that the data collected from simulations and low to mid Earth orbit still stands.

AI Acknowledgement

The author declares that no artificial intelligence (AI) tools were used in the preparation of this work.

Acknowledgements

The author thanks Dr. Devin Carroll and Kieran Tait for their supervision.

Works Cited

(1)Ely, T. A.; Burt, E. A.; Prestage, J. D.; Seubert, J. M.; Tjoelker, R. L. Using the Deep Space Atomic Clock for Navigation and Science. IEEE Trans. Ultrason. Ferroelectr. Freq. Control 2018, 65 (6), 950–961. https://doi.org/10.1109/TUFFC.2018.2808269

(2)Stern, S. A. The New Horizons Pluto Kuiper Belt Mission: An Overview with Historical Context. 2007. https://doi.org/10.1007/s11214-007-9295-y

(3)Asmar, S. W.; Armstrong, J. W.; Iess, L.; Tortora, P. Spacecraft Doppler Tracking: Noise Budget and Accuracy Achievable in Precision Radio Science Observations. Radio Sci. 2005, 40 (2). https://doi.org/10.1029/2004RS003101

(4)Bhaskaran, S. Autonomous Navigation for Deep Space Missions. In SpaceOps 2012 Conference; American Institute of Aeronautics and Astronautics: Stockholm, Sweden, 2012. https://doi.org/10.2514/6.2012-1267135

(5)Ely, T. A.; Seubert, J.; Bradley, N.; Drain, T.; Bhaskaran, S. Radiometric Autonomous Navigation Fused with Optical for Deep Space Exploration. J. Astronaut. Sci. 2021, 68 (1), 300–325. https://doi.org/10.1007/s40295-020-00244-x

(6)Schmittberger, B. L.; Scherer, D. R. A Review of Contemporary Atomic Frequency Standards. 2020. https://doi.org/10.48550/arXiv.2004.09987

(7)Fortier, T. M.; Luiten, A. N.; Margolis, H. S. Optical Atomic Clocks: Defining the Future of Time and Frequency Metrology. Optical 2026, 13 (1), 143. https://doi.org/10.1364/OPTICA.575770

(8)Bandi, T. N. A Comprehensive Overview of Atomic Clocks and Their Applications. Biol. Eng. Med. Sci. Rep. 2023, 9 (1), 01–10. https://doi.org/10.5530/bems.9.1.1

(9)Rochat, P.; Droz, F.; Mosset, P.; Barmaverain, G.; Wang, Q.; Boving, D.; Mattioni, L.; Belloni, M.; Gioia, M.; Schmidt, U.; Pike, T.; Astrium, E.; Emma, F. The Onboard Galileo Rubidium and Passive Maser, Status & Performance. 2005. https://doi.org/10.1109/FREQ.2005.1573898

(10)Burt, E. A.; Prestage, J. D.; Tjoelker, R. L.; Enzer, D. G.; Kuang, D.; Murphy, D. W.; Robison, D. E.; Seubert, J. M.; Wang, R. T.; Ely, T. A. Demonstration of a Trapped-Ion Atomic Clock in Space. Nature 2021, 595 (7865), 43–47. https://doi.org/10.1038/s41586-021-03571-7

(11)Kitching, J. Chip-Scale Atomic Devices. Appl. Phys. Rev. 2018, 5 (3), 031302. https://doi.org/10.1063/1.5026238

(12)Zhao, G.; Xia, J.; Liu, Y.; Zhou, Y.; Zhou, C.; Guo, F.; Wang, W.; He, D.; Feng, M.; Liang, T.; Ren, J.; Xu, Q.; Meng, J.; Gao, F.; Shen, Y.; Lu, X.; Lu, B.; Wang, Y.; Hu, X.; Tan, W.; Zou, H.; Chang, H. Laser Cooling Alkaline-Earth Atoms for Optical Clock in Chinese Space Station. Chin. Phys. Lett. 2025, 42 (6), 063701. https://doi.org/10.1088/0256-307X/42/6/063701

(13)Batori, E.; Almat, N.; Affolderbach, C.; Mileti, G. GNSS-Grade Space Atomic Frequency Standards: Current Status and Ongoing Developments. Adv. Space Res. 2021, 68 (12), 4723–4733. https://doi.org/10.1016/j.asr.2020.09.012

(14)Ely, T.; Prestage, J.; Tjoelker, R.; Burt, E.; Dorsey, A.; Herrera, R.; Kuang, D.; Murphy, D.; Robison, D.; Stuart, J.; Wang, R.; Seubert, J. Deep Space Atomic Clock Technology Demonstration Mission Results. 2022. https://doi.org/10.2514/1.A35334

(15)Seubert, J.; Ely, T. A.; Stuart, J. Results of the Deep Space Atomic Clock Deep Space Navigation Analog Experiment. J. Spacecr. Rockets 2022, 59 (6), 1914–1925. https://doi.org/10.2514/1.A35334
(16)Berthoud, P.; Pavlenko, I.; Wang, Q.; Schweda, H. The Engineering Model of the Space Passive Hydrogen Maser for the European Global Navigation Satellite System Galileo. In IEEE International Frequency Control Sympposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum, 2003. Proceedings of the 2003; IEEE: Tampa, Fl, USA, 2003; pp 90–94. https://doi.org/10.1109/FREQ.2003.1274993
(17)Droz, F.; Mosset, P.; Wang, Q.; Rochat, P.; Belloni, M.; Gioia, M.; Resti, A.; Waller, P. Space Passive Hydrogen Maser - Performances and Lifetime Data-. In 2009 IEEE International Frequency Control Symposium Joint with the 22nd European Frequency and Time forum; IEEE: Besancon, France, 2009; pp 393–398. https://doi.org/10.1109/FREQ.2009.5168208
(18)Cibiel, G.; Boizot, B.; Boy, J.; Carlotti, J.; Cambon, O.; Devautour-Vinot, S.; Candelier, V.; Lamboley, J.; Guibert, P.; Largeteau, A.; Inguimbert, C.; Piccheda, D. Ultra Stable Oscillators Dedicated for Space Applications: Oscillator and Quartz Material Behaviors vs Radiation. In 2006 IEEE International Frequency Control Symposium and Exposition; IEEE: Miami, FL, USA, 2006; pp 814–822. https://doi.org/10.1109/FREQ.2006.275494
(19)Guan, S.; Fu, G.; Wan, B.; Wang, X.; Chen, Z. Multi-Objective Optimization and Reliability Assessment of Multi-Layer Radiation Shielding for Deep Space Missions. Aerospace 2025, 12 (4), 337. https://doi.org/10.3390/aerospace12040337
(20)Tjoelker, R. L.; Prestage, J. D.; Burt, E. A.; Chen, P.; Chong, Y. J.; Chung, S. K.; Diener, W.; Ely, T.; Enzer, D. G.; Mojaradi, H.; Okino, C.; Pauken, M.; Robison, D.; Swenson, B. L.; Tucker, B.; Wang, R. Mercury Ion Clock for a NASA Technology Demonstration Mission. IEEE Trans. Ultrason. Ferroelectr. Freq. Control 2016, 63 (7), 1034–1043. https://doi.org/10.1109/TUFFC.2016.2543738
(21)Hao, Q.; Xue, W.; Xu, F.; Wang, K.; Yun, P.; Zhang, S. Efforts towards a Low-Temperature-Sensitive Physics Package for Vapor Cell Atomic Clocks. Satell. Navig. 2020, 1 (1), 17. https://doi.org/10.1186/s43020-020-00017-9
(22)Lutwak, R.; Rashed, A.; Varghese, M.; Tepolt, G.; LeBlanc, J.; Mescher, M.; Laboratory, C. S. D.; Serkland, D. K.; Geib, K. M.; Peake, G. M.; Römisch, S. THE CHIP-SCALE ATOMIC CLOCK – PROTOTYPE EVALUATION. 2007.
(23)Ely, T.; Burt, E.; Cheung, K.; Tjoelker, R. The Benefit of Space Clocks for the Deep Space Network. Radio Sci. 2025, 60 (8), e2025RS008244. https://doi.org/10.1029/2025RS008244
(24)Oelker, E.; Hutson, R. B.; Kennedy, C. J.; Sonderhouse, L.; Bothwell, T.; Goban, A.; Kedar, D.; Sanner, C.; Robinson, J. M.; Marti, G. E.; Matei, D. G.; Legero, T.; Giunta, M.; Holzwarth, R.; Riehle, F.; Sterr, U.; Ye, J. Demonstration of 4.8 × 10−17 Stability at 1 s for Two Independent Optical Clocks. Nat. Photonics 2019, 13 (10), 714–719. https://doi.org/10.1038/s41566-019-0493-4
(25)Chang, L.; Liu, S.; Bowers, J. E. Integrated Optical Frequency Comb Technologies. Nat. Photonics 2022, 16 (2), 95–108. https://doi.org/10.1038/s41566-021-00945-1

Democracy Under Pressure: Institutional Decline in the United States and Hungary
By Alex Numbers

I. Abstract

Since 2010, Hungary has experienced significant democratic backsliding under Prime Minister Viktor Orbán. One of the most important changes in Hungary has been the erosion of press freedom through increased government influence over the media. This influence has occurred through changes to media regulation, media ownership, and economic pressure on independent news organizations. At the same time, Hungary's democratic decline has involved many other institutional changes, including changes to the judiciary, constitution, and electoral system. Recent scholarship increasingly treats these changes as part of a broader process of institutional capture rather than as separate reforms. This raises the question: To what extent do restrictions on press independence contribute to democratic backsliding, and why do their effects differ between Hungary and the United States?

This paper argues that restrictions on press independence are not the only cause of democratic backsliding. However, they can worsen democratic erosion by making it harder to hold the government accountable, limiting the information available to citizens, and making political competition more difficult. Hungary shows how government control of the media can contribute to the decline of democratic institutions. The United States faces similar issues, especially political division, concentrated media ownership, and declining trust in institutions. The difference is that strong constitutional protections and independent institutions make it much harder for the government to control the press directly. The comparison therefore shows that the effects of press restrictions depend on the broader political and institutional environment. When other democratic institutions are already weakening, restrictions on the press can help sustain and deepen democratic backsliding.

II. Introduction

A free and independent press is widely regarded as an important part of a democratic society because it allows citizens to make well-informed political decisions, exposes corruption when necessary, and holds governments accountable (Bermeo; Haggard and Kaufman). Without access to independent information, citizens have a harder time evaluating political leaders and government policies. The press therefore does more than report political events. It provides one of the main ways citizens can learn about what their government is doing.

This is especially important when looking at democratic backsliding. Democratic backsliding happens when elected governments slowly weaken democratic institutions while still keeping the appearance of a constitutional government (Bermeo). Rather than taking over government all at once, democratic erosion can occur through changes to laws and existing institutions. Governments can weaken checks on their power, make it harder for political opponents to compete, and reduce the independence of important institutions without getting rid of elections.

Recent scholarship emphasizes that this process usually involves several institutions, not just one. Riedl, McCoy, Roberts, and Somer describe several institutional pathways through which democratic backsliding can occur, including legislative capture and executive aggrandizement. Enyedi and Mikola similarly describe Hungary as a case of institutional transformation in which Fidesz increasingly concentrated political and economic resources. This scholarship matters because it suggests that the decline of press freedom should not be studied separately from the broader political changes occurring at the same time.

The press can play an important role in this process. When journalists can investigate government officials and report on institutional changes, citizens are more likely to understand what is happening and speak out against it. But when political leaders can influence the media through regulations, ownership, or financial pressure, journalists find it harder to do their job. This does not mean that restrictions on the press automatically lead to democratic backsliding. Instead, they can make it easier for other forms of democratic erosion to continue.

Hungary provides one of the clearest examples of this process. After the collapse of communism in 1989, Hungary established a multiparty democracy with competitive elections, constitutional protections, and relatively independent institutions. However, after Viktor Orbán and his Fidesz party returned to power in 2010 with a two-thirds parliamentary majority, the government enacted major constitutional, judicial, electoral, and media reforms (Human Rights Watch). Recent scholarship continues to describe these changes as part of a broader concentration of political power rather than simply a change in media policy (Enyedi and Mikola).

The media was one of the areas most affected by these changes. The government gained greater influence over media regulation, while politically connected businesspeople acquired major media organizations. State advertising also became an important source of financial support for pro-government outlets. Recent research by Ágnes Urbán argues that these changes have created a systemic form of state intervention in Hungary's media market rather than a series of unrelated policies. Independent media still exists, but it operates in a market where government influence has grown increasingly important.

Hungary is not the only democracy where press independence has become a concern. The United States has strong constitutional protections for freedom of the press, but American media organizations face distinct problems, including political division, concentrated ownership, declining trust, and changes in how Americans consume news. Recent research also complicates the argument that political polarization itself causes democratic backsliding. Druckman, Green, and Iyengar find limited evidence of a direct relationship between affective polarization and democratic backsliding in the United States, even though polarization can make democratic functioning more difficult.

This distinction is important. The problems facing the American media system should not automatically be treated as evidence of the same type of democratic decline found in Hungary. Instead, the comparison allows us to examine how media problems can affect democracies differently depending on the strength of other institutions.

The existing scholarship therefore raises an important question. If media problems exist in both Hungary and the United States, why has Hungary experienced much greater democratic backsliding? Existing research has examined Hungary's institutional transformation and media system, while other scholarship has examined polarization and democratic strain in the United States. This paper brings these discussions together by comparing how political power interacts with media independence in the two countries.

This paper argues that restrictions on press independence are most damaging when they occur alongside broader institutional weaknesses. In Hungary, government influence over the media has reinforced other changes to democratic institutions by reducing government accountability and limiting citizens' access to independent information. In the United States, media polarization and concentrated private ownership create democratic challenges, but constitutional protections and institutional competition make government capture of the press much more difficult.

This comparison also has limitations. Hungary and the United States have very different political histories, institutional structures, and media systems. This paper does not argue that the two countries are experiencing the same level of democratic decline. Hungary has experienced much more major institutional backsliding than the United States. Instead, the comparison examines how different forms of pressure on the press operate under varying political conditions.

III. Literature Review

Political science and other fields have studied the connection between press freedom and democracy extensively. Most scholars agree that an independent press matters because it informs people, allows different ideas to be discussed, and helps hold those in power accountable. However, scholars disagree on how much media restrictions contribute to democratic backsliding compared with other political and institutional changes.

One major debate concerns whether democratic backsliding should be understood primarily as a problem of individual leaders or as a broader institutional process. Nancy Bermeo argues that modern democratic backsliding often occurs through "executive aggrandizement," in which elected leaders gradually expand executive authority while weakening institutional checks and balances. Similarly, Steven Levitsky and Daniel Ziblatt argue that democracies can weaken when elected leaders weaken democratic institutions and norms from within. Haggard and Kaufman also describe democratic regression as a process involving changes to political institutions and the distribution of power between governments and opposition groups.

More recent scholarship develops this argument by examining the different ways institutions can be captured. Riedl, McCoy, Roberts, and Somer identify legislative capture, plebiscitary override, and executive power grabs as different pathways through which democratic backsliding can occur. Enyedi and Mikola's research on Hungary provides a particularly important example of this broader approach. They argue that Orbán and Fidesz carried out a carefully organized project of institutional transformation and concentration of political and financial resources. This supports the argument that Hungary's media restrictions should not be

treated as an isolated cause of democratic decline.

The press remains important to this process because democratic institutions depend partly on citizens' access to information about how those institutions are changing. Journalists can investigate government actions, report on corruption, and help opposition groups communicate with voters. When citizens have less access to independent information, it can become harder for them to understand changes to the judiciary, constitution, elections, or government policy.

Another important concept in the literature is "media capture." This occurs when a government influences the media through ownership, regulation, money, or other forms of pressure rather than simply censoring every critical journalist. This distinction matters especially for Hungary. Independent journalism has not completely disappeared, but government influence has changed the conditions under which independent organizations operate.

Recent research by Urbán argues that the Hungarian media system has undergone systemic restructuring, including ownership changes, restrictions on competition, state advertising, and pressure on independent outlets. This supports the argument that media capture can happen through economic and market structures rather than only through direct censorship. It also helps explain why the Hungarian government does not need to control every individual journalist in order to influence the information available to the public.

At the same time, recent research suggests that the relationship between media problems and democratic decline is not always straightforward. A comparative review by McNeill and Valentine argues that journalists in both Hungary and the United States face different pressures, with political pressure playing a greater role in Hungary and private consolidation raising different concerns in the United States. This supports the need to distinguish government influence from private media ownership.

The United States provides a different case. The First Amendment provides strong constitutional protections for freedom of the press, and most major American news organizations are privately owned. However, a relatively small number of corporations and wealthy individuals control American media ownership. Political division has also created highly divided media environments in which citizens often receive very different information about the same political events.

Recent research on American polarization also complicates the relationship between media division and democratic backsliding. Druckman, Green, and Iyengar find little evidence that affective polarization directly causes democratic backsliding in the United States. However, they argue that polarization can still make democratic functioning more difficult. This distinction supports this paper's argument: media polarization can create problems for democracy without automatically becoming democratic backsliding.

The existing literature therefore provides a strong foundation for comparing Hungary and the United States, but it also shows why the comparison needs to be made carefully. The important difference is not simply whether media organizations are politically biased or concentrated. The more important question is whether political actors can use government institutions to systematically weaken independent journalism while also weakening other

democratic institutions.

This paper builds on this scholarship by focusing on the relationship between press independence and broader institutional decline. Rather than arguing that media restrictions caused Hungary's backsliding, the paper examines whether they can reinforce political changes already underway. The comparison with the United States helps show why similar problems with polarization or media concentration can produce different consequences when the surrounding institutions are stronger.

IV. Understanding Democratic Backsliding and Press Freedom

For this paper, press freedom means more than the absence of formal government censorship. A free press requires that journalists and news organizations can gather and publish information independently and criticize political leaders without undue interference. This includes freedom from government censorship, editorial independence, journalist safety, access to government information, diversity of media ownership, and the ability to investigate public officials (Ruohonen).

This broader definition matters because media systems in the United States and Hungary differ significantly. In Hungary, the government has influenced the media through regulation, ownership, and money. The United States does not have a similarly government-controlled media system. Instead, private companies or wealthy individuals own many major news organizations. Because of these differences, comparing the two countries means looking at how much freedom journalists actually have to work independently, not just whether the government formally censors them.

A free press is important to democracy for several reasons. Journalists can hold the government accountable by looking into corruption, misconduct, and controversial decisions. The news also gives people the information they need to decide how they feel about political leaders and their policies. An independent press gives opposition groups a way to reach voters and challenge the people in power. Finally, the press fosters political discussion by letting people hear different points of view.

These functions make press freedom especially important during periods of political change. Without access to independent information, citizens cannot effectively evaluate government reforms. As a result, weakening the press can make other forms of democratic erosion easier to maintain.

However, not every problem with the media necessarily restricts press freedom. Government censorship differs from political bias, and government control differs from private ownership. Misinformation and political polarization can also make productive political discussion harder, but that does not necessarily mean the government is restricting journalists. This distinction matters for the comparison in this paper. Hungary provides an example of government-supported media influence, in which political power has shaped the media's regulatory, economic, and ownership structures. The United States offers a different example: privately owned media organizations operate under strong constitutional protections yet face

problems stemming from concentrated ownership and political division.

The point of comparing the two countries is therefore not to decide which one has a "better" media system. Instead, the comparison looks at how different kinds of pressure on the media affect government accountability and political competition.

To compare Hungary and the United States, this paper examines several aspects of press freedom. First, it examines government interference, including whether officials or institutions pressure, restrict, or punish journalists. Second, it looks at who owns major news organizations and how much control they have. Third, it considers economic pressure, such as whether the government can use advertising or funding to influence the media. Fourth, it examines whether journalists can investigate and criticize political leaders without facing retaliation. Finally, it examines whether people have access to a variety of viewpoints and sources of information. These factors are then connected to democratic backsliding. The main question is whether press restrictions reduce government accountability, limit citizens' access to independent information, and make political competition less meaningful. This also helps explain why similar problems can have different effects in Hungary and the United States.

V. Hungary

Following the collapse of communist rule in 1989, Hungary developed a multiparty democratic system with competitive elections, constitutional protections, and independent institutions. Its media system included public broadcasters as well as privately owned newspapers, television stations, and radio organizations. Although political influence over the media existed before 2010, journalists had greater opportunities to criticize government officials and provide competing political viewpoints.

This period provides a useful baseline for evaluating the changes that followed Viktor Orbán's return to power.

In 2010, Orbán and Fidesz won a two-thirds parliamentary majority. This gave the government the power to make major constitutional and institutional changes without support from opposition parties. The government subsequently changed the constitution, the judiciary, the electoral system, and the media regulatory structure.

These reforms were not separate from the other changes happening in Hungary. They were part of a larger effort to restructure the country's political institutions. Recent research by Enyedi and Mikola supports this interpretation, describing Hungary's transformation as a broad concentration of political and financial resources rather than a single institutional change. As a result, the decline in press freedom is closely connected to broader democratic backsliding in Hungary.

The 2010 Media Laws created the National Media and Infocommunications Authority and the Media Council, which were given broad powers over Hungary's media. Human Rights Watch criticized these changes because the Media Council's structure gave the governing party significant influence over media regulation.

The concern was not simply that Hungary created a regulatory body. The larger issue was

that an institution responsible for overseeing the media could become politically dependent on the government. This gave the government another way to influence the media without having to censor every critical journalist directly.

Government influence also expanded through changes in media ownership. Businesspeople with close relationships to Fidesz acquired newspapers, television stations, radio stations, and other media organizations. This process reached a major point in 2018 with the creation of KESMA, the Central European Press and Media Foundation. The foundation brought together more than four hundred media outlets (Human Rights Watch; Reporters Without Borders).

Recent scholarship helps explain why this change matters. Urbán argues that Hungary's media restructuring cannot be understood as the result of a single policy. Instead, ownership restructuring, restrictions on competition, state advertising, and pressure on independent media have interacted over time. This means that the government's influence has been built into the broader structure of the media market.

Independent journalism did not disappear entirely. Independent news outlets could still operate, but they had to compete with pro-government organizations with much more funding and a larger reach.

This distinction is important because Hungary's media system is not simply one in which the government tells every journalist what to write. Instead, government influence has changed the conditions under which journalism operates. When pro-government organizations control a large share of the media market, independent outlets can keep reporting but reach fewer people and face greater financial pressure.

State advertising was another way the government could influence the media. Government agencies and state-owned companies gave much of their advertising money to pro-government outlets, while independent outlets received much less. Human Rights Watch has raised concerns about this unequal distribution of state advertising and the financial pressure it puts on independent media.

This put financial pressure on independent media organizations. Unlike direct censorship, the government did not have to stop a publication from criticizing it. Instead, reducing financial support made it harder for independent outlets to keep operating and reporting independently. Together, government regulation, concentrated ownership, and financial pressure made it harder to maintain a diverse media system. Independent journalists could still investigate the government, but they faced greater challenges getting information, accessing government officials, and keeping their organizations financially stable.

These changes matter because of the press's role in holding the government accountable. When journalists have fewer resources or less access to political information, people have fewer chances to learn about what the government is doing. This can make it harder for citizens to hold political leaders accountable through elections or public debate.

Evidence also shows that Hungarian journalists and media organizations have faced other forms of pressure. Human Rights Watch has reported that journalists have had limited access to

public officials and government information. Freedom House has also found that independent journalists and media outlets face government pressure and lack the same access to state resources as pro-government outlets.

This shows why press restrictions should be understood as more than direct censorship. A journalist does not have to be directly censored for their freedom to be affected. If journalists face financial pressure, have limited access to government officials, or fear retaliation, it becomes harder for them to investigate and report on government actions.

The evidence suggests that restrictions on the Hungarian press have contributed to democratic backsliding by reinforcing other changes to the country's institutions. The media environment affects how citizens learn about elections, government reforms, corruption, and opposition parties. When access to this information becomes more limited, citizens may find it harder to understand what the government is doing and to hold political leaders accountable. The press therefore links citizens to other democratic institutions. When that link weakens, citizens may have fewer opportunities to understand or challenge changes to the judiciary, the constitution, or the electoral system.

However, it would be inaccurate to argue that media restrictions alone caused Hungary's democratic decline. Hungary's government also changed other major institutions, and these changes helped concentrate political power. Enyedi and Mikola's research supports this broader interpretation of Hungary's decline.

This creates a more complicated relationship between press freedom and democratic backsliding. Political power made it easier for the government to influence the media, and media influence, in turn, made it easier for the government to sustain that power. Restrictions on the press were therefore both a result of Hungary's democratic backsliding and a mechanism that helped sustain it.

VI. The United States

The United States is a different example of democratic strain. The country still has strong constitutional protections for freedom of the press, competitive elections, and a private media system (U.S. Constitution, amendment. I). However, the U.S. has also faced growing political division, lower trust in institutions, and disagreements over democratic norms. These issues have raised concerns about how well democratic institutions can handle political conflict.
The American case, therefore, does not represent the same level of democratic backsliding as Hungary's. Instead, it allows us to examine whether problems within a legally protected media system can still affect democratic accountability.

Most American news organizations are privately owned rather than government-controlled. The First Amendment also provides strong protections against government restrictions on freedom of the press (U.S. Constitution, amendment. I).

However, a small number of large companies and wealthy individuals own many of the major media organizations in the United States. As a result, the press can be independent of the government but still influenced by the interests of the people and companies that own these

organizations.

This matters because private ownership does not necessarily mean the press is restricted. However, when a small number of people or companies own many media outlets, they can significantly influence what stories people see and how political issues are covered. Ultimately, this can shape the information that reaches the public.

Political division has also contributed to highly partisan media environments. Cable news, online publications, podcasts, and social media platforms often attract audiences with particular political viewpoints.

Unlike Hungary, this polarization is not mainly caused by government control over the media. Americans can generally choose from many different news sources. However, people may be more likely to choose sources that support their existing political beliefs.
This can weaken democratic debate because citizens may encounter completely different interpretations of the same political events. When people receive different information about government actions, elections, or political candidates, it becomes harder to have a shared political discussion.

At the same time, recent scholarship suggests that polarization should not automatically be treated as a cause of democratic backsliding. Druckman, Green, and Iyengar find limited evidence that affective polarization directly causes democratic backsliding in the United States. They argue instead that polarization can make democratic functioning more difficult, but other factors also matter.

This distinction strengthens the comparison with Hungary. If political division by itself were enough to cause democratic backsliding, the United States would be much closer to Hungary than it is. Instead, the American case suggests that polarization can weaken democratic debate without necessarily allowing the government to capture the press.

Rupert Murdoch's ownership of Fox News shows how much influence private media owners can have. Fox News has a strong political identity and a large audience, so its leadership can significantly shape how the public discusses and understands political issues.

However, this is not the same as the Hungarian government controlling the media. Fox News is privately owned and operates under laws that allow it to criticize political leaders and publish different kinds of information. Other news organizations can also criticize Fox News and give people different viewpoints. This makes the media environment in the United States different from Hungary, where the government has much more influence over the media.
The American case therefore demonstrates why media concentration and government censorship should be treated as separate issues.

Despite polarization and concentration, American journalists continue to investigate and criticize political leaders from both major parties. Investigative reporting has exposed government misconduct, corruption, and controversial decisions by elected officials.

These examples show why strong constitutional protections are still important. Political leaders can criticize news coverage they dislike, but they cannot easily use the government to control news organizations because legal and institutional limits make that difficult.

The continued presence of independent journalism therefore provides an important protection against the kind of media capture seen in Hungary.

However, the American media environment can still create democratic problems. Political division may increase distrust between citizens, while concentrated ownership can give a small number of organizations major influence over political information.

These problems matter because they show that a country does not need government censorship to have a media system that challenges democracy. However, they also show why media problems should not automatically be described as democratic backsliding.

The difference between the United States and Hungary lies in how far political power can be used to control the media environment. In the United States, private owners can influence political discussion, but the government does not control their competitors in the same way. In Hungary, political power has shaped regulation, ownership, and financial resources across much of the media system.

The United States therefore shows that media concentration and polarization can create democratic problems without necessarily producing the same institutional decline that has occurred in Hungary.

VII. Comparative Analysis and Evaluation

The comparison between Hungary and the United States shows that press freedom can be threatened in different ways. Both countries face media concentration and political division, but the government plays a very different role in each.

Both countries have media environments that can contribute to political division. People can choose news sources that match their political beliefs, and major media organizations can strongly influence how political issues are discussed. This can make it easier for people to stay within their own political viewpoints and deepen political divisions.

Both countries also have concentrated media ownership. A relatively small number of powerful organizations or individuals control a large share of the news market. This shows that media concentration is not limited to authoritarian or backsliding governments.

The biggest difference is the government's relationship with the media. In Hungary, the government has used regulation, state advertising, and connections with politically connected media owners to influence the media. These actions have made it harder for different viewpoints to have an equal place in the media and have helped strengthen the governing party's position. In the United States, media organizations remain largely independent of government control. Private ownership and partisan division can create problems, but the government does not exercise the same level of control over media regulation, ownership, or financing.

The countries also differ in the strength of their institutions. In Hungary, media changes occurred alongside changes to the judiciary, elections, and the Constitution. In the United States, strong constitutional protections and institutional competition still make it harder for the government to take control of the press.

The comparison therefore suggests that the effects of media problems depend on the

wider political environment.

In Hungary, media restrictions reinforce other forms of institutional erosion. When the government can influence the media while also changing electoral and judicial institutions, information restrictions can make those wider changes easier to sustain.

For example, if citizens receive less independent information about changes to the electoral system or judiciary, they may have fewer opportunities to understand how those changes affect political competition. The media therefore indirectly affects democratic backsliding by shaping the information available to citizens.

In the United States, media polarization can increase political distrust and division, but the press remains independent enough to make government control harder. Even when individual news organizations are very partisan, other organizations can challenge their reporting and provide different viewpoints.

This difference is important because it shows that political bias alone is not enough to explain democratic backsliding. Recent scholarship on American polarization supports this more cautious conclusion. Polarization can make democratic institutions work less effectively, but it does not automatically produce the type of institutional capture seen in Hungary.

An important counterargument also merits consideration. One could argue that Hungary's democratic decline primarily results from constitutional, judicial, and electoral changes, and that media restrictions mostly follow this broader political transformation. Research supports this view, showing that Hungary's backsliding involved a much broader concentration of political and economic power (Enyedi and Mikola).

This counterargument matters because it prevents exaggerating the press's role. Restrictions on the press did not, by themselves, create Hungary's democratic backsliding. Fidesz's political dominance and changes to other institutions were necessary parts of the process.

However, the evidence also suggests that media restrictions helped sustain the changes that had already taken place. Once political leaders have greater control over the media environment, citizens have fewer opportunities to learn about government misconduct or challenge institutional changes.

The relationship is therefore not one-directional. Democratic backsliding can create conditions for media capture, and media capture can, in turn, reinforce democratic backsliding. This helps explain why press freedom is best understood as an amplifier of democratic conditions. When other democratic institutions are strong, media problems may create political instability or polarization without necessarily causing institutional collapse. However, when political power is already becoming concentrated, restrictions on the press can reinforce that concentration by reducing accountability and limiting citizens' access to alternative information. This argument also explains why Hungary and the United States can experience media problems without experiencing the same political consequences. The two countries have some similar problems, including media concentration and political division. However, the relationship between political power and media institutions differs sharply.

In Hungary, government influence over the media is part of a broader pattern of institutional change. In the United States, media organizations can be highly influential and politically divided, but they remain largely independent of government control.

The evidence from both cases suggests that press freedom is neither the only cause of democratic backsliding nor simply a result of it. Instead, it can be one way democratic erosion deepens and continues over time.

This also suggests an important direction for future research. A larger comparative study could examine whether the same relationship exists in other democracies experiencing different levels of institutional decline. Comparing Hungary with countries such as Poland, Turkey, or other democracies experiencing political strain could help determine whether media capture consistently becomes more damaging when other institutions are already weakening. Future research could also examine how social media, online news, and disinformation change the relationship between media independence and democratic accountability.

VIII. Conclusion

This paper asked: To what extent do restrictions on press independence contribute to democratic backsliding, and why do their effects differ between Hungary and the United States? This paper argues that restrictions on press independence can substantially contribute to democratic backsliding, but they are neither an independent nor a sufficient cause. Instead, restrictions on the press are most important when they occur alongside other forms of institutional decline. In these situations, weakening press independence can reduce government accountability, limit citizens' access to independent information, and make it more difficult for opposition groups to compete.

Hungary demonstrates this process most clearly. Since 2010, the government has changed media regulation, encouraged the consolidation of media ownership among political allies, and used state advertising and other economic resources to influence the media environment. Recent scholarship supports the view that these changes were part of a broader restructuring of Hungary's political and economic institutions, not an isolated effort to control journalists. These changes matter because they occurred alongside major shifts in the judiciary, Constitution, and electoral system. Media restrictions did not cause Hungary's democratic backsliding on their own. However, they helped reinforce it by making it harder for citizens and journalists to challenge the government as it gained more political power.

The United States is different because, although its media system has problems with political division, concentrated ownership, and declining trust, these issues are not the same as government-backed media capture. The First Amendment protects freedom of the press, and independent news organizations can still investigate and criticize political leaders. Recent research on polarization also suggests that political division should not automatically be treated as a direct cause of democratic backsliding.

The comparison therefore shows why press freedom should not be measured only by whether a government formally censors journalists. A healthy democratic media system also

requires editorial independence, diverse sources of information, access to government information, and journalists' ability to investigate political leaders.

At the same time, the comparison shows that not every media problem should be treated as evidence of democratic backsliding. Political division and private media ownership can undermine democratic debate without necessarily leading to government control of the press. The most important difference between Hungary and the United States is the relationship between political power and media independence. In Hungary, the government has used its power to influence regulation, ownership, and financial resources in the media system, while other democratic institutions have also weakened. In the United States, strong constitutional and institutional protections make it much harder for political leaders to gain that level of control. This means that restrictions on the press are better understood as part of a larger process of democratic erosion. When democratic institutions are strong, media problems may hurt political debate without causing bigger problems for democratic institutions. But when political power is already becoming more concentrated, restrictions on the press can make it easier for that concentration of power to continue.

Ultimately, Hungary shows what can happen when media independence weakens at the same time as other democratic institutions. The United States, by contrast, shows that media polarization and concentrated ownership can still cause problems even when the press is legally independent. Protecting democracy requires more than just protecting the right to publish. It also means ensuring people can access independent information and that journalists can hold political leaders accountable.

The comparison shows that the press is not only affected by a democracy's health. It can also affect how accountable and competitive democratic institutions are. When independent journalism weakens and political power concentrates, restrictions on the press can make democratic backsliding easier to sustain. Future comparative research can help determine how consistently this relationship appears across different political systems and whether protecting media independence can prevent institutional decline before it becomes more difficult to reverse.

Works Cited

Bátorfy, Attila, and Ágnes Urbán. "State Advertising as an Instrument of Transformation of the Media Market in Hungary." East European Politics, vol. 36, no. 1, 2020, pp. 44–65. Taylor & Francis, doi:10.1080/21599165.2019.1662398.

Bermeo, Nancy. "On Democratic Backsliding." Journal of Democracy, vol. 27, no. 1, 2016, pp. 5–19.

Druckman, James N. "How to Study Democratic Backsliding." Political Psychology, vol. 45, suppl. 1, 2024, pp. 3–42. Wiley, doi:10.1111/pops.12942.

Druckman, James N., Donald P. Green, and Shanto Iyengar. "Does Affective Polarization Contribute to Democratic Backsliding in America?" The ANNALS of the American Academy of Political and Social Science, vol. 708, no. 1, 2023, pp. 137–163. SAGE Publications, doi:10.1177/00027162241228952.

Enyedi, Zsolt, and Bálint Mikola. "Legislative Capture in Hungary: Well-Managed Autocratization." The ANNALS of the American Academy of Political and Social Science, vol. 712, no. 1, 2025. SAGE Publications, doi:10.1177/00027162241307778.

Freedom House. Freedom in the World 2026: Hungary. Freedom House, 2026.

Haggard, Stephan, and Robert R. Kaufman. Backsliding: Democratic Regress in the Contemporary World. Cambridge UP, 2021.

Human Rights Watch. World Report 2026: Hungary. Human Rights Watch, 2026.

Levitsky, Steven, and Daniel Ziblatt. How Democracies Die. Crown, 2018.

Riedl, Rachel Beatty, Jennifer McCoy, Kenneth M. Roberts, and Murat Somer. "Pathways of Democratic Backsliding, Resistance, and (Partial) Recoveries." The ANNALS of the American Academy of Political and Social Science, vol. 712, no. 1, 2024, pp. 8–31.

Reporters Without Borders. "Hungary." World Press Freedom Index 2026, Reporters Without Borders, 2026.

Ruohonen, Jukka. "Reassessing Measures for Press Freedom." arXiv, 19 June 2021, arxiv.org/abs/2106.10427.

Urbán, Ágnes. "The Restructuring of the Hungarian Media System." Corvinus Journal of Sociology and Social Policy, vol. 15, no. 3S, 2024, pp. 25–49. doi:10.14267/CJSSP.2024.3.2.

United States. Constitution of the United States. Amendment I.

The Rise of Content Creation among Gen Z: the Attractions and the Precarity

By Ethan Y. Gu (student)[1*], Kanishka Elupala (mentor)[2]

Abstract

The creator economy has grown into a global industry worth nearly half a trillion dollars, yet content creation remains deeply precarious – characterized by income instability, algorithmic dependency, intense competition, and burnout affecting over 79% of creators. Despite these risks, content creation continues to attract Generation Z at extraordinary rates. This paper examines why content creation persists as a career choice among Gen Z despite its inherent precarity. Drawing on the gig economy framework and Gen Z's unique relationship with digital environments, I argue that two interconnected factors sustain this appeal: push factors (declining traditional job availability, necessity entrepreneurship, network effects, stigma avoidance) and pull factors (Gen Z's platform-native identity, flexibility, purpose-driven work, low barriers to entry, and the perceived promise of overnight success). This paper argues that few alternatives drive a generation towards their most enticing option: content creation.

Keywords

Social Studies, Content Creation, Gen-Z, Labor Market, Jobs, Precarity

1. Introduction

The creator economy has emerged as one of the most significant labor transformations of the twenty-first century. Content creation, the production and sharing of digital content across social media platforms, has evolved from a niche hobby into a global industry projected to reach $480 billion by 2027, with an estimated 50 million creators worldwide, according to Goldman Sachs.[1] For Generation Z, a cohort characterized by digital fluency and lifelong internet immersion, the line between social and professional platform usage has become increasingly fluid, normalizing content creation as both a reputable career choice and an accessible pathway to income.

The appeal of content creation lies in its flexibility, low barriers to entry, and potential for recognition. Requiring little more than a smartphone, digital entrepreneurship removes traditional startup requirements, and places the creator as the face of their startup brand. The COVID-19 pandemic accelerated this trend through necessity entrepreneurship, as labor market disruptions pushed individuals toward accessible alternatives.[2] However, this accessibility masks a deeply precarious reality. The creator economy exhibits striking parallels to gig work: unpredictable income, absence of benefits, intense competition, and algorithmic dependency. The top 1% of creators capture the majority of earnings, while 52% earn less than $10,000 annually, and changes in variables like culture and algorithms can destroy a creator’s livelihood overnight.[3, 4]

This paper examines why content creation remains attractive to Gen Z despite its inherent

1. The Harker School, 629 Paco Drive, Los Altos, California, 94024, USA
ethan.gu.0909@gmail.com

precarity. Drawing on the precarious nature of platform-dependent labor, the gig economy framework, and Gen Z's unique relationship with digital environments, I argue that content creation persists (in spite of precarity) due to a confluence of push factors (labor market constraints, network effects, and social normalization), pull factors (flexibility and perceived viability).

2. Background

A content creator is defined as "a person who creates digital content (such as videos, blogs, podcasts, etc.) and shares it on social media websites and applications or other online platforms"[5]. This ranges from an individual creating dance videos on TikTok to creating educational blogs and podcasts. Due to the fact that Gen-Z possesses a sort of "digital fluency"[6], and has basically lived their whole lives on the internet, this results in the line between "using platforms socially" and "using platforms professionally" being increasingly fluid.

In recent years, the creator economy – an economy of individuals who create digital content – has grown substantially, with a Goldman Sachs report estimating a staggering 50 million content creators and projecting the global creator economy to reach $480 billion by 2027. In addition, the International Advertising Bureau (IAB) estimated that there are 1.5 million creators earning a full-time equivalent [FTE] of $66,622 per year, 7 times the number recorded in 2020 (200,000 earning FTE $55,000). The creator economy has not only grown in size, but its viability as a source of income has also increased. However, it is important to note that there are different levels of content creation, from being a hobby to full-time and professional content creation.[3]

Content creation, especially on social media, plays a key role in the lives of Gen Z, who have never known a world without the internet, and in 2012, about 3 in 4 Gen Z individuals owned a smartphone.[7] Fast forward to the present, and suddenly, about 99% of individuals aged 18-29 say they own a smartphone, and with a vast majority of adults(84%) stating that they use Youtube. As a result, many of Gen Z connect with society through apps like Instagram instead of meeting in-person, leading to widespread exposure and self-comparison with content and influencers. This immersion potentially explains the normalization of content creation as a key role in society, and as a full-time job, with about 57% of Gen Z stating that they would be influencers if given the chance, and that being an influencer is a "reputable career choice."[8]

3. Why content creation is growing (and what makes it so attractive to Gen Z) – reasons and drivers.

To start, let's connect content creation to something better understood: entrepreneurship. More specifically, digital entrepreneurship. Although entrepreneurship is a fundamental field of research, its definitions can still be nebulous. Given this, the idea of a digital entrepreneur may be even more nebulous, with no general or conventional definition available. For the context of my research, I will use the following definition: A digital entrepreneur is an individual "creating

and delivering digital content to capitalize on opportunities, aiming to expand the audience to a scale that allows for generating sustainable revenue."[3]

This definition frames digital entrepreneurship through the lens of a content creator, allowing us to compare and contrast content creation from traditional entrepreneurship. Contrasting from a regular business venture, content creators usually maintain full equity over their firm, resulting in their business becoming their personal and public image. As a result, many individuals partaking in content creation use it as a means not only to share hobbies, but to also generate income through their content. This spectrum from casual to professional content creation gives it a wide market of individuals who want to create content on the website.[3]

Additionally, what brings many individuals over to content creation are the low barriers to entry. As one creator puts it, ''All you need is a smartphone and an account."[3] This low barrier to entry dramatically expands participation by removing traditional constraints such as startup capital, formal credentials, or institutional backing.

This trend of flexible, low-barrier work aligns with the broader rise of the gig economy, a labor market that is based on short-term, freelance, or task-based work rather than traditional employment. In a gig economy, workers are classified as contractors rather than employees, which moves responsibilities regarding health, paid leave, and other amenities to the workers themselves. Content creation shares many characteristics to the gig economy: creators are self-employed, produce content when they want, are paid based on interaction (per "sale" to a viewer), and need to manage their health and retirement funds without the help of the platform.[9]

Therefore, content creation allows individuals to work wherever, whenever. As mentioned above, very little is required to be able to create content, and analyzing through a gig economy mindset, content creation prioritizes flexibility, which allows individuals to create the content they want to create.[9] In other words, content creation allows for individuals to do whatever they want, whenever they want, post it, and potentially earn a profit.[10]

However, low barriers to entry would not be the only factor in the growth of content creation, as many creators still desire to make a profit. Multiple streams of revenue possible from content creation – including sponsorships, ad revenue, and subscriptions – make content creation more profitable, and due to the algorithms of social media websites, the more successful creators get more exposure. These two factors result in content creation appearing more viable or profitable than a traditional job.[9,11]

Content creation can also be used to gain exposure, potentially allowing individuals to gain access to new opportunities, such as partnerships and future job opportunities. In a study conducted on the impact of social media exposure on the job search for economists, where reputable economics researchers promoted the works of other economists looking for jobs, it was revealed that the more increased exposure on posts also increased the number of job interviews individuals received in comparison to a control group.[11] This idea of content creators as promoters ties back to companies using content creators as a form of marketing, further increasing its appeal.[12]

Another factor shaping the appeal of content creation, especially since the COVID-19 pandemic, is the rise of necessity entrepreneurship. When individuals are unable to secure employment in their preferred sectors, they often turn to alternatives like content creation, which offer low barriers to entry and the possibility of financial gain.[3]

In this context, content creation functions less as a purely passion-driven pursuit and more as a response to labor market disruptions. Its accessibility makes it an attractive fallback option, allowing individuals to generate income independently when traditional opportunities are limited, even if the likelihood of sustained success remains uncertain.[3,13]

An example of this would be during the COVID-19 pandemic. When individuals were driven out of work, they were forced to partake in jobs like live streaming on Twitch. A study was conducted to analyze the effects of Covid-19 on those already streaming and those who joined after. Interestingly, those who joined after were more motivated to stream for longer durations and more frequently than those who were already streaming before.[2] By definition, these individuals are necessity entrepreneurs, and are being pushed to content creation not because it is the best paying job out there, but because it is their best option at hand.

Research on Gen Z's job preferences reveals a consistent tension. On one hand, Gen Z values flexibility.[14] On the other hand, they also value high salary and potential for growth. A study on Gen Z job seekers, though focused primarily on international STEM students, identified four key tenets that Gen Z considers when choosing and staying with an employer: flexible hours, supportive environments, growth opportunities, and high salaries.[15] These findings suggest that Gen Z wants both autonomy and security. Similarly, another study finds that Gen Z is most likely to prioritize instrumental (Salary and Job Security) work values over cognitive, social and prestige work values.[16] Now, in an increasingly scarce job market, Gen Z faces a dilemma: security and stable pay or a more flexible and free schedule.

4. What makes content creation precarious – risks, instability, algorithmic control, gig economy similarities.

One way content creation can be considered precarious is through its gig-like nature. Precarity in this context refers to uncertain, unstable, and insecure forms of labor, where individuals lack predictable income, benefits, or legal protections. The gig economy – through ride-hailing, delivery, freelance platforms, and short-term contracts – develops this precarious nature by removing the liabilities of an employer. While gig work offers flexibility and independence, it shifts much of the risk from employers to workers: workers often lack paid sick leave, retirement contributions, health insurance, and stable earnings. For many, this creates a cycle of financial fragility disguised as independence.

Content creation operates in a remarkably similar way, particularly in terms of worker security. Creators have significant control over their schedules, the content they produce, and how they earn money, but platforms generally do not provide the protections associated with traditional employment. Creators can experience highly unpredictable income, lack employer-provided benefits, and have little protection against sudden changes in platform

policies or algorithms that can reduce their visibility and earnings. In this sense, content creation resembles gig work not simply because of its flexibility, but because the responsibility for managing financial and occupational risk falls largely on the individual worker. The freedom that makes content creation attractive therefore also contributes to its precarious nature: creators gain independence from traditional employers while simultaneously losing many of the protections provided by traditional employment.[9]

Furthermore, the reality of content creation is that success only comes from the top 1%. As shown in the IAB report, only 1.5 million of the 50 million estimated in the Goldman Sachs report are making a FTE revenue of $66,622.[1] In addition, due to the low barriers of content creation, the increasing supply of content leads to higher and higher levels of competition – where creators become interchangeable and usually don't stand out. As a result, many creators (52%) earn less than $10,000, and almost all podcasts are unprofitable.[3]

In addition, the ever-evolving environments and algorithms make it extremely difficult to stay relevant. As a result, constant promotion, posts and adaptation to current and relevant trends becomes paramount to maintaining relevance and revenue.[10] The precarity described here results in individuals working long hours with little returns, often leading to burnout. In fact, over 79% of creators have experienced this burnout, further revealing the precarious nature of this industry.[3] Putting that number into context, it is quite common for one creator to lose their fanbase based on one controversial action or simply a change in algorithms and cultures.[17] As soon as a trend or era of content ends, creators must find ways to rebrand, or potentially lose it all.

Elaborating on algorithmic changes, website regulation, specifically shadow banning, proves particularly problematic. Unlike outright bans or content removals, shadow banning operates invisibly, where content is hidden from users, essentially restricting the creator's visibility. This form of stealth content moderation leaves creators in a state of uncertainty, unable to determine if declining engagement stems from poor content, shifting audience interests, or deliberate algorithmic suppression. The ambiguity forces marginalized creators especially to develop "folk theories" through collaborative investigation with other users, attempting to prove shadow banning by monitoring engagement metrics and asking followers to verify visibility. For creators reliant on platform income, the financial consequences are direct and severe: reduced visibility translates to lost subscribers, brand deals, and revenue streams. Meanwhile, platforms deny shadow banning exists—a practice researchers term "black box gaslighting"—leaving creators without recourse or appeal mechanisms. This invisible form of algorithmic control adds yet another layer of precarity, as creators' livelihoods can be quietly undermined without explanation or accountability.[18]

5. Why it remains attractive despite the precarity – analysis (e.g., necessity entrepreneurship, lack of other options, or risk-taking mindset).

However, the question remains: why is content creation still so popular among Gen-Z when it is so precarious? In order to answer this question, I will divide my response into 2

separate sections: the push factors, the pull factors. These sections will include ideas already covered in the paper like necessity entrepreneurship, low barriers to entry, and flexibility, as well as new ideas, like the peer/network effect, stigmas, and the desire for fame.

i) Push Factors

From one perspective, Gen Z is not taking up content creation because they want to, but because they are forced to adapt to the ever-evolving job market landscape. Recent BLS reports indicate that the number of vacant jobs on the market have been steadily decreasing since its peak of 11.3 million in 2022, now sitting at about 7.6 million unfilled jobs.[19, 20] While this trend could also indicate the steady growth of the US economy – more people taking up employment and filling jobs – the unemployment level's steady rise – from 5.8 million in December 2022 to 7.3 million in March 2026 – tells a different story, indicating that many individuals looking for employment are staying unemployed for large periods of time.[21] Due to this slow hiring, Gen Z might have to take up other forms of employment to keep afloat as costs of living increase. Similar to a recession, necessity entrepreneurship could push people to take up jobs that are not ideal in terms of stability because it seems like the next best option. Therefore, the desire for income – no matter the kind – overshadows the importance of precarity in the market as Gen Z works around the limitations of the traditional job market.

The stigma behind unemployment also could push Gen Z to content creation. In stating that one is "self-employed," they work around the true nature of their precarious situation and avoid the shame that comes with unemployment. Even though content creation may not be a viable source of income, it acts as a placeholder until one finds a more solid source of income.

An additional peer pressure driving Gen Z towards content creation is the network effect. Essentially, the network effect is where the value of a product is derived from how large the market is.[22] In the case of social media, due to the large user base and audience potential, its inherent "value" increases for content creators, and as a result, more and more people would want to pursue content creation. Moreover, this cycle compounds on top of itself – more content means more users, more users means more creators, more creators means more content. Therefore, as the digital environment expands, its importance among entrepreneurs in terms of exposure increases to the point where Gen-Z peers feel that content creation is necessary to be seen and heard.[23]

ii) Pull Factors

On the other hand, content creation also appeals to Gen Z through the promise of flexibility. As discussed in earlier sections, content creation allows for flexible work hours, deviating from a strict 9-5 work schedule. Many of Gen Z desire to leave jobs within 2 years mainly due to some form of burnout or dissatisfaction.[14] In addition, the ability to create whatever content on whatever topic presents a higher level of attractiveness. Based on a 2026 Deloitte survey of Gen Z, about 96% believed that the purpose of their job is important to job satisfaction and well being.[24] Due to the shrinking number of available jobs, finding and

acquiring a job that fits this “purpose” becomes increasingly more difficult. However, content creation lets people record anywhere, at any time, and set their own schedules, helping them avoid burnout by taking breaks when needed. Thus, the freedom of creation may seem more enticing to some, as it allows those individuals to create for a purpose, further contributing to their job satisfaction. Due to the newfound freedom of this kind of job, Gen Z would perceive content creation as something much more enjoyable even though the profits and sustainability are uncertain.

Aside from flexibility, a study on the likelihood of students becoming entrepreneurs indicates another key factor. Out of all the students interviewed, there was an overarching consensus: exposure is the main reason they would become an entrepreneur.[23] The recognition aspect of entrepreneurship in this case can also be applied to digital entrepreneurship. Therefore, Gen Z could become digital entrepreneurs (content creators) for not just financial stability and flexibility, but also fame and recognition.

The "platform native" characteristic of Gen Z also plays a pivotal role in their preference for this kind of work.[25] Because they grew up using these specific applications to socialize, message friends, and upload personal updates, transitioning to broader content creation feels like one effortless step away. Platform algorithms systematically reinforce this perceived viability by relentlessly pushing successful content to users, constructing a glamorous ideal for the next generation.

Furthermore, the idea of quick success draws even more of Gen Z into content creation. Low barriers to entry allow for easy access into the industry, and the success pushed to individuals by the algorithm makes people believe that it is possible. Evidence from the 2020 pandemic include individuals like Khaby Lame, who amassed over 30 million followers over 2 months, and Bella Poarch, who amassed over 15 million followers over 1 month.[26,27] The belief that overnight fame could happen in content creation draws more and more of Gen Z, who highly value exposure, into the industry.

Finally, the inherently unpredictable nature of content creation acts as a powerful magnet rather than a deterrent. Because Gen Z is composed of platform natives who consume social media in real time, they witness this unpredictability firsthand. In the examples above, the creators achieved viral success overnight simply by executing a signature dance on TikTok or by using basic common sense to ridicule absurd "life hacks." To an outside observer, this rapid ascent appears almost entirely random and decoupled from traditional skill sets. However, this lottery-based system fosters a powerful psychological incentive: if algorithmic success is perceived as a game of chance rather than a strict meritocracy, then anyone – regardless of their socio-economic background or traditional qualifications – believes they have a genuine shot at hitting the jackpot.

6. Conclusion

This paper has examined the paradoxical persistence of content creation as a career choice among Generation Z despite its fundamentally precarious nature. Content creation

represents a distinct form of digital entrepreneurship characterized by low barriers to entry, flexibility, and potential rewards—yet these advantages are accompanied by significant risks: algorithmic dependency, income instability, absence of employment protections, and fierce competition where only the top 1% achieve sustainable success. Practices such as shadowbanning can quietly suppress content without notification or recourse, leaving creators in a state of uncertainty where livelihoods can be undermined without explanation. The constant need to adapt to evolving trends leads to burnout, with over 79% of creators experiencing exhaustion—a statistic underscoring the unsustainability of this labor model.

Despite these risks, content creation remains attractive to Gen Z for three interconnected reasons. Push factors—declining traditional job availability and rising costs—compel many toward content creation as necessity entrepreneurship. Pull factors—flexibility, purpose-driven work, low barriers, and the promise of overnight success—make it appear accessible and lucrative. Cultural factors—network effects, peer normalization, stigma avoidance, and Gen Z's platform-native identity—further entrench content creation as a logical career path.

Ultimately, there are many factors that could outweigh precarity in content creation from Gen Z's perspective. However, without regulatory interventions, platform accountability, or collective protections, the creator economy risks perpetuating exploitation masked as empowerment. Future research should explore policy interventions addressing creator precarity, including platform transparency, portable benefits, and collective bargaining mechanisms. As the creator economy continues to expand, understanding these tensions becomes increasingly critical for shaping a more equitable digital labor landscape.

In this paper, many of the claims made were mostly speculative, and did not have concrete evidence to back them up. The claims address the potential push/pull factors in the rise of content creation but do not address how much each push/pull factor affects this rise. Furthermore, all information was gathered through open access sources, so relevant information could be hidden behind paywalls and other barriers.

For further research, there are many pathways to choose from. Studies on how much precarity affects Gen Z, or Gen Z's tolerance for precarity and risk can be done, as well as in-depth analysis on precarity in the context of content creation. In addition, a deeper analysis into the economic factors (monetary incentives) of content creation vs. traditional jobs could be done – digging deeper into the demographic specifics of content creators (education, socioeconomic standing before content creation, region, etc.) and how that could affect their decision between a traditional job and content creation. Finally, more in-depth exploration of the algorithms and monetization of content can be done, helping us better understand how the websites themselves affect the precarity of this occupation.

AI Acknowledgement

AI was used for summarizing meeting transcripts, as well as some parts of brainstorming.

Works Cited

(1)Kornfeld, L.; Deighton, J. MEASURING THE DIGITAL ECONOMY. **2025**, 178.

(2)Wollborn, P.; Dornekott, D.; Holder, U. Entrepreneurial Efforts and Opportunity Costs: Evidence from Twitch Streamers. Int. Entrep. Manag. J. **2023**, 19 (3), 1209–1238. https://doi.org/10.1007/s11365-023-00849-2

(3)Edeling, A.; Wies, S. Embracing Entrepreneurship in the Creator Economy: The Rise of Creatrepreneurs. Int. J. Res. Mark. **2024**, 41 (3), 436–454. https://doi.org/10.1016/j.ijresmar.2024.07.003

(4)Duffy, B. E.; Pinch, A.; Sannon, S.; Sawey, M. The Nested Precarities of Creative Labor on Social Media. Soc. Media Soc. **2021**, 7 (2). https://doi.org/10.1177/20563051211021368

(5)content creator, n. meanings, etymology and more | Oxford English Dictionary. https://www.oed.com/dictionary/content-creator_n (accessed 2026-03-15)

(6)Ali, O. A. A. How Gen Z and E-Commerce SMEs Are Driving U.S. Economic Growth.

(7)Twenge, J. M. Generations: The Real Differences between Gen Z, Millennials, Gen X, Boomers, and Silents--and What They Mean for America's Future; Atria Books: New York (N.Y.), 2023.

(8)He, A.; Zink, N.; Moquin, E. How Brands Can Succeed at Influencer Marketing; Analyst Report; Morning Consult, 2023; p 32. https://pro.morningconsult.com/analyst-reports/influencer-marketing-trends-report?utm_source=mc_comms_outreach&utm_medium=mc_comms_referral&utm_campaign=influencer_report (accessed 2026-04-05)

(9)Vallas, S.; Schor, J. B. What Do Platforms Do? Understanding the Gig Economy. Annu. Rev. Sociol. **2020**, 46 (1), 273–294. https://doi.org/10.1146/annurev-soc-121919-054857

(10)MSc in Digital Marketing and Media, Yeshiva University, Katz School of Science and Health, NY, USA; Mou, A. J.; Faruk, M. I.; Master's in Information Technology Project Management, Saint Francis College, NY, USA. Content Creation as a Growth Strategy: Evaluating the Economic Impact of Freelance Digital Branding. Am. J. Sch. Res. Innov. **2024**, 03 (02), 28–51. https://doi.org/10.63125/mj667y36

(11)Qiu, J.; Chen, Y.; Cohn, A.; Roth, A. E. Social Media and Job Market Success: A Field Experiment on Twitter. SSRN Electron. J. **2024**. https://doi.org/10.2139/ssrn.4778120

(12)Mason, A. N.; Narcum, J.; Mason, K. Social Media Marketing Gains Importance after Covid-19. Cogent Bus. Manag. **2021**, 8 (1), 1870797. https://doi.org/10.1080/23311975.2020.1870797

(13)Fossen, F. M. Self-Employment over the Business Cycle in the USA: A Decomposition. Small Bus. Econ. **2021**, 57 (4), 1837–1855. https://doi.org/10.1007/s11187-020-00375-3

(14Waworuntu, E. C.; Kainde, S. J. R.; Mandagi, D. W. Work-Life Balance, Job Satisfaction and Performance Among Millennial and Gen Z Employees: A Systematic Review. Society **2022**, 10 (2), 384–398. https://doi.org/10.33019/society.v10i2.464

(15)F. Songong, E.; Goris, T. Pilot Study of Gen Z Candidates For Successful Employment in the U.S. Job Market. In 2025 ASEE Annual Conference & Exposition Proceedings; ASEE

Conferences: Montreal, Quebec, Canada, 2025; p 57045. https://doi.org/10.18260/1-2--57045
(16)Bierbrier, C. V. Work Value Preferences of Generation Z in the United States.
(17)YikesMC. Is the Worst Era of YouTube Finally Over?; 2026. https://www.youtube.com/watch?v=B8Dw2c0YHLU (accessed 2026-08-16)
(18)Delmonaco, D.; Mayworm, S.; Thach, H.; Guberman, J.; Augusta, A.; Haimson, O. L. "What Are You Doing, TikTok?" : How Marginalized Social Media Users Perceive, Theorize, and "Prove" Shadowbanning. Proc. ACM Hum.-Comput. Interact. **2024**, 8 (CSCW1), 1–39. https://doi.org/10.1145/3637431
(19)Job Openings and Labor Turnover Summary - 2026 M04 Results. **2026**.
(20)Infra-Annual Registered Unemployment and Job Vacancies: Total Economy: Unfilled Vacancies for the United States. https://fred.stlouisfed.org/series/LMJVTTUVUSM647S (accessed 2026-06-07)
(21)Unemployment Level. https://fred.stlouisfed.org/series/UNEMPLOY (accessed 2026-06-07)
(22)What Are Network Effects? | HBS Online. Harvard Business School. https://online.hbs.edu/blog/post/what-are-network-effects (accessed 2026-06-07)
(23)Friedman, B. A.; Aziz, N.; Keles, I.; Sayfullin, S. Predictors of Students` Desire to Be an Entrepreneur: Kyrgyzstan, Georgia, and the United States.
(24)2026 Gen Z and Millennial Survey. **2026**.
(25)Prensky, M. Digital Natives, Digital Immigrants. Horiz. **2001**, 9.
(26)Khabane lame's TikTok Statistics. Social Blade. https://socialblade.com/tiktok/user/khaby.lame (accessed 2026-06-07)
(27)Bella Poarch's TikTok Statistics. Social Blade. https://socialblade.com/tiktok/user/bellapoarch (accessed 2026-06-07)

Authors

Ethan Gu is a high school student from California, and is interested in studying economics in college. In his free time, Ethan enjoys listening to music, and playing volleyball with friend

Tooth Jewels and Enamel Health: Cosmetic Decoration or a Microbiofilm Trap?

By Kareena Bakhshi

Abstract

Tooth jewels, also called tooth gems, are small decorative crystals, or ornaments bonded to the surface of teeth. Recently, these jewels have grown in popularity due to social media, but its implications on enamel health remain understated. They have proven to promote plaque retention, thus driving demineralization of the tooth. The adhesives utilized in the bonding of tooth jewels could cause decreased enamel health. Namely, do-it-yourself (DIY) bonding agents can lead to a larger adhesive margin. Various literature suggests that jewels are not harmful, rather the risks arise from enamel alterations which promote plaque-retentive sites, biofilm-formation, and enamel demineralization. These complications can be amplified by factors like salivary function, adhesive choice, and oral hygiene. Newer evidence shows that preventive treatment, like fluoride, remineralizing agents, and professional application, could reduce risks and preserve enamel health. Because clinical studies are limited, this review draws on evidence from orthodontic brackets, due to their similarity in application to tooth jewels. The evidence suggests that brackets lead to raised or rough surfaces that increase plaque retention, and the risk of white spot lesions with inadequate oral hygiene, similar to tooth jewels. Thus, tooth jewels should not be treated as harmless; they should be considered cosmetic dental attachments that require professional placement, cautious hygiene, and safe removal.

Introduction

Tooth jewels fall into the practice of cosmetic dentistry in which a diamond or stone is attached to a tooth using the etch-and-bond process, similar to orthodontic brackets.[1] In which, 35%-37% phosphoric acid is applied to the tooth, and dissolves a layer of hydroxyapatite, which increases the surface roughness, thus allowing the jewel to bond to the tooth better.[2-3]

Tooth enamel is the outer protective tissue which protects teeth from damage and is highly durable. However, it is easily broken down through specific foods, acidic diets, and the adhesives utilized in the etch-and-bond process.[4] Because tooth jewels are directly attached to the surface of the tooth, it promotes the degeneration of the enamel.[1]

Dental enamel protects the tooth against mechanical, thermal, and chemical stress, however; unlike skin, it is unable to regenerate automatically after removal. This means that any procedure involving dental enamel modification warrants extreme caution and clinical consideration.

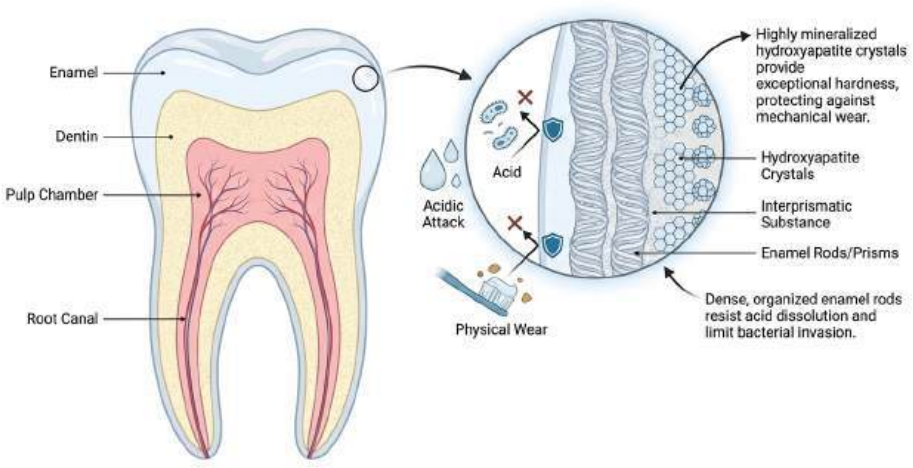

Fig 1: Enamel strength demonstrated through various components present within (Figure made with Figure labs).

The etch-and-bond process creates porosites on the enamel surface which leads to decalcification and a potential risk of enamel fracture.[3] This means that tooth jewels have the ability to alter the enamel's structure.[1-2] Once this structure has been altered, it allows for the formation of plaque retentive regions on the tooth, which can lead to plaque build-up and bacterial clusters as the user is unable to maintain proper dental hygiene.

This review examines how tooth jewels create a localized plaque-retentive region where bacteria can accumulate. The interface at which this occurs includes the tooth surface, the adhesive margin, the jewel, saliva, and microbiofilms. The most important risk behind tooth jewels may not be the center of the jewel, rather the margin around the jewel where the adhesive meets the enamel.

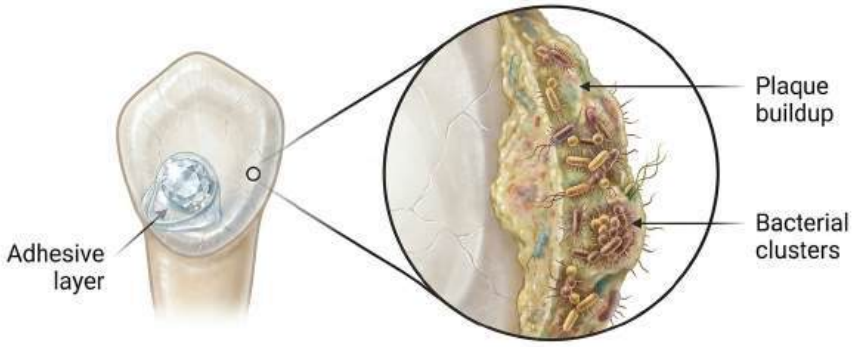

Fig 2: Bacterial and plaque build up through the installation of a tooth gem (Figure made with Figure labs).

Tooth Anatomy Relevant to Tooth Gems

The anatomy of a tooth is a huge factor in the adhesions of the jewel through the enamel. The enamel is the hardest tissue in the tooth and it allows teeth to process food.[4] Another component of the tooth is dentine, which is found underneath the enamel.[4] These components work hand-in-hand as enamel damage exposes dentine.[5] Next, the pulp, where the nerves and blood vessels can be found, are heavily affected by dental caries, which are caused when the enamel causes a breakage. [6-7]

Once these caries have formed, gingiva, the gum tissue, is oftentimes irritated by bacteria being trapped around tooth jewels, which results in bleeding, and in extreme cases increased tooth movement due to lack of supportive structures.[8] Another component of the mouth is saliva, with the ability to act as a buffer to mineralize and lubricate the mouth. However, when a tooth jewel is attached, the enamel loses its required minerals. To remineralize the tooth, saliva is used as it carries various ions that promote remineralization. Namely phosphate, calcium, and fluoride.[9] Lastly, the pellicle, where bacterial adhesion initially occurs, acts as a substrate for the formation of dental biofilms, but also covers all surface structures in the oral cavity. [10]

Fig 3: The anatomy of a tooth, showing both outer and inner structures (Figure made with Figure labs).

Tooth Jewels as a Mini Orthodontic Bracket

Tooth jewels are cosmetic dental attachments that share some similar bonding procedures with orthodontic brackets, thus creating adhesive margins that could influence plaque retention. Orthodontic appliances like these promote the formation of biofilms as they offer more binding sites for bacteria and impede oral hygiene.[19] Consequently, this disrupts microbial homeostasis, which leads to enhanced biofilm formation, thus leading to increased amounts of cariogenic bacteria in biofilms.[19] These disruptions allow the oral cavity to favor acidic bacteria such as *Streptococcus* species, all while decreasing the growth of health-associated bacteria, such as *Streptococcus gordonii* and *Actinomyces naeslundii*, which maintain homeostasis in the oral

cavity.[19-20] This change in the pH leads to the formation of white spot lesions, and thus leads to early demineralization.[16-17]

Fixed orthodontic appliances can create plaque-retentive regions. Due to the complex nature of these treatments, the enamel often creates plaque-retentive regions which promote bacterial formation, thus increasing the risk of enamel demineralisation. [19,21]

There are various techniques of bonding brackets along the tooth, some of them result in greater enamel demineralisation. As demonstrated by Raza et al., conventional methods of bonding led to greater demineralisation, as compared to the transillumination method. This method is utilised via guiding the UV light through the palatal side of the tooth, and passes through the enamel and dentine to cure the adhesive. The increased formation of biofilms could disrupt oral homeostasis by decreasing plaque pH, thus leading to gingival inflammation. [22]

Though both tooth jewels and orthodontic brackets, have increased plaque retention risks, there are various differences between these two types of dental attachments.[19,21] Firstly, orthodontic brackets are used for the purpose of moving teeth around, as compared to tooth gems whose sole purpose is aesthetics.[1] Alongside, since orthodontic brackets could cover a larger surface area than tooth gems, there is an increased risk of these brackets creating larger plaque-retentive areas, making the infectious region less localised.[18-19] Secondly, orthodontic treatments typically involve professional monitoring, whereas tooth jewel follow-ups may vary. Lastly, if tooth jewels are removed by the wearer it could increase enamel damage. [2,21-22]

These similarities make orthodontic research useful as a model to identify the mechanisms of enamel damage around tooth jewels. However, the differences in number, location, duration of use, attachment size, and maintenance limit a direct comparison between the two. Consequently, orthodontic findings cannot establish the frequency or severity of complications in those who use tooth jewels without proper clinical evidence.

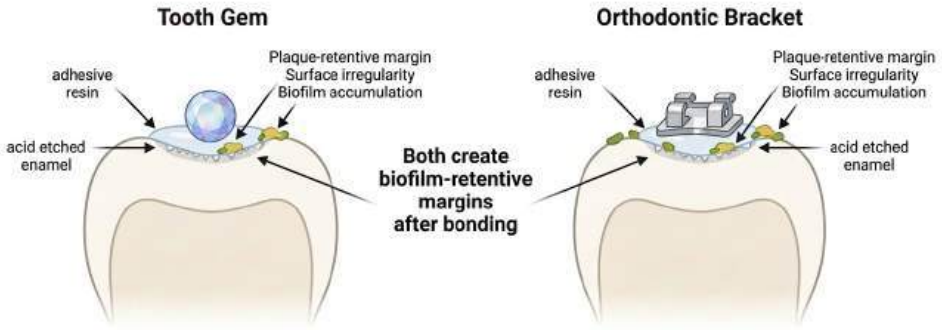

Fig 4: Side-by-side comparison of tooth gem vs. orthodontic bracket (Figure made with Figure labs).

Adhesive Comparison

There are various types of adhesives that can be used to apply tooth gems onto the tooth (Table 1). These adhesives are most commonly known as orthodontic adhesives and they differ

in both chemical composition and surface features. Some aspects of these adhesives include surface roughness (SR) and surface free energy (SFE) which determine levels of bacterial colonisation, plaque retention rates, and finally enamel demineralisation.[27]

Adhesive Used:	Surface Changes:	Biofilm Formation:	Enamel Effect:
Dental Grade	Controlled surface roughness, optimal polishing	Lower bacterial adhesion	Reduced risk of demineralisation
Excess Adhesive Used	Increased surface irregularities and exposed margins	Increased plaque retention and bacteria	Greater risk of enamel demineralisation
Excessively Etched Enamel	Microporosites, rough surface	Increased bacterial adhesion	Increased susceptibility to biofilms and demineralisation
DIY Adhesives	Variable composition and unknown properties, may lack biocompatibility	May increase plaque retention due to rough margins	Potential for enamel damage and demineralisation

Table 1: Adhesive Comparison

Though many tooth jewel users get this procedure under professionally controlled environments, some users utilize DIY dental kits. These kits are oftentimes purchased online and contain significant risks due to the lack of professional supervision. These kits use unsterilised equipment which could increase the risk of infection.[25] Alongside, these materials have not met regulatory standards, thus could lead to allergic reactions or toxicity.[25] Utilizing DIY dental kits can lead to various complications, such as accidentally etching off more enamel than expected. This exposes the tooth surface more and may even increase biofilm formation around orthodontic appliances like tooth jewels.[28] An example of this is utilising 37% phosphoric acid as it creates the most amount of microporous.[3]

The studies above examine how various bonding materials can influence enamel health. Specifically, Condò et al. researched about the properties of orthodontic adhesive resins and their effect on bacterial adhesion, while Jeon et al. found that rougher surfaces allowed for greater and more diverse bacterial attachment. Especially Jeon et al. noticed that there was fewer bacterial attachment to smoother surfaces than etched enamel, suggesting that bonding does not always increase bacterial adhesion. Jointly, these studies show that adhesive properties and surrounding enamel conditions should be considered greatly when evaluating plaque retention in the oral cavity. Alongside, this discussion was extended as Hennig et al. measured enamel demineralisation around various bracket-adhesion-systems; however, neither system was able to effectively prevent enamel demineralisation under bacterial exposure.

Gumline Placement and Periodontal Irritation

Though tooth jewels have an effect on the enamel, specific placement position of the jewel on the tooth is important in maintaining periodontal health. If the tooth jewel is placed too close to the gingival margin, it could interfere with plaque removal and increase the formation of biofilms.[32]As the plaque accumulates, bacteria trigger an inflammatory response in the gingival tissue, which results in plaque-induced gingivitis, whose symptoms are swelling and bleeding of the gums.[8] Gingivitis is generally reversible when appropriate measures are taken, however; consistent formation of biofilm may lead to severe periodontal diseases, and consequently tooth loss.[8]

The mechanisms responsible for this are similar to the ones observed around orthodontic appliances. Adhesive margins and surface irregularities increase the amount of binding sites available for bacteria, while also reducing the effectiveness of oral hygiene products.[19,27] As a result, cariogenic bacteria may form a more resistant biofilm. The location of the tooth jewel plays an important role in the determination of a clinical risk. If jewels are positioned on the mid-facial surface of the crown, they remain more accessible for toothbrushing and are less likely to interfere with plaque removal.[8,32] However, when these are placed near to the gingival margin, it creates a more localised plaque-retention area where bacteria are more likely to accumulate.[19,32] As a result, dentists should evaluate a patient's periodontal health prior to tooth jewel placement, so as to not compromise plaque removal. Additionally, Hennig et al. showed that cervical enamel demineralisation occurred 2.657 times more often than coronal demineralisation, while more severe lesions were 10.028 times more likely to occur cervically, which suggests that there is a larger risk of infection the closer the jewel is to the gingival margin. Next, educating the patient in terms of oral health and fluoride use should be presented in every tooth jewel appointment to minimize long-term complications.[12,19,26] Hence, precise acid-etching and controlled adhesive placement must be considered essential constituents of cosmetic dentistry to reduce enamel complications following this procedure.[24-25]

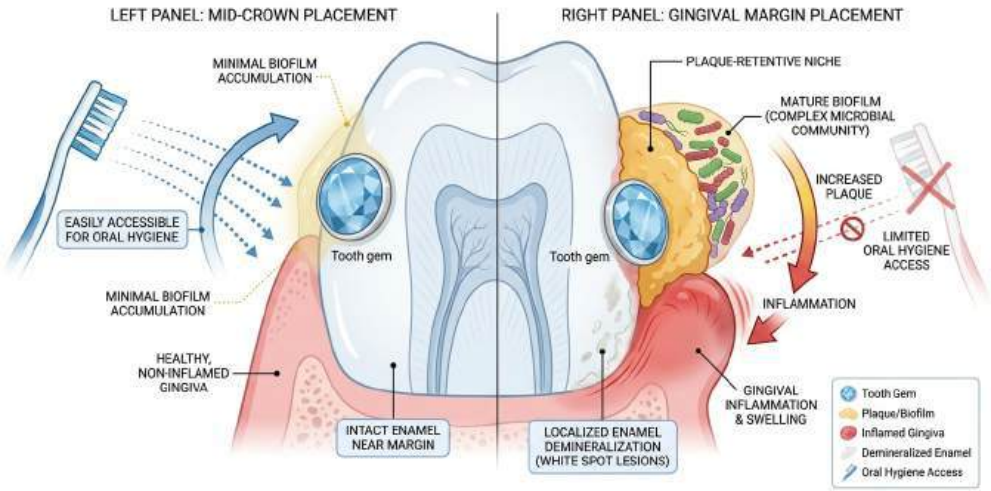

Fig 5: Different areas of tooth gem placement (Figure made with Figure labs).

Dental Plaque as a Biofilm

Through the tooth jewel, the surface shape and texture of the tooth is altered. Likewise, through the etch-and-bond process, it increases surface roughness, which could promote bacterial attachment through an increase in surface area.[10] This means that the rougher surface associated with tooth jewel attachment could promote certain bacteria such as *Streptococcus* sanguinis and *Streptococcus* mutans to express the glucosyltransferase genes gtfP, which increases the formation of biofilms, thus increasing cariogenic potential.[10,16]

However, surface roughness may not be the primary factor for bacterial adhesion. This is primarily caused by the varying measurement techniques of surface roughness which could imply that other factors could also influence biofilm formation. For example enamel chemistry and bacterial interactions could affect bacterial attachment. Additionally, these bacteria could interact with the pellicle, thus leading to the production of more biofilms.[17]

With the production of biofilm, plaque is also produced and it attaches better to rough, protected, and hard to clean surfaces, such as the area around these gems. Such complications have the ability to protect bacteria and ensure they come into contact with the enamel surface, thus increasing the potential risk of caries.[17] This leads to plaque creating both a microenvironment and functionally organised biofilms, which creates a microenvironment that promotes diseases like enamel demineralisation and various periodontal diseases through an infected gingiva.[8,18]

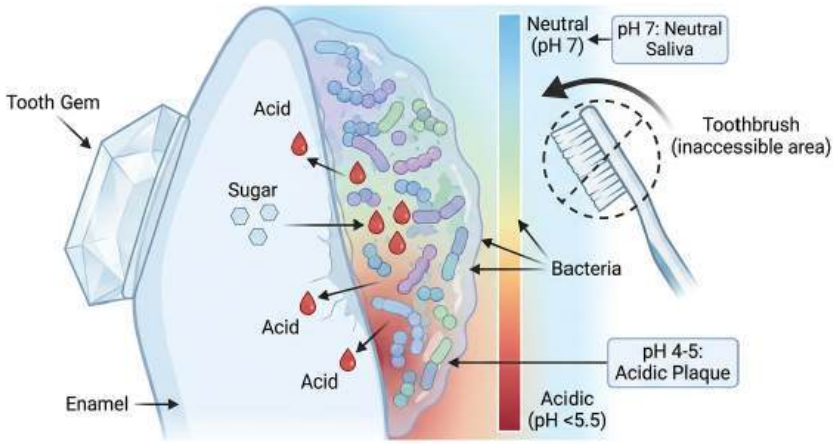

Fig 6: The Formation of Biofilm Around Tooth Gems (Figure made with Figure labs).

Tooth gems create an adhesive margin to previously smooth surfaces. Once this gem is bonded, it creates plaque-retentive zones that can remain even after brushing teeth; the smoother the enamel the more efficiently it can be cleaned. This occurs because bacteria can bond to surfaces which contain both greater surface area and surface roughness which protects them from shear forces generated by brushing of the teeth.[17]

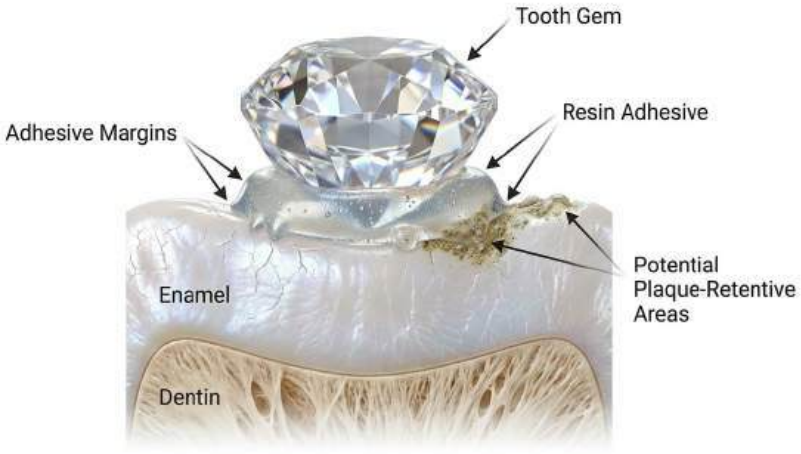

Fig 7: Adhesive Margins Created Around Tooth Via Tooth Gems (Figure made with Figure labs).

Enamel Chemistry, Hydroxyapatite, pH, and Minerals

One of the concerns following the application of tooth jewels is the risk of dental caries as enamel chemistry, mouth pH, hydroxyapatite, and tooth minerals all get altered. This process is known as demineralization which occurs when mineral ions are removed from the hydroxyapatite in the chemical reaction.[11] The formula for demineralization occurs as follows: $Ca_{10}(PO_4)_6(OH)_2 + 8H^+ \rightarrow 10Ca^{2+} + 6HPO_4^{2-} + 2H_2O$, once the pH rises the process of remineralization occurs as minerals are redeposited onto the hydroxyapatite.[11]

Demineralization primarily occurs when the oral environment becomes extremely acidic which decreases the overall pH, which may cause calcium and phosphate to leave the oral environment, thus promoting early demineralization. As this cycle continues on, this leads to dental caries.[12]

However, this process is not entirely irreversible as fluoridated hydroxyapatite (FHAp) can enhance hydroxyapatite stability, which can resist/prevent dental caries and improve remineralization.[13] Alongside, saliva also has the ability to inhibit demineralization and promote remineralisation, as it provides the tooth with the required minerals. [12]

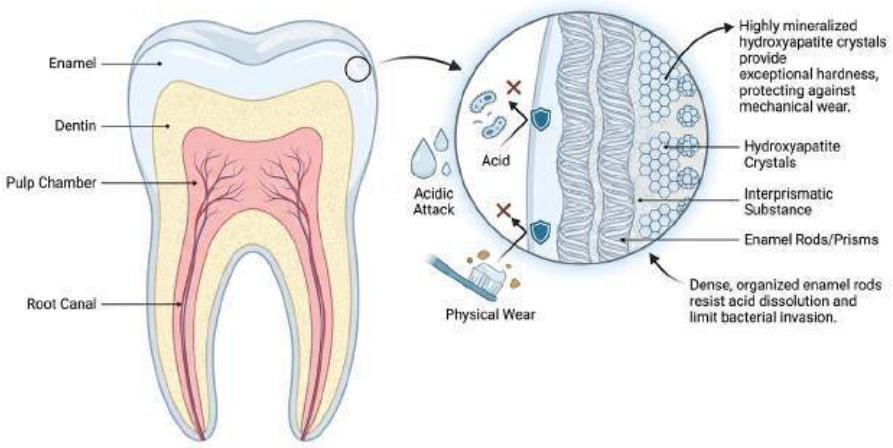

Fig 8: The process of enamel demineralisation (Figure made with Figure labs).

Dental caries is a disease that is driven by the consumption of sugar and results in the cycle of demineralization and remineralization of dental tissues.[14] As oral bacteria metabolizes sugar and produces acids, maintaining the pH lower than 5.5, the cycle continues.[15] To remineralize the enamel, it must first be made more resistant to acids, by substituting hydroxyl groups (OH-) in the hydroxyapatite lattice, suggesting that FHAp may be a good material for synthetic enamel.[13]

White Spot Lesions (Early Visible Enamel Damage)

One of the first visible signs of enamel damage are known as white spot lesions. On the enamel, these appear as white and opaque.[23] These lesions are often reversible when these lesions are first spotted and if fluoride is present in the environment.[14,26] This process is known as remineralization when there is increased growth of hydroxyapatite crystals, and with the presence of fluoride forms FHAp.[26] However, with tooth gems, as a portion of the tooth is covered, it could hide early changes in the enamel as the area could be near the adhesive margin. Furthermore, Julien et al. found that white spot lesions increased from 23.4% to 68.4% following treatment, which suggests that bonding can lead to early demineralisation. Alongside, the type of adhesive and materials utilized in the application of tooth gems can also play a huge role in the formation of white spot lesions. For example, acid-etching strips off the outer enamel, creating microscopic porosites which allow for the resin to bond with the tooth, however; if this procedure is conducted incorrectly it could increase surface area and roughness, which allows for more bacteria to enter the oral cavity.[13]

Fig 9: Side-by-side diagram of professionally placed vs. DIY tooth gem kit (Figure made with Figure labs).

Diet, Saliva, and Fluoride

There are various factors that can determine the dental risks in individuals associated with tooth jewels. Some of these factors include diet, salivary function, fluoride exposure, and oral hygiene can determine how effectively the oral cavity is able to counteract the increased amounts of plaque around the interface.[29] For example, increased amounts of carbohydrates in diet can promote the production of acid through cariogenic bacteria, such as *Streptococcus* and Lactobacillus species, where the saliva acts like as a natural buffer, by neutralizing the acids, and supplying calcium and phosphate ions, allowing for food debris and microorganisms to be cleared.[12,14,29] However, individuals with decreased saliva flow or a lower buffering capacity might be unable to make up for the acidic conditions that can develop near the binding site of the tooth jewel. Furthermore, the components of the salivary microbiota could influence an individual's susceptibility to dental caries, which means that patients with more cariogenic bacteria could experience greater risks following the jewel placement.[20]

Remineralising agents play an important role in the limitation of enamel damage. Following acid etching and adhesive placement, the enamel undergoes cycles of demineralisation and remineralisation. Fluoride is considered one of the most powerful types of preventative therapy as it promotes the formation of FHAp, which increases enamel resistance to acids.[13,26] Similarly, casein phosphopeptides-amorphous calcium phosphate (CPP-ACP) stabilizes the calcium and phosphate ions present in the oral cavity to redeposit them into enamel lesions. As a result, CPP-ACP has the ability to inhibit enamel demineralisation and promote remineralisation.[27-28] Though tooth jewels can increase localised plaque retention, using the appropriate amount of fluoride, remineralising agents, and maintaining good oral hygiene may reduce the chances of enamel changes progressing to serious periodontal diseases.

Removal

The final stage of a tooth jewel is the process of its removal, which is an extremely important process as it can have consequences on the preservation of the health of the enamel. Even if the preceding steps have been conducted correctly, inappropriate removal of the jewel can result in enamel loss, increased surface area roughness, thus leading to plaque accumulation, which may contribute to enamel demineralisation. As a result, the jewel removal should be conducted by a dental professional using proper debonding agents to minimize enamel damage. There are various debonding agents used in this process, most commonly resin-removal burs are used to detach the jewel. After this the remaining adhesive is removed using finishing burs and polished to increase enamel smoothness. Therefore, the smoothness of the enamel significantly differs depending on the method of removal used.[24]

Following jewel removal, the tooth should be thoroughly examined for residual adhesive, white spot lesions, discoloration, and enamel damage. In areas where white spot lesions are found, fluoride or other remineralising agents containing casein phosphopeptides can be prescribed to encourage remineralisation and increase enamel strength.[12-13,26] The current literature suggests that professional debonding procedures minimise irreversible enamel damage.

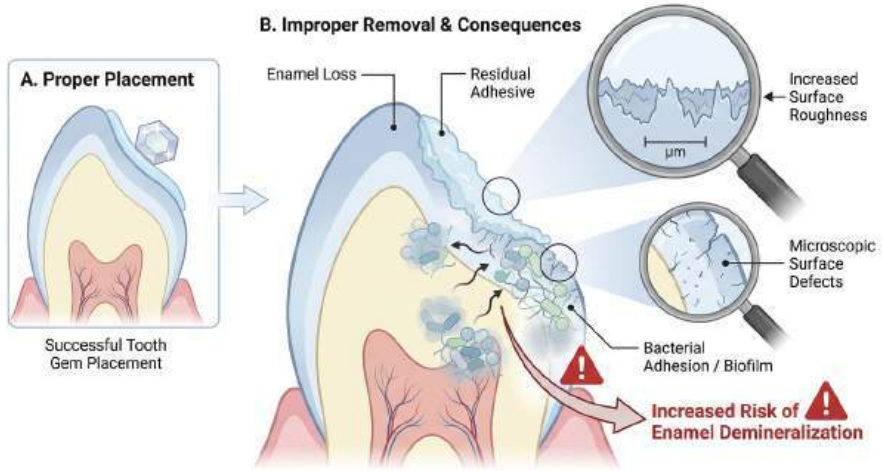

Fig 10: Proper vs. Improper Removal of Tooth Gem (Figure made with Figure labs).

Conclusion

This review examined the relationship between tooth jewel application and enamel health, focusing on gingival health, biofilm formation, and preventative strategies. Collectively, the literature suggests that tooth jewels can alter enamel surface topography through acid etching. These changes can form plaque retentive areas around the teeth, facilitating the formation of biofilm and plaque. While numerous studies have linked excessive acid etching to increased biofilm formation, the usage of orthodontic brackets as a model for understanding this process proves to be useful. However, an important difference is that these brackets are followed up by regular professional checks, whereas follow ups in tooth jewels remain vague in current literature.

The evidence reviewed indicates that tooth jewels may be associated with increased plaque retentive areas. However, significant variability exists across studies due to differences in adhesives, gingival health, participant oral hygiene, and jewel placement, making them difficult to cross compare. These factors can influence the formation of plaque, biofilms, and the process of demineralisation and remineralisation. As a consequence, the evidence in current literature does not support a singular risk profile for patients, showing that clinical outcomes may depend on the patient and the method of treatment. This inconsistency highlights the need for individualized risk assessments, proper material selection, proper clinical techniques, and regular monitoring.

Despite the growing body of research, several important questions remain unanswered. Future studies should focus on location of jewel placement, types of adhesives, and removal techniques. Such studies could address current gaps in research and influence clinical guidelines for tooth jewel application and bring to light new perspectives to this growing debate. Greater attention should also be given to protective factors that may reduce the risk of dental caries amongst all tooth jewel users.

Overall, current evidence suggests that tooth jewels may increase plaque-retentive areas when its placement, removal, and care is inadequate. However, these risks differ across patients due to various factors such as health, and further research and studies are required to determine its impact on oral health. Rather than viewing tooth jewels and orthodontic appliances as universally harmful or beneficial, future research should focus on understanding the specific conditions where their usage minimises enamel health while increasing the risks of dental caries.

Works Cited

(1) Senneby, A., et al. “Acid Tolerance Properties of Dental Biofilms in Vivo.” BMC Microbiology, vol. 17, no. 1, 2017, article 165. https://doi.org/10.1186/s12866-017-1074-7.

(2) Brizuela, M., and B. Bordoni. “Anatomy, Head and Neck, Dental Pulp.” StatPearls, StatPearls Publishing, 2026.

(3) Sterzenbach, T., et al. “Bioadhesion in the Oral Cavity and Approaches for Biofilm Management by Surface Modifications.” Clinical Oral Investigations, vol. 24, no. 12, 2020, pp. 4237–60. https://doi.org/10.1007/s00784-020-03646-1.

(4) Chang, J., et al. “Biomimetic Enamel-like Crystals: A Versatile Platform for Unraveling the Basic Mechanisms of Demineralization and Remineralization.” ACS Applied Materials & Interfaces, vol. 17, no. 36, 2025, pp. 50505–18. https://doi.org/10.1021/acsami.5c13544.

(5) Hong, X., et al. “Clinical Effects of Different Etching Modes for Universal Adhesives: A Systematic Review and Meta-Analysis.” Annals of Palliative Medicine, vol. 10, no. 5, 2021, pp. 5462–73. https://doi.org/10.21037/apm-21-890.

(6) Raza, M., et al. “Comparison of Demineralization around Orthodontic Brackets Cured by Conventional Method and Transillumination Technique—An In Vitro Evaluation.” Journal of Orthodontic Science, vol. 9, no. 1, 2020, article 16. https://doi.org/10.4103/jos.JOS_3_20.

(7) Anil, A., et al. “Demineralization and Remineralization Dynamics and Dental Caries.” Dentistry, edited by L.-C. Rusu and L. Cosmina Ardelean, vol. 12, IntechOpen, 2022. https://doi.org/10.5772/intechopen.105847.

(8) Pitts, N. B., et al. “Dental Caries.” Nature Reviews Disease Primers, vol. 3, no. 1, 2017, article 17030. https://doi.org/10.1038/nrdp.2017.30.

(9) Rathee, M., and A. Sapra. “Dental Caries.” StatPearls, StatPearls Publishing, 2026.

(10) Featherstone, J. “Dental Caries: A Dynamic Disease Process.” Australian Dental Journal, vol. 53, no. 3, 2008, pp. 286–91. https://doi.org/10.1111/j.1834-7819.2008.00064.x.

(11) Lacruz, R. S., et al. “Dental Enamel Formation and Implications for Oral Health and Disease.” Physiological Reviews, vol. 97, no. 3, 2017, pp. 939–93. https://doi.org/10.1152/physrev.00030.2016.

(12) Marsh, P. D. “Dental Plaque as a Biofilm and a Microbial Community—Implications for Health and Disease.” BMC Oral Health, vol. 6, suppl. 1, 2006, article S14. https://doi.org/10.1186/1472-6831-6-S1-S14.

(13) Dionysopoulos, D., et al. “Dentin Hypersensitivity: Etiology, Diagnosis and Contemporary Therapeutic Approaches—A Review in Literature.” Applied Sciences, vol. 13, no. 21, 2023, article 11632. https://doi.org/10.3390/app132111632.

(14) Veerabhadrappa, S. K., et al. “DIY Tooth Filling Kits.” British Dental Journal, vol. 239, no. 10, 2025, p. 668. https://doi.org/10.1038/s41415-025-9385-9.

(15) Teutle-Coyotecatl, B., et al. “Effect of Surface Roughness of Deciduous and Permanent Tooth Enamel on Bacterial Adhesion.” Microorganisms, vol. 10, no. 9, 2022, article 1701. https://doi.org/10.3390/microorganisms10091701.
(16) Jani, B., et al. “Effectiveness of Different Etching Agents on Enamel Surface and Shear Bond Strength: An In Vitro Evaluation.” Cureus, 2024. https://doi.org/10.7759/cureus.54008.
(17) Ghaleb, L., et al. “Evaluation of Enamel Surface Integrity after Orthodontic Bracket Debonding: Comparison of Three Different System.” BMC Oral Health, vol. 24, no. 1, 2024, article 358. https://doi.org/10.1186/s12903-024-04138-4.
(18) Condò, R., et al. “In Vitro Evaluation of Structural Factors Favouring Bacterial Adhesion on Orthodontic Adhesive Resins.” Materials, vol. 14, no. 10, 2021, article 2485. https://doi.org/10.3390/ma14102485.
(19) Hennig, C.-L., et al. “Influence of Different Bracket Adhesive Systems on Enamel Demineralization—An In Vitro Study.” Journal of Clinical Medicine, vol. 12, no. 13, 2023, article 4494. https://doi.org/10.3390/jcm12134494.
(20) An, J.-S., et al. “Managing Oral Biofilms to Avoid Enamel Demineralization during Fixed Orthodontic Treatment.” Korean Journal of Orthodontics, vol. 53, no. 6, 2023, pp. 345–57. https://doi.org/10.4041/kjod23.184.
(21) Jeon, D.-M., et al. “Orthodontic Bonding Procedures Significantly Influence Biofilm Composition.” Progress in Orthodontics, vol. 21, no. 1, 2020, article 14. https://doi.org/10.1186/s40510-020-00314-8.
(22) Kinane, D. F. “Periodontal Diseases.” Nature Reviews Disease Primers, vol. 3, no. 1, 2017, article 17039. https://doi.org/10.1038/nrdp.2017.39.
(23) Trombelli, L., et al. “Plaque-Induced Gingivitis: Case Definition and Diagnostic Considerations.” Journal of Clinical Periodontology, vol. 45, suppl. 20, 2018. https://doi.org/10.1111/jcpe.12939.
(24) Julien, K. C., et al. “Prevalence of White Spot Lesion Formation during Orthodontic Treatment.” The Angle Orthodontist, vol. 83, no. 4, 2013, pp. 641–47. https://doi.org/10.2319/071712-584.1.
(25) Reynolds, E. C. “Remineralization of Enamel Subsurface Lesions by Casein Phosphopeptide-Stabilized Calcium Phosphate Solutions.” Journal of Dental Research, vol. 76, no. 9, 1997, pp. 1587–95. https://doi.org/10.1177/00220345970760091101.
(26) Amaechi, B. T. “Remineralization Therapies for Initial Caries Lesions.” Current Oral Health Reports, vol. 2, no. 2, 2015, pp. 95–101. https://doi.org/10.1007/s40496-015-0048-9.
(27) Heller, D., et al. “Saliva and Serum Protein Exchange at the Tooth Enamel Surface.” Journal of Dental Research, vol. 96, no. 4, 2017, pp. 437–43. https://doi.org/10.1177/0022034516680771.
(28) Belstrøm, D., et al. “Salivary Microbiota in Individuals with Different Levels of Caries Experience.” Journal of Oral Microbiology, vol. 9, no. 1, 2017, article 1270614. https://doi.org/10.1080/20002297.2016.1270614.

(29) Yoshida, Y., et al. “The Influence of a Glucosyltransferase, Encoded by gtfP, on Biofilm Formation by Streptococcus sanguinis in a Dual-Species Model.” APMIS, vol. 122, no. 10, 2014, pp. 951–60. https://doi.org/10.1111/apm.12238.
(30) Huq, N., et al. “The Interactions of CPP–ACP with Saliva.” International Journal of Molecular Sciences, vol. 17, no. 6, 2016, article 915. https://doi.org/10.3390/ijms17060915.
(31) Velliyagounder, K., et al. “The Microflora Diversity and Profiles in Dental Plaque Biofilms on Brackets and Tooth Surfaces of Orthodontic Patients.” Journal of Indian Orthodontic Society, vol. 53, no. 3, 2019, pp. 183–88. https://doi.org/10.1177/0301574219851160.
(32) Farooq, I., and A. Bugshan. “The Role of Salivary Contents and Modern Technologies in the Remineralization of Dental Enamel: A Narrative Review.” F1000Research, vol. 9, 2021, article 171. https://doi.org/10.12688/f1000research.22499.3.
(33) Bhatia, S. “Tooth Jewellery—Its Knowledge and Practice among Dentists in Tricity, India.” Journal of Clinical and Diagnostic Research, 2016. https://doi.org/10.7860/JCDR/2016/18578.7441.
(34) Ng, T. C.-H., et al. “Updates on Caries Risk Assessment—A Literature Review.” Dentistry Journal, vol. 12, no. 10, 2024, article 312. https://doi.org/10.3390/dj12100312.

Human Chimerism: Advances in Genetic Testing for Diagnosis

By Maria Victoria Cardelli De All

Abstract

The presence of two or more genetically distinct cell populations in a single individual is known as human chimerism. This condition can affect transplants, forensic investigations, genetic diagnosis, and precision medicine. In this review, we will assess how advances in genetic testing improve the diagnosis of human chimerism, by surveying current publications on standard and developing diagnosis methods. Techniques like short tandem repeat (STR), polymerase chain reaction (PCR), fluorescence in situ hybridization (FISH), as well as Next Generation Sequencing (NGS) are compared based on their limitations, clinical applications, and detection thresholds. The studies indicate that qPCR, digital PCR, and targeted NGS can improve low-level detection in transplant monitoring, but tissue selection remains essential in suspected natural chimerism. Nevertheless, limited clinical awareness and the absence of standardized diagnostic protocols are some of the challenges that continue to delay diagnosis. Overall, in this review we emphasize the progress made in molecular diagnosis, as well as identifying areas to improve on for future research on the accurate diagnosis of human chimerism.

Keywords

Biology, Genetics, Human chimerism, Genetic diagnosis, Next-generation sequencing

Introduction

Traditionally, genetic testing assumes that every individual has a single unique genome inherited at fertilization. However, human chimerism defies this belief by proving that one individual can have two or more genetically different cell populations originating from different genetic sources. Despite the fact that it is considered a rare condition, human chimerism has significant involvement in different branches of medicine, for example, transplantation, genetic diagnosis, forensic science, and precision medicine. Affected individuals, based on the type of chimerism and distribution of genetically different cells in the body, can have a variety of clinical manifestations or remain completely asymptomatic (Conlin et al.; van Bever et al.; Sheets et al.).

One of the major issues associated with human chimerism is the potential of undiagnosed cases; however, case reports cannot show the number of undiagnosed cases. In different individuals, chimerism is discovered accidentally during routine medical examinations due to the inability of traditional diagnostic methods to detect low-level or tissue-specific chimerism. Late identification could complicate the monitoring of transplants, interfere with parentage and forensic testing, as well as result in inaccurate genetic diagnoses. With the increased awareness of human chimerism, researchers noted the necessity to have more precise and reliable diagnostic methods that would allow detecting diverse cell populations at earlier stages (Yu et al.; Sheets et al.; Cho et al.).

Advances in molecular genetics improved the diagnosis of human chimerism significantly. Traditional methods like short tandem repeat (STR), polymerase chain reaction (PCR), and fluorescence in situ hybridization (FISH) were supplemented with new technologies like Next Generation Sequencing (NGS), which allowed for increased sensitivity and the ability to detect low-level chimerism that could not be detected with traditional methods. These advances increased the opportunities for early and accurate diagnosis, while showing the limitations of traditional diagnostic methods (Picard et al.; Zhang et al.).

This overview evaluates how advances in genetic testing have improved the diagnosis of human chimerism. Based on the analysis of updated literature on traditional and emerging molecular diagnosis methods, this review compares their limitations, strengths, as well as clinical applications while considering the challenges that persist and the future outcomes for improving the diagnosis of this genetic condition.

The question guiding this review is: how do advances in genetic testing improve human chimerism diagnosis, and how does their usefulness differ between natural chimerism and post-transplant monitoring? Blouin and Askar's review focuses on hematopoietic transplantation; this overview also considers tissue-dependent diagnosis in natural chimerism, without claiming that the distinction has never been discussed before (Blouin and Askar).

Review Method

This is a narrative review, not a systematic review or a statistical analysis. The provided bibliography was supplemented with targeted searches of PubMed-indexed publications and publisher pages using chimerism with STR, qPCR, digital PCR, NGS, SNP array, nanopore, and disputed maternity. Primary case reports and assay comparisons were selected for their relevance to diagnosis; full texts were accessed where available, while restricted articles were limited to their abstracts or indexed article text. This selection was not exhaustive, and access restrictions limit the comparison.

I. Results and Discussion

Biology of Human Chimerism

Definition of Human Chimerism

Human chimerism is defined as a condition in which some cells or tissues in a person's body contain at least two different sets of DNA (National Cancer Institute). These diverse cell populations can have a natural origin or develop later in life. Genetically different cell lines in natural human chimerism can develop during pregnancy or embryonic development. In contrast, artificial chimerism can be developed as a result of medical procedures, for example, hematopoietic stem cell transplantation, when donor cells are established in the recipient (National Cancer Institute).

This distinction is crucial, since not all situations involving genetic differences between cell populations can be considered chimerism. It should also be differentiated from mosaicism. Mosaicism is defined as emerging from a single zygote, in which genetic changes develop after fertilization. In contrast, constitutional chimerism involves genetically different cell populations from separate fertilization products, while acquired chimerism involves transferred cells. Consequently, despite the fact that both conditions can produce genetic variations within one individual, the biological origin of that variation is different (Conlin et al.).

During genetic diagnosis, this distinction becomes crucial. When an unusual DNA profile is detected, establishing whether it is contamination, mosaicism, constitutional chimerism, or acquired chimerism can require additional investigation. The existence of several explanations for unusual DNA results proves that detection of a second genetic profile should also involve explaining its biological origin (Conlin et al.; Charalsawadi et al.).

Table 1. Chimerism and mosaicism

Feature	Constitutional chimerism	Mosaicism
Cell ancestry	Distinct fertilization products contribute cell lineages.	Genetically different descendants of one zygote.
Mechanism	Embryonic contributions from different fertilization products.	Postzygotic mutation or chromosome changes within a lineage.
Evidence of origin	Informative parental alleles and genome-wide patterns may reveal separate contributions (Conlin et al.; Charalsawadi et al.; Shin et al.).	Allelic and chromosome patterns are interpreted within a single-zygote origin (Conlin et al.).
Shared difficulty	Proportions may vary between tissues; one sample may miss a lineage (Yu et al.; van Bever et al.).	Tissue representation and abundance affect what a sample reveals (Conlin et al.).

This table compares constitutional chimerism with mosaicism. Acquired chimerism instead involves transferred cells, as discussed in the text.

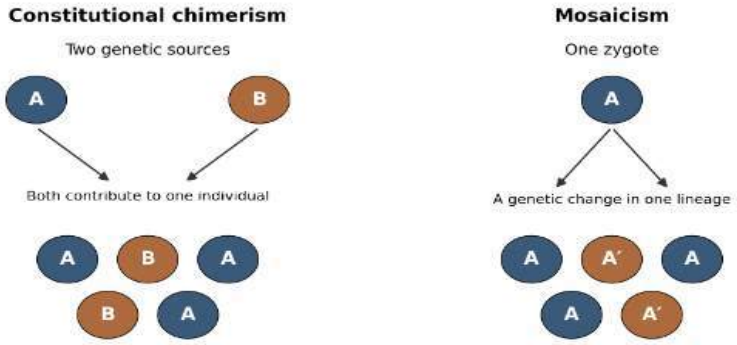

Figure 1. Different genetic sources in constitutional chimerism and a single source in mosaicism. Original conceptual illustration; the symbols show cell ancestry rather than measured proportions. Acquired chimerism is not illustrated (Conlin et al.; Yu et al.; Shin et al.).

Types of Human Chimerism

Tetragametic chimerism develops when two independently fertilized ova produce separate embryos, which later merge during early embryonic development. As a result, instead of

developing as two individuals, their cells merge into a single individual with cell populations from both embryos (Yu et al.).

Tetragametic chimerism can be particularly challenging to diagnose, especially when both original embryos have the same chromosomal sex. An example of this challenge can be seen in the report by Cho and colleagues of a phenotypically normal woman with 46, XX/46, XX chimerism. This case shows how same-sex chimerism can be missed when chromosome testing alone does not reveal the different genetic populations (Cho et al.).

Microchimerism occurs when a comparatively small, genetically different cell population exists inside another individual. Maternal and fetal cells can be exchanged during pregnancy, creating long-term populations of genetically different cells. In contrast to tetragametic chimerism, when genetically different ancestry can contribute to substantial portions of the body, microchimerism occurs at very low levels (Maloney et al.).

Chimerism can also be created artificially. Following allogeneic hematopoietic stem cell transplantation, the hematopoietic system of the recipient contains donor-derived cells, creating an acquired chimeric state. Researchers do differentiate complete donor chimerism from mixed chimerism, when donor and recipient cell populations remain detectable at the same time. Mixed chimerism is expected to change over time and can be classified as increasing, stable, or decreasing (Blouin and Askar).

Complete donor chimerism means that recipient DNA was not detected above the tested assay's threshold in that sample. It does not mean that every cell in the body is donor-derived (Blouin and Askar).

Blood chimerism refers to the coexistence of genetically different cells in the blood. Although it can occur naturally, for example, through cell exchange associated with twin development, it can also be acquired. In contrast to tetragametic chimerism, it can be limited to hematopoietic cells. Significantly, this distinction is important since the genetic profile gained from blood may not reflect the genetic composition of other tissues (National Cancer Institute; Yu et al.).

Blood chimerism describes the tissue in which a mixture is found, rather than a separate cause. It can therefore overlap with natural, transfusion-associated, or transplant-associated chimerism (National Cancer Institute).

Transfusion-associated chimerism is another acquired form, but transfusion does not necessarily lead to a persistent donor cell population. Jackman and colleagues confirmed only one long-term case in a cohort that included 324 transfused trauma patients. This should be distinguished from the donor engraftment expected after allogeneic HSCT (Jackman et al.).

Formation of Human Chimerism

The different types of chimerism describe how genetically different cell populations appear within an individual. However, their formation is based on different biological mechanisms. Understanding these mechanisms is important, since the origin of the second cell population can influence its distribution and be later detected (Yu et al.; Maloney et al.).

One of the mechanisms of human chimerism formation is embryonic fusion. Tetragametic chimerism develops when two separate ova are fertilized separately, creating two genetically different embryos, which later merge during early embryonic development. The original cell populations from both embryos then merge into one individual, creating a coexistence space of two genetically different populations (Yu et al.).

Since the two populations appear early in development, their cells can merge into different parts of the developing body. Consequently, it can create genetic differences between tissues, producing a risk of being undiagnosed since the proportion of each genetic population detected can depend on where a sample is taken. Thus, embryonic fusion provides a biological explanation for how tetragametic chimerism can produce unexpected or inconsistent DNA (Yu et al.; van Bever et al.).

Cell exchange during pregnancy is another mechanism, occurring when cells are exchanged during pregnancy without two embryos merging into one. Since maternal and fetal exchange can introduce genetically different cells into another individual, it can produce microchimerism. In contrast to tetragametic chimerism, this mechanism does not require the merger of two entire embryos into one body. Instead, a comparatively small external cell population appears in the other individual. If genetically different cells appear only at low levels, their diagnosis becomes even more dependent on more accurate and sensitive genetic testing methods (Maloney et al.).

Health Consequences and Implications

Chimerism does not have one health consequence shared by all affected individuals. The van Bever patient had bilateral ovotestes, whereas the Cho patient was phenotypically normal (van Bever et al.; Cho et al.). Maternal microchimerism was also found in healthy adults in the Maloney study; its detection alone does not mean that it causes disease (Maloney et al.). In transplantation, low-level recipient DNA must be interpreted with disease-specific testing since recipient cells are not necessarily malignant (Wu et al.).

II. Challenges in Diagnosing Human Chimerism

Why Chimerism is Often Missed

Chimerism does not necessarily produce an evident phenotype. Therefore, an individual can have more than one genetically different cell population without showing any symptoms or physical characteristics directly related to chimerism. This challenge becomes important for chimerism forms like same-sex tetragametic chimerism. The lack of evident traits is one reason why chimerism can be discovered through unexpected laboratory results. One clear example of this challenge can be related to the van Bever case, in which a 19-year-old phenotypic male was first examined for a painless scrotal mass and was then diagnosed with 46, XX/46, XY tetragametic chimerism through further tissue and genetic examination (Cho et al.; van Bever et al.).

The distribution of genetically different cell populations also represents another challenge. Chimerism does not necessarily appear in equal proportions or populations throughout the body. As a result, the genetic profile obtained from one biological sample may not be representative of the genetic composition of the individual. Another example of this challenge was reported by Yu and colleagues. Initial molecular analysis revealed inconsistencies in maternity testing. Further examination revealed that her blood contained one cell population, while additional genetic contributions were revealed through testing other tissues. Consequently, the results supported a chimerism diagnosis. Such examples show how during laboratory analysis, one sample may not be enough to represent the whole genetic composition of an individual (Yu et al.).

Although the right tissue is selected, the concentration of the secondary genetic population can produce another challenge for correct chimerism diagnosis. Traditional genetic methods have a limit below which a signal cannot be confidently detected. As a result, the same biological sample can be interpreted differently depending on the sensitivity of the method used. To illustrate this challenge, let us take a hypothetical post-transplant sample. A sample containing less than 1% of a certain genotype can be interpreted as complete donor chimerism by a method with a detection threshold of approximately 1–5%. However, the analysis of the same sample using a method capable of detecting chimerism in a low-level genetic population (0.01-0.1%) can reveal the presence of mixed chimerism. Although the biological sample is the same, the laboratory classification can differ based on the detection threshold of the method (Picard et al.; George et al.).

Consequences of Misdiagnosis

The challenges of recognizing chimerism can have many consequences beyond the diagnosis itself. Ignoring the possibility of multiple genetic populations can lead to an unusual DNA profile interpreted as evidence of different biological relationships or inconsistencies with genetic findings. A clear example of something like this would be parentage testing, since these analyses are based on the idea that the tested DNA completely represents the inherited genetic profile of the individual (Yu et al.; Sheets et al.).

The case reported by Yu and colleagues represents this challenge, since initial testing revealed inconsistencies with her biological relationship to her children. Additional tissue testing revealed different cell populations and how they can explain the unexpected results through tetragametic chimerism. This example shows how misdiagnosis can affect the interpretation of other genetic analyses (Yu et al.).

Another challenge would be transplantation, due to post-transplantation chimerism. After hematopoietic stem cell transplantation, the genotype of donor and recipient coexists, and their proportions can change over time. Post-transplantation chimerism can be defined as a dynamic process, rather than a fixed state. These challenges show why advances in genetic testing are important for the diagnosis of human chimerism (Blouin and Askar).

III. Traditional Genetic Testing Methods

Over time, different molecular techniques have been developed and adapted for the purpose of identifying genetic differences that can separate one cell population from another. Methods like PCR are used alongside techniques such as STR, SNP, INDEL, FISH, smMIPS, and TaqMan-based methods. Many of these techniques can be used together, combining PCR for amplification and different genetic markers needed to distinguish cell populations (Gineikiene et al.; George et al.; Picard et al.; Wu et al.; Johnson et al.).

PCR is an amplification process, whereas STRs, SNPs, and INDELs are genetic marker types. They can therefore be combined within one assay rather than treated as competing methods (Picard et al.).

Polymerase Chain Reaction (PCR)

PCR allows selected regions of DNA to be amplified before they are analyzed, becoming an important part of molecular chimerism analysis. This method is really useful when the quantity of genetic material associated with a certain population is small; the amplification allows creating enough copies of targeted sequences for further detection and analysis (Gineikiene et al.).

There are many forms of PCR, like real-time PCR, digital PCR, multiplex PCR, and allele-specific approaches. Since there are many forms, PCR can be adapted depending on the needs of the researchers. For example, multiplex PCR has been used for the detection of chimerism following hematopoietic stem cell transplantation. It is necessary to analyze the DNA using informative genetic markers or other detection methods. Amplification alone is insufficient for chimerism diagnosis because the informative genetic targets and interpretation are required (Gineikiene et al.; George et al.).

Short Tandem Repeat (STR) Analysis

STRs represent DNA segments that include short sequences repeating many times. As the number of repeats varies between individuals, it makes STR suitable for distinguishing one DNA from another. Also, we can use STR loci if an individual has genetically distinct cell populations, which helps to reveal alleles associated with more than one genetic profile. STR fragments can be analyzed according to their size, allowing researchers to distinguish genetic differences using profiles. It can be analyzed through multiple examinations, increasing the possibility of finding differences between two populations (Chen et al.; Charalsawadi et al.).

However, STR has one limitation connected with genetic populations present in a low concentration. Stutter peaks can be created as artifacts during PCR amplification and produce additional signals resembling a true STR allele, interfering with the interpretation of smaller genetic populations (Chen et al.; Picard et al.).

SNPs and Quantitative PCR

A single-nucleotide polymorphism is a variation that occurs at individual positions in the DNA sequence. The SNP can be used to distinguish genetic material by detecting different nucleotides at informative positions in two genetically distinct individuals. Informative SNPs can serve as a marker to identify the presence of more than one genotype in chimerism. qPCR expands this approach since it allows estimating the amount of targeted DNA during amplification. Fluorescence signals are measured during the process, providing quantitative information about genetic sequences being amplified. Gineikiene and colleagues reported quantification to 0.1% with their SNP-specific qPCR system using 100 ng of DNA per reaction. This threshold is specific to that evaluated assay and not to all SNP- or INDEL-based tests (Gineikiene et al.).

Digital PCR

Digital PCR partitions a reaction into many separate compartments and estimates target abundance from positive compartments. George and colleagues reported a limit of detection of 0.01% using two combined wells and a limit of quantification of 0.01% using three wells, with a 25% coefficient-of-variation criterion; their discussion permits DNA loading up to 250 ng per well (George et al.). This result is dependent on the target, DNA amount, and number of wells, so it should not be presented as the sensitivity of every digital PCR assay.

In a comparison involving 281 samples from 28 transplant recipients, Valero-Garcia and colleagues reported median warning intervals before relapse of 63 days with digital PCR and 45.5 days with qPCR (Valero-Garcia et al.). This supports investigating earlier monitoring signals, but does not establish a survival benefit or equivalent performance in natural chimerism.

Other Genetic Approaches

INDEL markers, FISH, TaqMan probes, and single-molecule molecular inversion probes are some of the other approaches used. They help to provide additional information for chimerism diagnosis. Ultrasensitive methods have become a major point of research since they allow diagnosing low-concentration populations. Studies using this approach have already appeared, with single-molecule molecular inversion probe capture and high-throughput sequencing of copy number deletion polymorphism (Wu et al.; Johnson et al.).

FISH helps to locate cells with distinguishable chromosome markers in tissue sections, while SNP arrays help to establish genome-wide patterns and the origin of genetic contributions (Johnson et al.; Conlin et al.; Shin et al.). These methods provide answers to different questions from bulk DNA quantification. Sex-chromosome FISH cannot generally distinguish two XX populations, and the cited array cases do not establish a universal microchimerism detection limit.

IV. Advances in Genetic Testing

Next-Generation Sequencing (NGS)

NGS provides the ability to study much larger amounts of genetic information simultaneously without restriction to a small number of selected markers and allows the examination of many genetic variants in a sample. STR-based analysis examines selected loci, while targeted NGS can examine larger marker panels. As it increases the depth and ability to evaluate many genetic differences, validated NGS assays can identify low-level populations that could be harder to detect using methods with higher detection thresholds (Picard et al.).

NGS can solve one problem of chimerism testing: detecting low-level populations that some STR assays miss. For example, in the case of hematopoietic stem cell transplantation, NGS's ability to identify small amounts of donor or recipient DNA could provide important information following treatment (Pedini et al., "Evaluation"; Zhang et al.).

Picard and colleagues found quantification thresholds of 0.3% for AlloSeq HCT and 0.5% for NGStrack, despite lower-level detection with AlloSeq (Picard et al.). Zhang and colleagues compared 56 serial samples from six transplant patients and reported an NGS turnaround of approximately 19.6–24.8 hours (Zhang et al.). In that laboratory, estimated assay costs per sample were $91–151 for NGS, $99 for qPCR, and $27 for STR; these are study-specific figures rather than current universal prices (Zhang et al.). These results show why marker panels, precision, and turnaround should be compared alongside sensitivity; samples from six patients are not evidence from 56 independent patients.

Third-Generation Sequencing

Third-generation nanopore sequencing can provide faster single-sample testing in a defined transplant-monitoring setting. Pedini and colleagues evaluated a 34-INDEL assay using 30 ng of DNA, with an acceptable range of 5–60 ng, and reported sensitivity of 0.5% (Pedini et al., "Third-Generation"). The workflow required about three hours for one sample, but 20–24 hours for batches; clinical specimens came from only three patients. It may therefore complement urgent monitoring, but does not establish faster diagnosis of unsuspected natural chimerism or better detection of very small populations.

Table 2. Selected assay evidence

Method	Input example	Threshold evidence	Time example	Use and validation limits
STR-PCR	1–5 ng (Picard et al.).	Reported sensitivity 1–5%; separate LOD/LOQ not specified in comparison (Picard et al.).	~15 h (Picard et al.).	Allele identification; stutter interference examined in 30 donors (Chen et al.).
SNP qPCR	100 ng per reaction (Gineikiene et al.).	Reported quantification to 0.1%; separate LOD not confirmed from accessible text (Gineikiene et al.).	Assay-specific time not confirmed.	Informative SNPs; examples from six transplant patients (Gineikiene et al.).
ddPCR	DNA loading up to 250 ng/well discussed (George et al.).	LOD 0.01% with two combined wells; LOQ 0.01% with three, at 25% CV (George et al.).	Assay-specific time not confirmed.	Molecular counting; informative target and sufficient copies required (George et al.).
AlloSeq HCT / NGStrack	10 ng / 20–40 ng per monitoring mix (Picard et al.).	Reported quantification thresholds 0.3% / 0.5%; lower detection distinct (Picard et al.).	~24–48 h; workflow-dependent (Picard et al.).	Transplant mixtures; low-range precision assessed (Picard et al.).
smMIP NGS	Clinical specimens selected with at least 500 ng available; exact reaction input not confirmed (Wu et al.).	Synthetic-mixture detection 1 in 10,000; separate formal LOQ not given in abstract (Wu et al.).	Not specified in accessible abstract.	Molecular identifiers; synthetic performance is not clinical diagnostic sensitivity (Wu et al.).
Nanopore ONtrack	30 ng; acceptable range 5–60 ng (Pedini et al., "Third-Generation").	Reported sensitivity 0.5%; separate formal LOD/LOQ not specified (Pedini et al., "Third-Generation").	~3 h single; 20–24 h batch (Pedini et al., "Third-Generation").	30 samples; clinical specimens from three patients (Pedini et al., "Third-Generation").
FISH	Intact nuclei; count cells examined.	No transferable LOD/LOQ established by this tissue-method report (Johnson et al.).	Not specified in accessible abstract.	Cellular localization; informative chromosome probes required (Johnson et al.).
SNP array	DNA; matched parental samples can aid origin assessment.	No universal low-fraction LOD/LOQ established by the cited cases (Conlin et al.; Shin et al.).	Not confirmed from accessible case reports.	Genome-wide origin evidence; not a substitute for transplant assay validation (Conlin et al.; Charalsawadi et al.; Shin et al.).

LOD = limit of detection; LOQ = limit of quantification; CV = coefficient of variation. Input amounts and turnaround times belong to the cited assays and workflows. Unconfirmed entries reflect limited access to sources, not proof that the original article omitted the detail. Quantitative performance primarily concerns known transplant mixtures or synthetic dilutions, rather than clinical sensitivity for natural chimerism. Costs were not standardized across studies.

V. Current Limitations

Tissue selection remains one of the biggest limitations during diagnosis for chimerism since distinct populations might not be equally distributed throughout the body; if the appropriate tissue is not tested, any method will fail to detect chimerism. Also, finding genetically distinct cell populations does not mean explaining their biological origin, which is why additional investigation is required to determine the type of chimerism (Yu et al.; Sheets et al.; Charalsawadi et al.).

Technical and practical limitations also create a problem during diagnosis. Stutter peaks, for example, created during PCR amplification can interfere with STR analysis. Individually, qPCR needs appropriate informative markers to distinguish between donor and recipient DNA. NGS can be useful, but its equipment, reagent costs, and batching requirements may restrict implementation. The studies reviewed here do not provide a standardized cost comparison. Furthermore, much of the quantitative-method evidence concerns transplantation-associated

chimerism, while natural-chimerism diagnosis is illustrated largely by case reports. In consequence, even though current technologies have improved, further research is still needed to determine how these advances can be more consistently applied to natural chimerism (Chen et al.; Gineikiene et al.; Picard et al.; Blouin and Askar).

VI. Future Implications

One important step considering future implications would be the development of standardized diagnostic protocols that answer when additional tissues should be tested following an unusual genetic result or when a higher-sensitivity method should be introduced. Protocols like this should be validated to determine whether they improve diagnosis when conventional testing produces conflicting results or when a second genetic population is suspected to occur at a very low level. More research on natural chimerism would also be an important step to take, since it raises awareness and could determine whether highly sensitive methods developed primarily for transplantation can be adapted to cases where a second genotype is unknown (Yu et al.; Sheets et al.; Wu et al.).

A proposed workflow would first check sample identity, contamination, and transfusion or transplant history, then confirm an unexpected genetic result with an independent method. For suspected natural chimerism, additional appropriate tissues and informative parental markers may clarify distribution and origin (Yu et al.; Sheets et al.; Charalsawadi et al.). For transplantation, testing should use validated donor–recipient markers, consistent sampling, and serial measurements interpreted with disease-specific findings (Blouin and Askar; Valero-Garcia et al.). This is a synthesis of the reviewed evidence, not a validated clinical protocol. Future validation should compare matched tissues and assays using predefined detection and quantification criteria.

Areas like assisted reproductive technologies (ART), such as IVF and embryo transfer, could be explored further to gain a better understanding of their association with chimerism. Future studies would need appropriate comparison groups before drawing conclusions about these associations. The Sheets report concerns preventing mistaken interpretations of parentage results, rather than preventing chimerism itself. Nevertheless, there is no current evidence that IVF or ART could become a method to prevent chimerism (Sheets et al.).

Conclusion

Because of genetically distinct populations that may be present at low levels, human chimerism is a challenging diagnosis. Its uneven distribution throughout the body and sometimes asymptomatic appearance make this condition challenging to recognize. Advances in genetic testing have improved the ability to identify it. STR-PCR remains an established approach in post-transplant monitoring, while validated quantitative PCR, digital PCR, and targeted sequencing assays can detect smaller genetic contributions under specified conditions (Picard et al.; Gineikiene et al.; George et al.).

On the other hand, technological sensitivity alone cannot guarantee a diagnosis. Variables like tissue selection, interpretation of low-level genetic results, cost, and the limited amount of specialized research are all matters that slow down the accurate diagnosis. However, advances in

genetic testing now allow us to detect, distinguish, and quantify genetically distinct cell populations (Yu et al.; Sheets et al.; Zhang et al.).

For natural chimerism, the main outcome is a diagnosis supported by tissue distribution and genetic origin. For transplant-associated chimerism, the principal resolution is reproducible quantification and interpretation over time. These differences should guide test selection rather than assuming that higher analytical sensitivity alone guarantees an accurate diagnosis.

AI Acknowledgement

During the preparation of this work, the author used OpenAI ChatGPT to brainstorm research topics, improve language and readability, organize sources, and prepare comparison tables and an original conceptual figure. No patient-level data analysis was performed. The author is responsible for reviewing and editing the resulting content and for the final submitted manuscript.

Authors

Victoria Cardelli is an international student at Wilbraham & Monson Academy with an interest in biology, genetics, and medicine. She has participated in Stanford and University of Chicago summer programs and conducted research on human chimerism. Her academic interests focus on improving genetic diagnosis and advancing equitable access to healthcare.

Works Cited

Blouin, A. G., and M. Askar. “Chimerism Analysis for Clinicians: A Review of the Literature and Worldwide Practices.” Bone Marrow Transplantation, vol. 57, no. 3, 2022, pp. 347–359. https://doi.org/10.1038/s41409-022-01579-9.

Charalsawadi, C., et al. “Case Report: Molecular Analysis of a 47,XY,+21/46,XX Chimera Using SNP Microarray and Review of Literature.” Frontiers in Genetics, vol. 13, 2022, article 802362. https://doi.org/10.3389/fgene.2022.802362.

Chen, D. P., et al. “Systematic Analysis of Stutters to Enhance the Accuracy of Chimerism Testing.” Annals of Clinical & Laboratory Science, vol. 38, no. 3, 2008, pp. 264–272. https://pubmed.ncbi.nlm.nih.gov/18715856/.

Cho, S., et al. “A Case of 46,XX/46,XX Chimerism in a Phenotypically Normal Woman.” International Journal of Legal Medicine, vol. 134, no. 6, 2020, pp. 2045–2051. https://doi.org/10.1007/s00414-020-02296-y.

Conlin, L. K., et al. “Mechanisms of Mosaicism, Chimerism and Uniparental Disomy Identified by Single Nucleotide Polymorphism Array Analysis.” Human Molecular Genetics, vol. 19, no. 7, 2010, pp. 1263–1275. https://doi.org/10.1093/hmg/ddq003.

George, D., et al. “Detection and Quantification of Chimerism by Droplet Digital PCR.” Chimerism, vol. 4, no. 3, 2013, pp. 102–108. https://doi.org/10.4161/chim.25400.

Gineikiene, E., et al. “Single Nucleotide Polymorphism-Based System Improves the Applicability of Quantitative PCR for Chimerism Monitoring.” The Journal of Molecular Diagnostics, vol. 11, no. 1, 2009, pp. 66–74. https://doi.org/10.2353/jmoldx.2009.080039.

Jackman, R. P., et al. “Lack of Persistent Microchimerism in Contemporary Transfused Trauma Patients.” Transfusion, vol. 59, no. 11, 2019, pp. 3329–3336. https://doi.org/10.1111/trf.15518.

Johnson, K. L., et al. “The Use of Fluorescence in Situ Hybridization (FISH) on Paraffin-Embedded Tissue Sections for the Study of Microchimerism.” BioTechniques, vol. 29, no. 6, 2000, pp. 1220–1224. https://doi.org/10.2144/00296st01.

Maloney, S., et al. “Microchimerism of Maternal Origin Persists into Adult Life.” The Journal of Clinical Investigation, vol. 104, no. 1, 1999, pp. 41–47. https://doi.org/10.1172/JCI6611.

National Cancer Institute. “Chimerism.” NCI Dictionary of Cancer Terms, https://www.cancer.gov/publications/dictionaries/cancer-terms/def/chimerism. Accessed 7 Oct. 2026.

Pedini, P., et al. “Evaluation of Next-Generation Sequencing and Crystal Digital PCR for Chimerism Monitoring of Post-Allogeneic Hematopoietic Stem Cell Transplantation.” Transplantation and Cellular Therapy, vol. 27, no. 1, 2021, pp. 89.e1–89.e10. https://doi.org/10.1016/j.bbmt.2020.09.023.

Pedini, P., et al. “Third-Generation Nanopore Sequencing for Post-Transplant Chimerism Monitoring.” HLA, vol. 107, no. 6, 2026, article e70770. https://doi.org/10.1111/tan.70770.

Picard, C., et al. “New Methods for the Quantification of Mixed Chimerism in Transplantation.” Frontiers in Immunology, vol. 14, 2023, article 1023116. https://doi.org/10.3389/fimmu.2023.1023116.

Sheets, K. M., et al. “A Case of Chimerism-Induced Paternity Confusion: What ART Practitioners Can Do to Prevent Future Calamity for Families.” Journal of Assisted Reproduction and Genetics, vol. 35, no. 2, 2018, pp. 345–352. https://doi.org/10.1007/s10815-017-1064-6.

Shin, S. Y., et al. “Identification of the Mechanism Underlying a Human Chimera by SNP Array Analysis.” American Journal of Medical Genetics Part A, vol. 158A, no. 9, 2012, pp. 2119–2123. https://doi.org/10.1002/ajmg.a.35476.

Valero-Garcia, J., et al. “Earlier Relapse Detection after Allogeneic Haematopoietic Stem Cell Transplantation by Chimerism Assays: Digital PCR versus Quantitative Real-Time PCR of Insertion/Deletion Polymorphisms.” PLoS One, vol. 14, no. 2, 2019, article e0212708. https://doi.org/10.1371/journal.pone.0212708. Corrected version republished 1 Mar. 2019; correction, https://doi.org/10.1371/journal.pone.0213966.

van Bever, Y., et al. “Multiparameter Investigation of a 46,XX/46,XY Tetragametic Chimeric Phenotypical Male Patient with Bilateral Scrotal Ovotestes and Ovulatory Activity.” Sexual Development, vol. 12, no. 1–3, 2018, pp. 145–154. https://doi.org/10.1159/000479946.

Wu, D., et al. “Ultrasensitive Detection of Chimerism by Single-Molecule Molecular Inversion Probe Capture and High-Throughput Sequencing of Copy Number Deletion Polymorphisms.” Clinical Chemistry, vol. 64, no. 6, 2018, pp. 938–949. https://doi.org/10.1373/clinchem.2017.284737.

Yu, N., et al. “Disputed Maternity Leading to Identification of Tetragametic Chimerism.” The New England Journal of Medicine, vol. 346, no. 20, 2002, pp. 1545–1552. https://doi.org/10.1056/NEJMoa013452.

Zhang, A., et al. “Engraftment and Measurable Residual Disease Monitoring after Hematopoietic Stem Cell Transplantation: Comparison of Two Chimerism Test Strategies, Next-Generation Sequencing versus a Combination of Short-Tandem Repeats and Quantitative PCR.” The Journal of Molecular Diagnostics, vol. 26, no. 4, 2024, pp. 233–244. https://doi.org/10.1016/j.jmoldx.2024.01.007.

An Assessment of LSTM and XGBoost's Comparability in Financial Prediction

By Vuong Trac Truong

Abstract

Financial prediction has become a common testbed within the field of machine learning, yet two widely used approaches- LSTM and XGBoost- do not have a clearly established superior. The problem here is that changing the model often changes much of the experiment and its conditions. Studies vary in lookback windows, tuning effort, and test markets, making it difficult to determine whether performance differences stem from the experimental setup or the models themselves. This review therefore examines ten sources and applies three questions to studies reporting original results: whether the lookback window was tested rather than assumed, whether tuning effort was comparable, and whether more than one market was used. No study reviewed here satisfies all three criteria, and only two directly compare LSTM and XGBoost. Even so, the closest result places the models within 0.0004 F1 of each other- a negligible difference. However, that study tunes the lookback window only for the LSTM. The literature, therefore, does not establish either model as superior. Rather, it reveals what a future comparison would need before the question can be answered definitively.

Introduction

Predicting stock prices has become a common application of machine learning, with approaches including but not limited to traditional statistical methods and deep neural networks (Sezer et al.). Among others, one method is the recurrent neural network, which processes input data in order and carries information from one time step to the next. However, standard recurrent models can struggle to retain information over longer stretches of time, as gradients can shrink during training and reduce the influence of earlier observations. Long short-term memory (LSTM) networks can counteract this observed weakness through the use of memory cells and gates that control what information is retained throughout (Waqas and Humphries). Because of this, temporal order is an important parameter in time-series forecasting, as the design philosophy helps determine what kinds of relationships the model can find (Lim and Zohren).

This review has selected LSTM and XGBoost due to the difference in how they approach price history. On one hand, an LSTM processes historical data as a sequence in chronological order, whereas XGBoost receives the same history, albeit as separate features in a table. Tree-based models such as XGBoost have historically performed strongly on tabular data, a fact that is corroborated by Grinsztajn et al., where tree-based models outperformed extensively tuned deep-learning models across forty-five data sets, whilst Shwartz-Ziv and Armon found that boosted trees, even with less tuning, could match or even outperform deep models. That is not to say other methods, such as traditional statistical models or transformers, are irrelevant. However, they fall outside the scope of this review.

Indeed, pre-existing studies have not yet settled whether each method is superior in this instance. For instance, Bhandari et al. claim LSTM performs well for next-day stock prediction,

but this is only in relation to itself, rather than another model family. Likewise, Oukhouya et al. compare LSTM, XGBoost, and a hybrid model across six different markets, yet some of the claims in their discussion do not seem to reflect their results. Kang finds that LSTM and XGBoost perform almost identically on a three-class movement task and also shows that the optimum lookback window changes with model size, rather than being a fixed value. Saud and Shakya, meanwhile, treat lookback length as the main experimental variable and reach a substantially different optimum. Other review papers seem to organize this literature by model and application (Sezer et al.), but give less attention to whether differences in experimental design make the design philosophies themselves difficult to compare.

This review therefore asks to what extent previous literature can ascertain the efficacy of LSTM compared to XGBoost in stock forecasting, taking into account lookback window, tuning effort, and market selection. This paper argues that existing comparisons between LSTM and XGBoost are, by design, insufficient to conclude whether one model is better. To do this, we examine ten sources: four report original equity forecasting results, and only two compare the two architectures directly. We select relevant evidence for each of the three design choices, then use the combined results to determine which experiments maintain rigor and which are limited by design in answering this specific question. This will be a secondary data analysis with qualitative analysis, rather than an empirical study, since every figure discussed here will be as reported by the original authors. First, we will cover what each model does and its properties. We then analyze the evidence for sequential memory and the evidence for the tabular-data account against those choices, before turning to the head-to-head studies that sit between them. Finally, we synthesize the studies to identify which findings remain valid and what that means for how future comparisons should be designed.

Scope

We analyzed ten sources, largely categorized into three distinct groups: four empirical studies, two benchmarks of tree ensembles against deep models on tabular data, and four surveys/method papers that establish what each model does and how the field has used it.

The three design choices named above were applied to each study reporting original results, in the form of three questions:

- Was the lookback window tested, or assumed?
- Was tuning effort reported for both model families?
- Was more than one market used?

Background

A recurrent neural network, or RNN, processes price history in order, carrying information from one day into the next. Standard RNNs, however, struggle to retain information over longer periods because the influence of earlier data weakens during training in what is known as the vanishing gradient problem (Waqas and Humphries). LSTMs reduce this through a memory cell and three gates (input, forget, and output) that control what is stored, removed, or

passed forward. Because these gates are learned during training, this LSTM can selectively retain or discard information from different points within the supplied lookback window, allowing useful earlier signals to influence later predictions. Although Waqas and Humphries discuss LSTMs in hydrology and not finance, the mechanism is identical; only the type of time-series data changes.

Indeed, the amount of past data given to the model, otherwise known as the lookback window, is an important part of time-series forecasting. A 30-day lookback window, for example, means the model uses the previous 30 days to make its prediction. An LSTM processes those days as a sequence and compresses their information into an internal state, with the network learning which patterns are worth retaining (Lim and Zohren).

XGBoost handles the same window differently. It does not process the days sequentially. It combines many decision trees, each of which makes decisions by splitting the input data according to individual features. These trees are not built independently; they are added one after another, with each new tree attempting to correct the errors left by the existing model, which is what the "boosting" in gradient boosting refers to (Chen and Guestrin). To give XGBoost a 30-day price history, the sequence must be flattened into 30 separate input columns. The model might use day 3, then day 27, then day 8 when making a prediction. The days still retain their position within the window, but XGBoost has no built-in mechanism that processes them in chronological order.

This makes the lookback window an important setting shared by both models, as the studies reviewed below demonstrate. For an LSTM, it determines the length of the sequence the network learns from; for XGBoost, it determines how many past-day features the trees can use. Giving both models the same window creates what this review calls input parity: both receive the same historical information. Input parity alone, however, does not guarantee a fair comparison. Fixing one window without testing alternatives assumes that both architectures respond to its length in the same way.

Sequential Memory

The main argument for LSTMs is how well their structure matches time-series data. Prices arrive in order, and information from previous days may remain useful for different lengths of time. Through its gates, an LSTM can learn what should be retained or forgotten as the sequence progresses (Waqas and Humphries). A tree given the same flattened window has no equivalent mechanism.

The same reasoning shows up in how widely recurrent models are used in financial prediction. Sezer et al., reviewing studies published between 2005 and 2019, found recurrent models, mostly LSTMs, in more than half of the time-series papers surveyed. Their findings are less decisive than this frequency might suggest, however. Deep-learning models exceeded traditional machine-learning approaches in most studies, performed similarly in many others, and lost in two.

The study by Bhandari et al. shows where this evidence becomes more limited. It

forecasts the next-day closing price of the S&P 500 using market, macroeconomic, and technical inputs, and finds that a single-layer LSTM with approximately 150 neurons performs better than the multilayer alternatives tested. Adding layers reduced out-of-sample accuracy, and dropout did not recover the lost performance. The authors suggest either overfitting or greater architectural complexity as possible explanations, but do not distinguish between the two. They also do not revisit these explanations after dropout, a regularization method intended to reduce overfitting, fails to improve performance. The result is still useful on its own terms, since it shows that added complexity did not improve performance in this study.

Against the three questions used in this review, though, the study leaves several questions. It defines an input window but never reports its value or varies it during tuning, and it tests only one index. Because only one model family is included, the question of whether tuning effort was comparable between LSTM and XGBoost does not apply. Most importantly, there is no competing architecture: the comparison is LSTM against other LSTM configurations. The study therefore provides useful evidence about how LSTM complexity affects S&P 500 forecasting, but it cannot establish whether LSTM performs better than a tree-based model.

Sequential memory may still make LSTMs well suited to price data. The studies reviewed here simply do not isolate that potential advantage from the other experimental choices that change alongside the architecture.

Tabular Data

The tabular benchmarks introduced above are worth examining in more detail, because the strength of their evidence also limits how far it can be applied to price forecasting. Grinsztajn et al. assembled 45 tabular datasets and compared tree-based models versus extensively tuned deep-learning models, including a Transformer they describe as a "best case" for deep learning on tabular data. On medium-sized datasets of roughly ten thousand samples, tree models consistently performed better and also trained faster.

Importantly, this advantage cannot easily be explained by a lack of tuning. The authors spent approximately 20,000 compute hours searching for effective hyperparameters and found that tree models remained ahead at every search budget. Their explanation instead concerns the types of patterns each architecture handles well. Neural networks tend to struggle when relationships change sharply across distinct regions of the data, and when many input variables contain little useful information. Both are common characteristics of tabular datasets.

Shwartz-Ziv and Armon reach a similar conclusion. Across eleven tabular datasets, XGBoost outperformed several deep-learning architectures designed specifically for tabular data whilst also requiring less tuning. They also found, though, that combining the deep models with XGBoost performed better than XGBoost alone, which suggests the two approaches may capture partly different patterns within the same data.

Neither result transfers cleanly to price forecasting. Grinsztajn et al. specifically required their datasets to contain independent and identically distributed observations, meaning that one row should not depend upon the rows before it, and therefore excluded time-series data.

Shwartz-Ziv and Armon do not impose the same restriction, but their datasets are still treated as static tables, and none of the models compared treat the rows as an ordered sequence. Neither benchmark, then, tests the feature that gives LSTMs their main theoretical advantage: the ability to process information in temporal order.

There is also a second mismatch. Grinsztajn et al. require columns to represent different kinds of features, giving signal data (where each column records the same signal from a different sensor) as an example of what this criterion excludes. A flattened price window would fail the same rule for a related reason. It may look like ordinary tabular data, with one row containing many numerical columns, but its columns represent the same price signal at different time lags, not distinct measured features. The price one day ago and the price twenty days ago are two positions within the same price history, not two different types of information.

A flattened price window would therefore conflict with two of Grinsztajn et al.'s stated selection criteria: its observations are temporally dependent, and its columns represent the same type of signal across several lags. This does not weaken the benchmark itself, but it does limit how directly its findings can be applied to price forecasting. These studies provide strong evidence that tree-based models perform well against deep learning on conventional tabular data. They do not establish that the same advantage remains once temporal order becomes part of the prediction problem.

Head-to-Head Comparisons

Only two studies reviewed here directly compare LSTM and XGBoost. Oukhouya et al. forecast daily closing prices for six international indices, representing six different markets, using grid search with blocked time-series cross-validation for tuning. Their hybrid model, which applies XGBoost to the errors left by the sequence model, performs best on most reported test measures. It records the lowest RMSE and highest R2 across all six indices, whilst XGBoost leads on two other measures and LSTM on one. Two issues make the individual comparison harder to interpret.

First, the discussion does not match the paper's own results table. The authors state that XGBoost performs particularly well on DAX whilst LSTM performs well on NASDAQ. XGBoost, however, leads none of the reported measures on DAX. Its only individual wins occur on two NASDAQ measures, where LSTM leads none of the reported measures. The two indices highlighted in the discussion appear to be reversed.

Second, although both models are tuned, the resulting model capacities are difficult to compare. Reported XGBoost tree depths reach 60 for one index and 30 for two others, alongside training errors close to zero. The paper does not discuss whether this indicates overfitting, but such high-capacity trees combined with near-zero training error are consistent with that possibility. This LSTM, meanwhile, is trained for a fixed 100 epochs. The study therefore reports tuning for both architectures but lacks a common measure of tuning effort or model capacity that would show whether the comparison is balanced.

Kang produces a much closer result. Across Korean equities from 2006 to 2024, this

LSTM achieved an F1 score of 0.4312 and XGBoost 0.4316, compared with a baseline of 0.1852. A gap of 0.0004 is not evidence of a meaningful advantage for either architecture.

The inputs, however, are not identical. This LSTM receives raw price and volume data, whilst XGBoost receives engineered technical indicators and undergoes 70 rounds of tuning. Kang is also the only head-to-head study reviewed here to systematically test lookback length, sweeping six window sizes across five LSTM model sizes. The results show that the best-performing window changes with model size, supporting treating lookback length as an experimental variable rather than a setting fixed in advance.

Two limitations remain. The window sweep is performed only for the LSTM, so the same question is never tested for XGBoost. Kang also predicts one of three price-movement classes and not the next price itself, making it a trend-prediction task under the distinction used by Sezer et al. The result is therefore useful for comparing the architectures, but it does not directly answer which performs better for price regression.

Two further studies provide evidence about lookback length without including a tree-based model. Saud and Shakya test seven windows ranging from 2 to 30 days on two Nepalese bank stocks. They conclude that GRU performs better overall, yet their results show this depends on the window tested. LSTM produces lower average MAPE at lookback periods of 2 and 5 days, and GRU performs better only once the window exceeds 5 days. This matters because the authors themselves identify short windows as preferable, and report the lowest average LSTM MAPE at a lookback of 2 days for both stocks. Their overall model ranking therefore depends in part on averaging across window settings where the relative performance of the architectures has already changed.

More importantly for this review, their preferred LSTM window differs substantially from Kang's. Saud and Shakya identify a best window of 2 days, whilst Kang reports a best setting of 100 days under its tested configuration. These should not be treated as a direct contradiction, since the studies use different markets, targets, and experimental arrangements. The roughly fifty-fold difference instead strengthens a narrower point: there is little basis for assuming that one fixed lookback window will stay appropriate across studies or model configurations.

Discussion

Across the studies reviewed, the lookback window clearly plays a large role in performance, and controlling it is more difficult than setting the same value across all models. Indeed, Kang shows that the best-performing window changes with model size, while Saud and Shakya go deeper, showing how the relative ranking of LSTM and GRU changes with the window itself, with LSTM performing better on shorter windows and GRU on longer ones, highlighting that fixing a window in advance can influence which method would appear ‘better’.

Notably, neither study identifies a universal optimal lookback window. Saud and Shakya favor 2 days, whereas Kang chooses 100 days under different experimental conditions. However, these results are not necessarily contradictory, given the different markets, targets, and model

configurations between the two studies. Instead, the difference suggests that the best window may depend on the surrounding experiment itself. Giving two models the same amount of historical data may provide input parity, but it cannot guarantee a fair comparison if that window is more efficient for one model than another. Therefore, a stronger comparison would fine-tune window lengths for both methods.

In the case of tuning, the significance appears to be even more ambiguous. Outside of finance, Grinsztajn et al. show that tree-based models remain strong even after roughly 20,000 compute hours of hyperparameter search, while Shwartz-Ziv and Armon find that XGBoost generally requires less tuning than the deep-learning models they test. Here, these studies provide strong evidence for XGBoost on tabular data, but neither tests financial time series.

Within finance studies, Kang and Oukhouya et al. tune the models they compare. However, the tuning effort is difficult to compare across methods, as epoch counts, search iterations, and tree depth refer to different parts of training, and thus cannot be treated as equivalent measures of effort. As a result, a study can tune both models without it being truly balanced, given the ambiguity of tuning effort.

Market selection is even less well supported. Oukhouya et al. are the only authors in this review to test across multiple markets, but their strongest overall results come from the hybrid model rather than either LSTM or XGBoost alone. The remaining financial studies are each limited to a single national market: Korea in Kang, Nepal in Saud and Shakya, and the United States in Bhandari et al. Results from one market may still be useful, but they cannot show whether an observed advantage comes from the architecture itself or from characteristics of that particular market.

Altogether, no single empirical study reviewed abides by the three criteria set. Because of this, how commonly reported conclusions are severely limited in how they can be interpreted. For example, Bhandari et al. cannot support a comparison between LSTM and tree-based models, as tree-based models were not included in their investigation. Likewise, Oukhouya et al. make market-specific claims that were not supported by their own results table, and Saud and Shakya's conclusion that GRU performs better than LSTM is not true in every case, as LSTM exhibits greater performance at shorter lookback windows. Tabular benchmarks are also poor direct evidence for time-series financial forecasting, since Grinsztajn et al. exclude temporally dependent data by the experiment's design and Shwartz-Ziv and Armon treat their datasets as static tables.

The closest result to a definitive, direct answer lies in Kang. Under the conditions of the study, the two models perform within 0.0004 F1 of each other. Here, the gap is too small to claim any meaningful advantage or performance difference between the two methods. At the same time, the result also remains tied to the conditions of said experiment, so it cannot answer the broader question by itself.

Conclusion

Existing comparisons between LSTM and XGBoost often offer conflicting views over

which model is more suitable for stock price prediction, but the studies reviewed here rarely vary only the method used by themselves. The two models interpret price history differently, with LSTM carrying information forward through gates in chronological order (Waqas and Humphries), whereas XGBoost receives a flattened input without a function for processing data sequentially (Chen and Guestrin). Because of this, the lookback window becomes an important shared variable when comparing the two.

Much of the evidence surrounding either model, however, does not make any direct claims regarding that comparison. Recurrent models seem to make up more than half of the financial time-series studies surveyed by Sezer et al., but the Bhandari et al. study reviewed here includes no competing models and does not note any lookback window. On the other hand, the tabular benchmarks by Grinsztajn et al. and Shwartz-Ziv and Armon provide strong evidence for tree-based models on tabular data, but neither tests temporal order as part of the prediction problem. Evidence on either side, therefore, appears to leave the direct comparison unresolved.

Regardless, the main issue is that existing studies do not separate differences by experimental setup, and those by the method. In the future, a stronger comparison would need to test the lookback window for both methods, report tuning effort in a comparable manner, and evaluate the models across multiple markets.

Overall, the reviewed studies do not establish that either LSTM or XGBoost is generally superior for stock prediction. Instead, it shows why the existing evidence is not sufficient to settle the question. A benchmark designed specifically around these comparison conditions would be needed to answer the question more definitively.

Works Cited

Bhandari, H. N., et al. "Predicting Stock Market Index Using LSTM." Machine Learning with Applications, vol. 9, 2022, article 100320, https://doi.org/10.1016/j.mlwa.2022.100320.

Chen, T., and C. Guestrin. "XGBoost: A Scalable Tree Boosting System." Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Association for Computing Machinery, 2016, pp. 785–794, https://doi.org/10.1145/2939672.2939785.

Grinsztajn, L., et al. "Why Do Tree-Based Models Still Outperform Deep Learning on Typical Tabular Data?" Advances in Neural Information Processing Systems, edited by S. Koyejo et al., vol. 35, Curran Associates, 2022, pp. 507–520.

Kang, S. "Stock Price Prediction Using Triple Barrier Labeling and Raw OHLCV Data: Evidence from Korean Markets." arXiv, 2025, arXiv:2504.02249, https://doi.org/10.48550/arXiv.2504.02249.

Lim, B., and S. Zohren. "Time-Series Forecasting with Deep Learning: A Survey." Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, vol. 379, no. 2194, 2021, article 20200209, https://doi.org/10.1098/rsta.2020.0209.

Oukhouya, H., et al. "Forecasting International Stock Market Trends: XGBoost, LSTM, LSTM-XGBoost, and Backtesting XGBoost Models." Statistics, Optimization & Information Computing, vol. 12, no. 1, 2024, pp. 200–209, https://doi.org/10.19139/soic-2310-5070-1822.

Saud, A. S., and S. Shakya. "Analysis of the Look Back Period for Stock Price Prediction with RNN Variants: A Case Study on the Banking Sector of NEPSE." Procedia Computer Science, vol. 167, 2020, pp. 788–798, https://doi.org/10.1016/j.procs.2020.03.419.

Sezer, O. B., et al. "Financial Time Series Forecasting with Deep Learning: A Systematic Literature Review: 2005–2019." Applied Soft Computing, vol. 90, 2020, article 106181, https://doi.org/10.1016/j.asoc.2020.106181.

Shwartz-Ziv, R., and A. Armon. "Tabular Data: Deep Learning Is Not All You Need." Information Fusion, vol. 81, 2022, pp. 84–90, https://doi.org/10.1016/j.inffus.2021.11.011.

Waqas, M., and U. W. Humphries. "A Critical Review of RNN and LSTM Variants in Hydrological Time Series Predictions." MethodsX, vol. 13, 2024, article 102946, https://doi.org/10.1016/j.mex.2024.102946.

Food Insecurity's Impact on Hypotension and its Effects on Dementia By Ronin Back

Abstract

There have been a plethora of studies on how hypertension and medical obesity can affect the probability of developing dementia later in life and its severity. Yet, there is far less research on how hypotension and being underweight can affect the probability of developing dementia and worsening the symptoms and what can cause these symptoms. This paper aims to answer that question. Anorexia/being underweight can often cause low blood pressure, which leads to white matter lesions and axonal deterioration. These symptoms can be linked to malnutrition which is driven by food insecurity.

The author employed a secondary data analysis from the data which is gathered from IPUMS National Health Interview Survey including data on age, reported dementia, educational levels, and levels of food security. Using published papers that break down how low blood pressure in adults can predict the probability that they would experience dementia or cerebral damage. If an individual experiences malnutrition early in life, it is likely to lead to an increase in brain lesions, and data analysis from countries/regions that have a malnourished population, accounting for confounding variables support this conclusion. As well as employing data analysis of malnutrition and food insecurity to investigate the link between high rates of food insecurity and increased rates of dementia in the population.

There is sufficient correlative evidence that is statistically significant to support the claim which is consistent with but does not establish the proposed pathway. This allows for a more well-rounded look into the causes of dementia and various factors that can worsen dementia once one has developed symptoms.

Keywords: Dementia, Food Insecurity, Malnutrition, Alzheimer's disease

Introduction

There have been a range of studies reviewing and analyzing how hypertension (HTN) - when an individual experiences a systolic blood pressure (BP) of 130 mm Hg or greater and a diastolic BP of 80mm Hg or higher (Iqbal and Jamal, 2023)- and medical obesity can affect the probability of developing dementia later in life and its severity. However, there is a relatively large gap regarding hypotension, when the blood pressure is abnormally low, which would read as below 90/60 mm Hg; malnutrition and anorexia affect the probability of developing dementia as one ages. Still, there is evidence that anorexia increases the amount of neurofilament light chain protein (NfL), a biomarker for neuronal injury, found in the brain, which serves as an indicator for neuronal death (Dobrescu et al., 2019). There is further evidence that orthostatic hypotension (OH) - a sudden drop in blood pressure when standing up from a supine position (Ringer, 2025) later in life can be used to help determine if there will be a higher probability of an individual developing dementia by observing rapid decreases in systolic blood pressure (SBP) (Wolters, 2016). This is meaningful because there is a link to malnutrition and its probability of

developing OH (Kocyigit , 2018). Once one has experienced Alzheimer's Disease (AD) or Vascular Dementia (VaD) there is a higher chance that they would begin to experience more extreme symptoms as well as increasing their probability of undergoing mixed dementia because once one begins to experience dementia their appetites decrease leading to increased rates of anorexia making it harder for the body to regulate its blood flow leading to hypotension and OH (Sarhan, 2025). Still, there is a higher probability of experiencing one of the other two main types of dementia as measured from a study of a sample population finding only roughly 7% of the individuals with AD, VaD and Mixed Dementia had mixed dementia. (Lin et al., 2017)This leads to the body being unable to have blood perfuse the brain, which causes neuronal and white matter lesions in addition to increasing plaque buildup within the brain. The damage is specifically located around the hippocampus, thalamus, and frontal cortex (Cheng et al., 2024). This damage can be mimicked by the buildup of plaque cells in the brain which is often caused by a failure of the glial cells to regulate substances in the brain causing the damage that often manifests as AD(Abdulkhaliq et al., 2026).

This paper aims to investigate the correlative links and mechanisms between malnutrition and its driving effects on dementia using peer-reviewed papers and empirical data gathered from preexisting surveys issued by globally recognized health institutes.

In recent neurobiology papers analyzing malnutrition and hypotension's effects on the brain, there have been studies showing how hypotension affects the likelihood of developing dementia later in life through measurement of SBP immediately after standing (Ma, Blacker, and Hofman, 2023) and how hypotension can cause lesions in the white matter of the brain and hypoperfusion of the hippocampus, Frontal Lobe, and Thalamus. All of which are large contributors to AD, (VaD), and Mild Cognitive Impairment. (Anwer, et al., 2026). Yet there's more evidence that malnutrition and anorexia are both contributors to cognitive decline. This can be seen in how malnutrition and anxiety can drive neurodegeneration, with malnutrition playing the largest role but not linking it to hypotension (Wang et al., 2024). This research - while demonstrating malnutrition's effects on a population - is limited in scope because of the geographical limitations of the survey portion and is also limited because it doesn’t demonstrate how hypotension, brought about by this malnutrition, is a driving factor.

Analyzing studies on population health, there have been strong causal findings between food insecurity, and its driving force behind malnutrition and its impacts on cognitive functions. After studies conducted in China observing the aging population, they found that malnutrition served as an independent factor that drove dementia rates up among the aging population (Li et al., 2026). This idea was further supported by another finding showing that while there are other factors that are driven by malnutrition, specifically in older age, can still drive increased rates of cognitive decline (Feng et al. 2022). Still, there is strong evidence to support the idea that food insecurity can drive malnutrition. A study performed in Greece found that out of a population of surveyed individuals were 50.4% food secure and 70.2% of the sample was malnourished, showing heavy overlap between these two factors (Gkiouras et al. 2020). These findings were further supported by a study demonstrating that those experiencing food insecurity would often

experience caloric deficiency and when able to get food, they would often purchase food of poor quality that would also be lacking in nutrients driving overall rates of malnutrition (Young et al., 2023).

When an individual experiences elevated rates of food insecurity, they are more probable to experience heightened risks of malnutrition which has a link to hypotension and heightened rates of dementia diagnoses. Therefore, if an individual undergoes increased rates of food insecurity, there should be a link to higher rates of individuals experiencing dementia diagnoses.

This research is significant because it can demonstrate a stronger link between food insecurity and dementia rates, allowing for more efficient detection and prevention of cognitive impairment, most common in the form of dementia. It also allows for more effective treatment of cognitive decline once it has started since malnutrition has also been linked to worsened cognitive capabilities and acceleration of cognitive decline in those with dementia (Sanders et al., 2017).

Materials & Methods

Figure 1 Mean Age in which an Individual Would Experience Dementia as a Function of Food Security

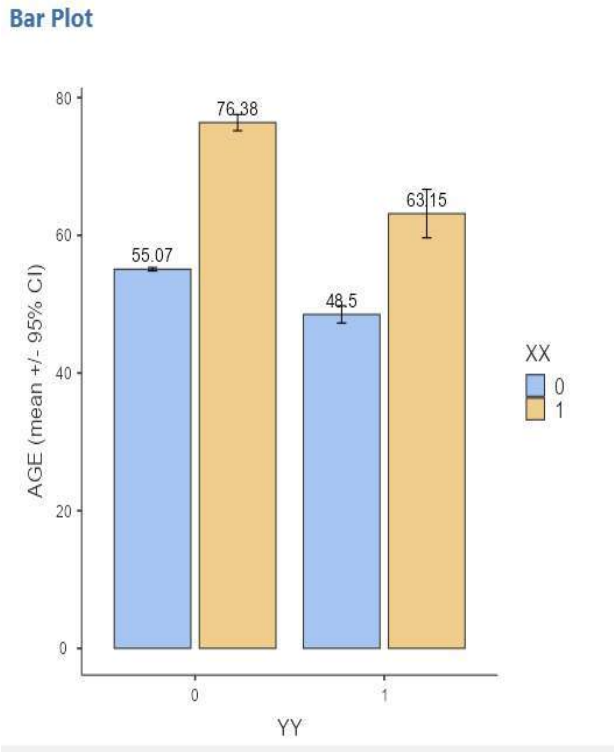

Figure 2 Correlation between DataforC and Poverty

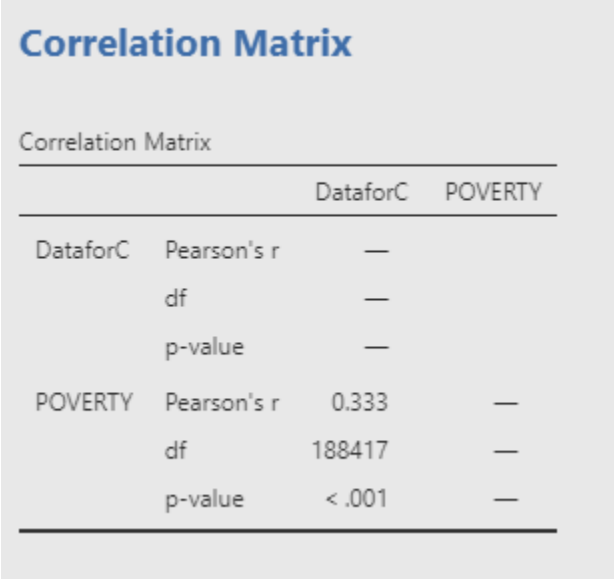

Correlation Matrix

Correlation Matrix

		DataforC	POVERTY
DataforC	Pearson's r	—	
	df	—	
	p-value	—	
POVERTY	Pearson's r	0.333	—
	df	188417	—
	p-value	< .001	—

Figure 3 Dementia Diagnoses Correlation to Poverty

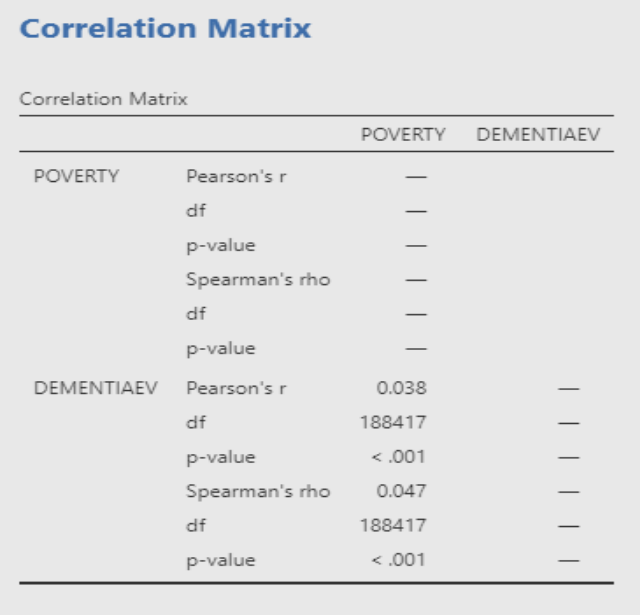

Correlation Matrix

Correlation Matrix

		POVERTY	DEMENTIAEV
POVERTY	Pearson's r	—	
	df	—	
	p-value	—	
	Spearman's rho	—	
	df	—	
	p-value	—	
DEMENTIAEV	Pearson's r	0.038	—
	df	188417	—
	p-value	< .001	—
	Spearman's rho	0.047	—
	df	188417	—
	p-value	< .001	—

Figure 4 Association Between Dementia and Food Insecurity

Correlation Matrix

Correlation Matrix

		DEMENTIAEV	DataforC
DEMENTIAEV	Pearson's r	—	
	df	—	
	p-value	—	
	Spearman's rho	—	
	df	—	
	p-value	—	
DataforC	Pearson's r	0.030 ***	—
	df	188417	—
	p-value	< .001	—
	Spearman's rho	0.049 ***	—
	df	188417	—
	p-value	< .001	—

In order to properly test for a link between Dementia and food insecurity, we extracted data from the IPUMS National Health Interview Survey with a date range from 2020 to 2024. These data points include FSSTATDET, DEMENTIAEV, AGE, SEX, POVERTY, and EDUCATION, as well as identifying variables such that individuals could be distributed to specific files. FSSTATDET is a survey issued to individuals that asks them questions about their food availability and security. DEMENTIAEV - DV for future reference - is a categorical variable that reports whether or not the specific individual have been diagnosed with dementia; AGE reports the subjects age, SEX is another categorical variable reporting an individual's sex; POVERTY is the variable that reports and individuals income in relation to the poverty line where 1.00 is at the poverty line whereas a value like 2.00 reports someone as being 2 times above the poverty line; and EDUCATION reports and individuals reported educational levels ranging from failing to graduate middle school to the MD or PhD level.

To test if there was a correlative relationship between Dementia and food insecurity in figure 1, we converted DV into XX, a binary variable where 0 represented no Dementia diagnosis, whereas a 1 represented an individual who had been diagnosed with dementia, and that transformation could be understood by the calculative application Jamovi, which would analyze the data downloaded and imported from IPUMS NHIS taking over 188,417 points from 2020-2024. Furthermore, the data for FSSTATDET was altered into YY, binary as well, where very low food security is translated into 1 to represent high rates of food insecurity, whereas 0 represents everything else, since any score above very low represents limited food or diversity, not true food insecurity. These data points were interred into the Jamovi Bar plot with mean AGE as the Y axis variable.

To attempt and measure correlative relationship between poverty and food insecurity in figure 2 with a correlation matrix, the author employed a transformation of FSSTATDET and

altered it into the variable DataforC. To do this, the author subtracted the value of FSSTATDET from its maximum value, and therefore linking with a lower score to having lower access to food, functionally as the system approached 0, the levels of food security would decrease. This way, as the value of DataforC grows smaller so does income allowing for an effective demonstration of correlation if one exists as it will work such that they would both increase or decrease together instead of working in an inverse pattern which would produce misleading data.

In figure 3, the correlation matrix is used to measure dementia's diagnoses rates against rate of poverty within the sample group measuring for predictability of the datum points based on the data of the other group, attempting to predict the outcomes of the other group as derived from their own outcomes. In the 4th figure, we attempt to gauge the correlation between food insecurity from DataforC and weigh it against DV in a correlative matrix to measure for datum's predictive values against the other variable within the matrix using Jamovi's computing capabilities, like used in the other figures and graphs employed within the paper.

Results

When the data is analyzed, we find that there is a correlative relationship between dementia risk and lack of access to food. Even when the data is taken with confounding variables taken into account, we still find that the relationship does exist. While the correlative relationship is weak, it is statistically significant while this may simply be driven by the size of the population measured. There is a measurable link between POVERTY and DV, POVERTY and DataforC, and DataforC and DV . The correlation matrix between POVERTY and DV had a Pearson's r of 0.038 showing that there is a relationship between the two variables. When this is performed between POVERTY and DataforC there is a Pearson's r of 0.333. When the relationship between DataforC and DV the Pearson's r value drops to 0.049. Yet, between figures 2, 3, and 6, the p value for each of these is less than 0.001. When measuring figure one, the data indicates that individuals who experience food insecurity are far more prone to development of dementia. The data shows that the mean age of individuals within the group to be 76 as opposed to the average age of 63, in the population that experiences food insecurity.

Discussion

The hypothesis that food insecurity has a correlative relationship with dementia, with this impact being amplified in an aging population, has statistically significant correlative data to support the claim. When Figure 6 is observed, the Pearson's r is 0.049, showing a weak but positive relationship between the two variables, the value is further supported by the p value being lesser than 0.001 meaning that the data is reliable and not occurring due to chance. Yet, the Spearman's rho is so small the correlation drawn with this number alone cannot be uniquely useful. Still, perhaps one of the greatest inhibitors in the outcomes of this method of determining if there is a correlation based off of a matrix is that anyone can be underweight or at risk of food insecurity meaning that while there are the target individuals in the data, there is also a huge excess of random and confounding variables. These issues are primarily addressed in analysis

like in figure one, where we directly measure mean ages against nothing but those that serve as analogues in the opposite groups. This means that even though the Pearson's R is nearly insignificant, it does still exist and when measuring for age, there is a substantial difference.

This evidence is also unlikely to be derived from a difference in lifespan brought about by socioeconomic status which is represented by the rates of food insecurity. When correcting food security and mortality rates, a study was conducted analyzing the correlation, tracking a group of participants for decades and any death before 80 years old was considered a premature death. In the study, the researchers found that at 50, even if individuals experience extreme food insecurity, they were set to live 28 more years as opposed to the 32.5 of a healthy person. (Ma et al, 2024). With a certainty level within 5% being so close to the values represented on the bar graph, it is far easier to see the 13 year gap between the two groups while even with this data taken into account, the difference in life expectancy would be about 4.5 years which undermines the idea that survivorship bias skews the numbers in a specific direction.

A probable partial cause for this large difference, as well, is the brain's autoregulation of blood flow shifts when an individual grows older, where the shape changes from a sigmoid curve to a linear line (Moretti et al., 2008). This furthers the hypothesis that food insecurity and dementia are linked with the possible mechanism of hypotension as it is caused by failure of cerebral autoregulation in older adults and is important in determining future risk of dementia. (Torre, 2012). However, the mechanism - while in accordance with previous studies - is not proven in this paper and instead it has evidence supporting the hypothesized link to dementia diagnosis.

An unexpected limitation of this study was the distribution of studies globally. All of the data gathered from the study was solely limited to the United States where we lack the ability to survey large groups that experience extreme food insecurity. This limits insight into the issue of malnutrition in the United Kingdom, South America, Africa, and Australasia. The World Alzheimer's Report in 2015 found that the highest prevalence of AD globally was in East Asia and Western Europe and determined that the distribution of new dementia cases was 49% in Asia (including Australia), 25% in Europe, and 8% in Africa (Prince, 2013). This demonstrates that we were severely limited in analysis of a significant majority of cases. Yet, this review not only struggles to provide the average prevalence of dementia and malnutrition in countries that are excluded from the study, but almost any study will struggle to be accurate in low- to middle-income countries. This is a multifactor problem driven not only by cultural stigma around diseases like dementia or AD, but also by a lack of education around it in many areas, leading to lower rates of self-reporting. This issue is further compounded by the test that is actually administered to determine if an individual has an implicit bias towards Western individuals due to cultural bias and oversights, as well as being designed under the premise that the individual is educated in a Western manner. (Mezzich et al., 99)

The theory that in many countries where malnutrition and hunger is a serious issue are experiencing higher rates of dementia than traditional data would suggest is further supported by the evidence that countries like Africa where the rates of dementia are underreported. Even if

the representation of data is increasing since it is still inaccurate ranging from reported rates of 2.3 to 20% although the World Health Organization found that there would be a 207% increase in dementia cases from 2017 to 2050 (Akinyemi, 2021). This provides an area for further investigation into how malnutrition could drive dementia in regions that were previously underrepresented with the new and extended evidence that individuals are more likely to experience dementia from malnutrition. Yet, this is why getting the data from specific countries that have higher levels of development and infrastructure to support self reporting are important because countries that are generally less developed are less accurate due to socially and economically disadvantaged individuals tend to fail in perceiving and report the presence of illness or health-deficits because an individual's assessment of their health is directly derived from their social experience (Subramanyam, 2008). A state like Japan - a country with the highest prevalence of dementia was apparent, with a prevalence of 24.8 per 1000 (AIHW, 2026). One study that was conducted in Japan found high rates of malnutrition at 75.6%, whilst only 24.4% were at risk or well-nourished (Kamo et al., 2017). These high results can also be observed in a 2022 cross-sectional study on malnutrition conducted in long-term care in Japan using the MUST and Global Leadership Initiative on Malnutrition, whereby 34.3% of those with dementia had severe malnutrition. This bolsters the idea that malnutrition can serve as a driving factor behind dementia as well, showing other large scale statistical analyzes. . A review found that the lowest prevalence of malnutrition was in France, as only 6.8% of long-term care residents were considered to be malnourished in a pilot group (Perry et al., 2023). This data is significant because it allows for further investigation into how we could better understand neurodegeneration within long-term care residents and how malnutrition can drive it in previously under investigated fields. (Bajeux et al., 2021) As dementia is most associated with the fastest-growing and aging population group, more studies are required globally to determine the prevalence of malnutrition in long-term care facilities and its impacts on worsening cognitive health to intervene with public health initiatives.

Conclusion

Our view of the correlative link between dementia and high rates of food insecurity and malnutrition is driven by malnutrition's impact on the circulation of blood flow to the brain. The study identifies how lack of access to food is a driver behind dementia risks with the proposed impact on hypotension and its impacts on neuronal death with white matter lesions with the process being put forth as the proposed mechanism. However, more research could be done to demonstrate the mechanical link between pure malnutrition and its driving impacts on hypotension and how to better measure rates of malnutrition in other countries to better formulate hypotheses from larger and previously excluded populations. Still, there is also legislative opportunity to better combat the issues analyzed in this paper. It is possible to pass legislation that would better target elderly individuals such that healthcare can better account for them receiving the proper nutrients (Qian et al., 2023). It is also possible to pass legislation making nutritious food more affordable and accessible (Narasimmaraj et al., 2026).

Works Cited

Abdulkhaliq, Altaf A., et al. "Amyloid-B and Tau in Alzheimer's Disease: Pathogenesis, Mechanisms, and Interplay." Cell Death & Disease, vol. 17, no. 1, Jan. 2026, https://doi.org/10.1038/s41419-025-08186-8. Accessed 8 Oct. 2026.

Akinyemi, Rufus O., et al. "Dementia in Africa: Current Evidence, Knowledge Gaps, and Future Directions." Alzheimer's & Dementia, vol. 18, no. 4, Sept. 2021, https://doi.org/10.1002/alz.12432.

Anwer, Mehwish, et al. "Vascular Contribution to Cognitive Impairment and Dementia (VCID): Proceedings of 2025 Workshop of the Jackson Laboratory." Mammalian Genome, vol. 37, no. 1, Mar. 2026, https://doi.org/10.1007/s00335-026-10210-x. Accessed 8 Oct. 2026.

Australian Institute of Health and Welfare. Dementia in Australia. AIHW, Australian Government, 30 July 2026, www.aihw.gov.au/reports/dementia/dementia-in-aus/contents/how-many-people-have-dementia/prevalence-of-dementia#prev_age_sex. Accessed 8 Oct. 2026.

Bajeux, Emma, et al. "Integrated Care for Older People in France in 2020: Findings, Challenges, and Prospects." International Journal of Integrated Care, vol. 21, no. 4, 2021, https://doi.org/10.5334/ijic.5643.

Cheng, Yingzhe, et al. "Hypotension With Neurovascular Changes and Cognitive Dysfunction: An Epidemiological, Pathobiological, and Treatment Review." Chinese Medical Journal, vol. 138, no. 4, May 2024, pp. 405–18, https://doi.org/10.1097/cm9.0000000000003103.

de la Torre, Jack C. "Cardiovascular Risk Factors Promote Brain Hypoperfusion Leading to Cognitive Decline and Dementia." Cardiovascular Psychiatry and Neurology, vol. 2012, Dec. 2012, pp. 1–15, https://doi.org/10.1155/2012/367516.

Dobrescu, Sandra Rydberg, et al. "Anorexia Nervosa: 30-Year Outcome." The British Journal of Psychiatry, vol. 216, no. 2, May 2019, pp. 97–104, https://doi.org/10.1192/bjp.2019.113. Accessed 8 Oct. 2026.

Emre, Kocyigit Suleyman, et al. "The Relationship Between Cognitive Frailty, Physical Frailty and Malnutrition in Turkish Older Adults." Nutrition, vol. 126, Oct. 2024, p. 112504, https://doi.org/10.1016/j.nut.2024.112504. Accessed 8 Oct. 2026.

Feng, Long, et al. "Malnutrition Is Positively Associated With Cognitive Decline in Centenarians and Oldest-Old Adults: A Cross-Sectional Study." eClinicalMedicine, vol. 47, May 2022, p. 101336, https://doi.org/10.1016/j.eclinm.2022.101336.

Gkiouras, Konstantinos, et al. "Malnutrition and Food Insecurity Might Pose a Double Burden for Older Adults." Nutrients, vol. 12, no. 8, Aug. 2020, p. 2407, https://doi.org/10.3390/nu12082407. Accessed 8 Oct. 2026.

Iqbal, Arshad, and Syed Jamal. "Essential Hypertension." National Library of Medicine, 20 July 2023, www.ncbi.nlm.nih.gov/books/NBK539859/. Accessed 8 Oct. 2026.

Kamo, Tomohiko, et al. "Coexisting Severe Frailty and Malnutrition Predict Mortality Among the Oldest Old in Nursing Homes: A 1-Year Prospective Study." Archives of Gerontology

and Geriatrics [Netherlands], vol. 70, 2017, pp. 99–104. PubMed, https://doi.org/10.1016/j.archger.2017.01.009. Accessed 8 Oct. 2026.

Kocyigit, S. E., et al. "Malnutrition and Malnutrition Risk Can Be Associated With Systolic Orthostatic Hypotension in Older Adults." The Journal of Nutrition, Health and Aging, vol. 22, no. 8, Oct. 2018, pp. 928–33, https://doi.org/10.1007/s12603-018-1032-6.

Li, Lian, et al. "Association Between Malnutrition and Cognitive Decline in Chinese Older Adults: The Mediating Role of Physical Frailty." Frontiers in Public Health, vol. 14, June 2026, https://doi.org/10.3389/fpubh.2026.1843902. Accessed 8 Oct. 2026.

Lin, Chu Sheng, et al. "Hospitalization and Associated Factors in People With Alzheimer's Disease Residing in a Long‐Term Care Facility in Southern Taiwan." Geriatrics & Gerontology International, vol. 17, no. S1, Apr. 2017, pp. 50–56, https://doi.org/10.1111/ggi.13033.

MEZZICH, JUAN E., et al. "The Place of Culture in DSM-IV." The Journal of Nervous & Mental Disease, vol. 187, no. 8, Aug. 1999, pp. 457–64, https://doi.org/10.1097/00005053-199908000-00001.

Moretti, Rita. "Risk Factors for Vascular Dementia: Hypotension as a Key Point." Vascular Health and Risk Management, Volume 4, Apr. 2008, pp. 395–402, https://doi.org/10.2147/vhrm.s2434.

Narasimmaraj, Prihatha R, et al. "Health Effects of Reductions in Food Assistance Benefits in the USA." The Lancet, vol. 407, no. 10529, Feb. 2026, pp. 671–72, https://doi.org/10.1016/s0140-6736(26)00036-x. Accessed 8 Oct. 2026.

Perry, Emma, et al. "Prevalence of Malnutrition in People With Dementia in Long-Term Care: A Systematic Review and Meta-Analysis." Nutrients, vol. 15, no. 13, June 2023, pp. 2927–27, pmc.ncbi.nlm.nih.gov/articles/PMC10343750/. Accessed 8 Oct. 2026.

Prince, Martin, et al. "The Global Prevalence of Dementia: A Systematic Review and Metaanalysis." Alzheimer's & Dementia, vol. 9, no. 1, Jan. 2013, pp. 63–75.e2.

Qian, Haobing, et al. "Food Insecurity, Memory, and Dementia Among US Adults Aged 50 Years and Older." JAMA Network Open, vol. 6, no. 11, Nov. 2023, p. e2344186, https://doi.org/10.1001/jamanetworkopen.2023.44186. Accessed 8 Oct. 2026.

Ringer, Matthew, et al. "Orthostatic Hypotension." Ncbibooks, StatPearls Publishing, 17 Jan. 2025, www.ncbi.nlm.nih.gov/books/NBK448192/?report=classic. Accessed 8 Oct. 2026.

Sarhan, Mutaz, et al. "The Pathophysiology of Mixed Alzheimer's Disease and Vascular Dementia." Theranostics, vol. 15, no. 18, Sept. 2025, pp. 9793–818, https://doi.org/10.7150/thno.118737. Accessed 8 Oct. 2026.

Wang, Ling-ying, et al. "Prevalence of Mild Cognitive Impairment and Its Association With Malnutrition in Older Chinese Adults in the Community." Frontiers in Public Health, vol. 12, Aug. 2024, https://doi.org/10.3389/fpubh.2024.1407694.

Wolters, Frank J., et al. "O2‐09‐03: Orthostatic Hypotension and the Long‐Term Risk of Dementia: A Population‐Based Study." Alzheimer's & Dementia, vol. 12, no. 7S_Part_5, July 2016, https://doi.org/10.1016/j.jalz.2016.06.444.

Young, Angela Odoms, et al. “Food Insecurity, Neighborhood Food Environment, and Health Disparities: State of the Science, Research Gaps and Opportunities.” The American Journal of Clinical Nutrition, vol. 119, no. 3, Dec. 2023, pp. 850–61, https://doi.org/10.1016/j.ajcnut.2023.12.019. Accessed 8 Oct. 2026.
Zhang, Meibiao, et al. “Cerebral Infarction: Epidemiology, Classification, Mechanisms, Diagnosis, and Management.” MedComm, vol. 7, no. 4, Apr. 2026, https://doi.org/10.1002/mco2.70704.

Hues of Attraction: Color Cohesion, Gendered Cues, and Scarcity Framing in Early-Stage Fashion Brand Engagement on Instagram By Juana Colaroff

Abstract

Fashion startups founded within the past one to two years face limited budgets, low conversion rates, and competition from larger, established retailers. Color is one tool these young companies may be underusing. While color psychology is well studied in traditional retail, little research examines how palette cohesion, saturation, and gender-coded colors relate to consumer engagement on platforms like Instagram. This exploratory study analyzes the Instagram accounts of eight early-stage fashion startups, coding each brand's color palette and recording its follower count, profile scarcity language, and like and comment engagement rates. Within this small sample, none of the three color features examined lined up consistently with engagement. Brands with more saturated palettes did not show higher engagement, and when palette cohesion was measured by counting hue families, seven of the eight brands, including the brand with the lowest engagement, had highly cohesive palettes, so cohesion did not distinguish stronger from weaker accounts. Gender-coded palettes were not consistently associated with comment engagement either. Only one of the eight brands used scarcity language in its profile. Because the sample is small, color coding relied on visual judgment, and no statistical testing was possible, these results should be read as preliminary. They suggest that no single, simple color rule explains engagement for early-stage brands, and that other factors, such as audience, posting behavior, and follower quality, may matter as much as color. Future research should use larger samples, software-based color measurement, and post-level data, and could explore AI-driven personalization of color presentation.

1. Introduction

1.1 Color and Early-Stage Fashion Brands

Fashion startups founded within the last one to two years face many challenges when competing against larger, more established retail brands. According to LendingTree's analysis of U.S. Bureau of Labor Statistics data, roughly one in five new U.S. businesses closes within its first year ("22.1% of New US Businesses").

One factor these brands may overlook is color. In a review of color research in marketing, Singh reports that people form judgments about a product within about 90 seconds of first seeing it, and that an estimated 62 to 90 percent of that assessment is based on color alone (783). Because this figure comes from a literature review rather than a single experiment, and is widely repeated in popular marketing writing, it is best read as an indication of color's importance rather than a precise measurement. Research on color and brand personality also suggests that specific hues carry associations; for example, red has been linked to perceptions of excitement and blue to perceptions of competence (Labrecque and Milne). For an early-stage brand, misjudging demand, including demand for particular colors and styles, can leave unsold inventory and cause

financial losses (Fisher et al.), a risk that is especially damaging for a company that is still building capital.

Color also shapes how consumers form impressions of a brand. Businesses deliberately select colors for logos, packaging, and marketing materials to build recognition and emotional attachment (Maghraby et al.). This matters more as fashion shopping moves online: clothing and apparel is the leading shopping category on Instagram, and consumers are reported to be about 12 percent more likely to buy clothing there than on any other social platform ("Instagram Shopping Statistics"). A brand's color strategy is no longer only a question of what looks good in a storefront window; it is also a question of what makes a person stop scrolling. Because this study measures Instagram engagement rather than sales, however, it can speak only to that first step, attention and interaction, and not to purchasing.

1.2 Theoretical Framework

This study draws primarily on the effect heuristic, which describes how people often rely on immediate emotional reactions rather than deliberate reasoning when forming judgments. Finucane et al. show that this heuristic allows people to make fast judgments based on how something makes them feel rather than through careful analysis of its features or benefits. The heuristic may be especially relevant on Instagram, where a user decides within moments whether to stop on a post or keep scrolling. If color produces a quick positive or negative reaction, that reaction could shape whether a user engages with a post before considerations such as price or fit come into play.

The effect heuristic was originally described in the context of decisions, and the most direct test would measure purchase intention or actual purchases. Those data are not publicly available for private startups. Instead, this study uses engagement rate as an indirect, observable trace of quick affective responses. A like requires a single tap and is usually given within seconds of seeing a post, which makes it the closest publicly visible record of an immediate reaction; a comment requires more effort and is treated as a sign of deeper involvement (Sabate et al.; Swani et al.). This proxy has clear limits: engagement does not always lead to purchases, and a like can reflect many things besides an emotional response to color. Engagement rate is therefore treated here as evidence of attention and interaction, not of buying behavior.

A second framework, the scarcity principle, holds that consumers place higher value on products they perceive as limited or hard to obtain. A meta-analysis of 131 studies found that scarcity cues generally increase purchase intentions, although the size of the effect depends on the type of cue and product (Barton et al.). A separate meta-analysis found that consumers are more likely to buy a scarce product when it involves social signaling and seasonality (Ladeira et al.), two features common in fashion. On Instagram, a brand can signal scarcity through a bio that mentions a "limited edition," a countdown in a Story, or a caption announcing an exclusive drop. Because scarcity language is not a color variable, combining it with the color analysis would make it impossible to tell which factor was associated with engagement. Scarcity is

therefore examined separately, as a descriptive question about whether these brands use the principle at all, and it is not used to explain the color results.

1.3 Research Questions and Hypotheses

This study compares the color strategies of eight fashion startups launched within the past one to two years. Its primary research question is: How is a brand's Instagram color strategy, meaning its palette cohesion, saturation, and use of gender-coded colors, associated with like and comment engagement rates for early-stage fashion brands? Because consumer color preference is shaped by many factors, including complexion, occasion, and gender (Adeboye et al.), no single color strategy is expected to work for every audience. The aim is to identify patterns that are worth testing in larger studies.

Because the data were collected at the brand level, as averages across each brand's ten most recent posts, the hypotheses are stated at the brand level:

- **Hypothesis 1.** Because strong visual cues capture attention during scrolling (Kanuri et al.), brands whose feeds use more highly saturated palettes will show higher like engagement rates than brands with low-saturation palettes.
- **Hypothesis 2.** Because comments signal deeper involvement (Sabate et al.; Swani et al.), brands whose palettes combine gender-coded color cues with strong visual alignment across the feed will show higher comment engagement rates.
- **Hypothesis 3.** Brands that maintain a cohesive color palette across their feeds will show higher engagement rates than brands with less cohesive palettes.

A secondary, descriptive question was also examined: Of the eight sampled brands, which ones use visible scarcity language in their Instagram profiles?

2. Literature Review

2.1 Color, Emotion, and Gendered Color Cues

Color can act as a fast and often unconscious signal. In a review of theoretical and empirical work, Elliot describes how perceiving color can influence emotional, cognitive, and behavioral responses, frequently without the person being aware of it. Elliot also stresses that these effects depend heavily on context, meaning the same color can produce different reactions in different situations.

In fashion retail, color is one of the first things a consumer notices about a garment, and a positive reaction to a palette may carry over to how the consumer evaluates the product itself. Research on apparel suggests that color is a meaningful factor in consumer buying behavior for clothing (Casas and Chinoperekweyi).

Clothing color preferences are also tied to personal and social characteristics. In a survey of 186 Nigerian undergraduates, Adeboye et al. found that complexion and occasion were the factors respondents ranked highest when choosing garment colors, and that gender also shaped

choices: respondents most associated blue with men's clothing and pink with women's, and saw black and white as worn by both. These associations come from one cultural context, and color-gender associations are not fixed; they can vary across cultures, generations, and audiences. Even so, they suggest that when early-stage brands choose a palette, they may be signaling, intentionally or not, who the clothing is for. In this study, a palette's gendered association is treated as one feature of its overall color strategy rather than as a separate explanatory variable.

2.2 Color Dynamics in Social Media Engagement

On Instagram, color is one of several visual features associated with how users respond to posts. In an analysis of tourism photographs on Instagram, Yu et al. found that users tended to respond more to brighter and more saturated images, and that certain hues were linked to higher post popularity. Because users scroll quickly, visual properties such as color can register before captions or brand names are read, and color research in marketing treats color as one of the most immediate elements of how marketing content is perceived (Labrecque et al.).

Researchers have examined several color dimensions, including lightness, saturation, and color complexity (Yu et al.; Kanuri et al.). Using Facebook posts from two organizations, Kanuri et al. found that images with greater color complexity, meaning more variation in color across pixels, tended to attract more attention and engagement. Color congruence has also been studied: in a laboratory experiment using Instagram brand posts, Luarn et al. found that matching colors between a post's main theme and its background improved viewers' aesthetic experience and, in turn, their brand attitudes and intentions. These findings point in somewhat different directions, one favoring visual complexity and the other visual consistency, which is part of why this study examines both saturation and cohesion.

Engagement metrics also differ in the level of involvement they reflect. Studies of brand posts on social media have analyzed likes and comments as separate measures of popularity, with likes requiring minimal effort and comments requiring more active participation (Sabate et al.; Swani et al.). This study follows that distinction, treating likes as a measure of low-effort, immediate interest and comments as a measure of more deliberate engagement.

2.3 The Digital Color Psychology Research Gap

Despite extensive research on color and consumer behavior, much of the foundational work was built around traditional contexts such as storefronts, packaging, and print advertising. Labrecque et al. note that while fields such as neuroscience and visual cognition have advanced the understanding of color perception, comparatively little research has situated color within marketing, and marketing has since moved heavily onto social platforms. More recently, Zutshi observed that research on how color choice influences engagement and user behavior on social media remains limited (1699). The Instagram color studies cited above focus on tourism content, laboratory experiments, or large organizations. Less is known about early-stage fashion brands, which have small audiences, limited budgets, and little brand recognition to rely on. This study

responds to that gap with an exploratory analysis of eight currently operating Instagram-based fashion startups, using the platform's own publicly visible engagement data.

3. Methodology

3.1 Research Design

This study uses an exploratory, descriptive design: it compares a small, purposively selected set of brands to identify patterns that could be tested in larger studies. It is not designed to test whether color causes differences in engagement, and no statistical significance testing was performed.

3.2 Sample Selection

The sample consists of eight fashion startups, each launched within the past one to two years and each using Instagram as its main channel for consumer engagement. The brands were purposively chosen to meet three criteria: (1) the brand must be an early-stage company; (2) the brand must maintain an active, publicly viewable Instagram account with a visually identifiable color strategy; and (3) the brand must have a clearly stated or inferable target demographic, since gender-coded color cues are part of Hypothesis 2.

Three of the brands were marketed toward women (LEONÍ Studio, By Viktorija, and The Trenza), three toward men (Kalpraag, CACHÍ, and Twenty Four), and two as unisex or gender-neutral (ABLANKPAGE and 12LUNES). This selection allowed color strategies to be compared within and across gendered market segments, since gender is one of the factors associated with garment color preference (Adeboye et al.). Figures 1 through 8, which document each brand's Instagram profile and coded color palette, are provided in Appendix A.

3.3 Data Collection

Data was collected in August 2026 through direct review of each brand's public Instagram account. For each of the eight brands, the brand's handle, target demographic, and dominant color palette were recorded. Each brand's profile bio was also reviewed for explicit scarcity or limited-availability language, such as references to a "drop," a "limited edition," or a restricted release.

Engagement data were recorded by reviewing each brand's ten most recent posts and recording the number of likes and comments on each. Because Instagram does not always display exact like counts publicly, some values are manual counts of visible engagement. These values were averaged for each brand to produce an observed average likes value ($\bar{x}L$) and an observed average comments value ($\bar{x}C$). Follower counts were recorded from each brand's profile at the time of coding. Because the brands differ greatly in size, engagement was normalized by follower count:

- Like engagement rate = ($\bar{x}L \div$ followers) $\times$ 100

- Comment engagement rate = ($\bar{x}$C ÷ followers) × 100

Likes were treated as a measure of low-effort, immediate interest and comments as a measure of higher-effort, more deliberate engagement (Sabate et al.; Swani et al.).

3.4 Color Coding Criteria

All color coding was done by visual inspection of each brand's ten most recent posts, following the color dimensions most often studied in social media engagement research: lightness, saturation, and color complexity (Yu et al.; Kanuri et al.). Three features were coded.

- **Saturation.** *Low*: most posts are dominated by neutrals (black, white, gray, beige), muted or desaturated tones, or very pale tints. *Medium*: most posts mix muted tones with clearly colored elements, or use moderately rich colors that are not vivid. *High*: most posts are dominated by vivid, intense hues. Brands whose posts fell between two categories received a combined label (for example, Low-to-Medium). A short descriptive label, such as "Minimalist Monochromatic," was added to summarize each palette.
- **Gendered association.** Where a palette carried a clear gendered association, such as soft pinks read as feminine, that judgment was informed by the color-gender associations reported by Adeboye et al. and is treated with caution, since such associations vary across cultures.
- **Palette cohesion.** Cohesion was defined as the number of distinct hue families among each brand's coded dominant colors, shown as the palette swatches in Figures 1 through 8: *High* = one or two hue families; *Moderate* = three; *Low* = four or more. The families used were neutrals (black, white, gray, cream, beige, oat, and dark espresso brown), blues, greens, reds and pinks, oranges and yellows (including ochre, terracotta, and gold), and purples. Cohesion was not part of the original coding plan. It was defined and coded after the engagement data had been recorded, so Hypothesis 3 should be read as a post hoc, exploratory test rather than a pre-planned one.

All coding was done by a single researcher. The limitations of this approach are discussed in Section 6.

4. Results

4.1 Engagement by Color Strategy

Table 1 summarizes the color strategy and observed engagement data for each of the eight brands. Follower counts range from 1,958 to 152,000, so raw average likes and comments are not directly comparable: a large account will naturally accumulate more likes than a small one regardless of its color strategy. Table 1 therefore reports each brand's like and comment engagement rates alongside the raw averages. The Instagram profile screenshots from which each classification was drawn are provided in Appendix A.

Table 1. Color strategy, palette cohesion, follower count, and engagement rate by brand

Brand	Followers	Saturation & Palette Type	Palette Cohesion	Avg. Likes ($\bar{x}$L)	Like Engagement Rate	Avg. Comments ($\bar{x}$C)	Comment Engagement Rate
By Viktorija	5,377	Low-to-Medium (Light Romantic Cues)	High (2)	457	8.50%	12	0.223%
CACHÍ	3,807	Low-to-Medium	High (2)	66	1.73%	3	0.079%
Twenty Four	152,000	Low (Cool Dark Monochromatic / High Contrast)	High (1)	2,226	1.46%	24	0.016%
12LUNES	1,958	High (Vibrant Streetwear)	Moderate (3)	25	1.28%	5	0.255%
The Trenza	12,200	Medium (Warm Palette)	High (2)	60	0.49%	1	0.008%
LEONÍ Studio	11,000	Low (Minimalist Monochromatic)	High (1)	29	0.26%	3	0.027%
ABLANKPAGE	10,600	Low (High Contrast)	High (1)	27	0.25%	0	0.00%
Kalpraag	75,800	High (Rich Jewel Tones)	High (2)	17	0.02%	0	0.00%

Note. Brands are ranked by engagement rate, high to low. Saturation, palette type, and cohesion are based on visual inspection rather than a colorimetric tool; the number in parentheses is the count of hue families in the brand's coded palette (see Section 3.4 and Appendix A). Average likes and comments are calculated across each brand's ten most recent posts; engagement rates are those averages divided by follower count. Follower counts were recorded in August 2026.

Normalizing by follower count substantially changed the picture. Twenty Four's 2,226 average likes rank highest in absolute terms, but relative to its 152,000 followers, that translates to a like engagement rate of 1.46%, third in the sample. By Viktorija, with 5,377 followers, has a like engagement rate nearly six times higher (8.50%). Kalpraag's 17 average likes look modest in raw form, and scaled against its 75,800 followers, its 0.02% like engagement rate is the lowest in the dataset, even though Kalpraag uses the most vividly saturated palette in the sample (emerald green, royal blue, and midnight). A similar reversal appears in the comment data: 12LUNES, not By Viktorija or Twenty Four, has the highest comment engagement rate (0.255%), narrowly ahead of By Viktorija (0.223%), while Twenty Four's relatively large raw comment count (24) corresponds to one of the lowest comment engagement rates (0.016%).

Engagement rate has a weakness of its own: it assumes that every follower is a real, active account. Very low engagement on a large following, as with Kalpraag, can signal inactive or purchased followers, which would lower engagement rate regardless of color. This possibility cannot be ruled out with publicly available data and is an alternative explanation for Kalpraag's result.

4.2 Use of Scarcity Language

Of the eight brands, only one, Twenty Four, includes explicit limited-drop language in its profile bio. The other seven brands do not reference scarcity anywhere in their profiles. Twenty Four also has the largest follower count in the sample, but with only one brand using scarcity language, this study cannot assess whether scarcity messaging is associated with reach or engagement. What can be observed is that seven of the eight brands are not visibly applying a principle that the literature links to higher purchase intention (Barton et al.; Ladeira et al.), which may represent an unused opportunity.

5. Discussion

5.1 Hypothesis 1: Saturation

Hypothesis 1 predicted that brands with more saturated palettes would show higher engagement rates. The data from this sample do not support this prediction. The two highest like engagement rates, By Viktorija (8.50%) and CACHÍ (1.73%), came from Low-to-Medium saturation palettes, while the most saturated palette in the study, Kalpraag's, was associated with the lowest like engagement rate (0.02%). The other High-saturation brand, 12LUNES, fell in the middle (1.28%). Within these eight brands, saturation alone does not appear to explain differences in engagement. Eight brands are not enough to conclude that saturation is unimportant more generally, and this result sits alongside studies that have found brighter, more saturated, or more colorful images associated with higher engagement (Yu et al.; Kanuri et al.). One possible explanation is that those studies compared individual images, while this study compared brand-level averages, which may hide differences between individual posts.

5.2 Hypothesis 2: Gendered Cues and Visual Alignment

Hypothesis 2 predicted that brands combining gender-coded color cues with strong visual alignment would show higher comment engagement rates. The comment data do not show a consistent pattern. Viktorija's soft, pastel "Light Romantic Cues" palette, consistent with the association between pink and women's clothing reported by Adeboye et al., had the second-highest comment engagement rate (0.223%). However, the highest comment engagement rate (0.255%) belonged to 12LUNES, a unisex brand whose palette of cobalt blue, ultraviolet, graphite, and charcoal is neither gender-coded nor highly cohesive; it is the only brand in the sample coded as Moderate cohesion. Meanwhile, Kalpraag and ABLANKPAGE, whose palettes carry clear masculine or neutral associations and were both coded as High cohesion, recorded 0.00% comment engagement rates. Within this sample, then, neither gender coding nor palette alignment tracked comment engagement in a consistent way, so Hypothesis 2 is not supported. These comparisons also rest on very small numbers, between 0 and 24 average comments per post, so a few comments more or less would change the rankings.

5.3 Hypothesis 3: Palette Cohesion

Hypothesis 3 predicted that brands with more cohesive palettes would show higher engagement rates. Cohesion was defined and coded after data collection (see Section 3.4), so this is an exploratory test. Under the hue-family rule, seven of the eight brands were coded as High cohesion and one, 12LUNES, as Moderate. The High-cohesion group includes both the brand with the highest like engagement rate (By Viktorija, 8.50%) and the brand with the lowest (Kalpraag, 0.02%), whose emerald green, royal blue, and midnight palette falls into only two hue families. The single Moderate-cohesion brand, 12LUNES, had the highest comment engagement rate. Hypothesis 3 is therefore not supported in this sample: palette cohesion, as measured here, did not distinguish stronger from weaker accounts.

This result may say as much about the measure as about cohesion itself. Most early-stage brands in the sample appear to use small, controlled palettes, so a simple count of hue families leaves little variation to compare. A finer measure, such as software-based color variance across posts, might separate brands that a family count treats as equally cohesive. It is also possible that, once a brand has a reasonably consistent palette, other factors, including its audience, posting frequency, paid promotion, follower quality, and the products themselves, explain more of the difference in engagement than color does.

5.4 Scarcity

Separately from the color findings, the scarcity results show that most of these startups are not using one of the low-cost tools discussed in Section 1.2. Because only one brand used scarcity language, this study cannot say whether scarcity would help these brands, only that almost none of them are currently using it.

5.5 Practical Implications

For emerging fashion brands with limited capital, the main practical lesson of this exploratory study is caution. Neither vivid color nor a tightly limited palette was, on its own, associated with higher engagement among these eight brands, so founders should not assume that following a single color rule will lift their results. Instead, brands could treat color choices as something to test against their own data, for example by tracking engagement rate before and after a change in palette. The scarcity findings point to a separate opportunity that this study did not test: because scarcity cues appear to be more persuasive when products involve seasonality and social signaling (Ladeira et al.), limited seasonal color releases could be one practical way for a fashion startup to apply the scarcity principle. Whether such releases increase engagement or sales for early-stage brands would need to be tested.

6. Limitations

6.1 Sample Size and Level of Analysis

The most significant limitation is sample size. With only eight brands, the dataset cannot support statistical significance testing, so the results should not be read as generalizable.

Viktorija accounts for much of the variation in like engagement rate, which means the observed patterns are heavily influenced by a single brand. Data was also collected at the brand level, averaging ten posts per brand, so this study cannot examine how color in an individual post relates to that post's engagement. Future research should expand the sample and collect post-level data to allow statistical testing of the relationships proposed.

6.2 Subjectivity in Color Coding

Classifying a palette as Low, Medium, or High saturation, and counting its hue families, relies on visual judgment, so different coders could classify the same palette differently. All coding in this study was done by one researcher, and cohesion was defined after the engagement data were recorded, which increases the risk of seeing an expected pattern. The hue-family count used for cohesion is also coarse: it placed seven of the eight brands in the same category, which limits what it can reveal. A second, independent coder, or software that reports hue, saturation, and brightness values directly from image data, would make the classifications more reliable and more precise. The gendered color associations used in coding also come from a single survey in one cultural context (Adeboye et al.) and may not apply to the international audiences of these brands. Because some like counts were counted manually, minor measurement error is also possible.

6.3 Engagement as a Measure

Engagement rate is an indirect measure. It does not capture purchase intention or sales, so these findings cannot show whether color strategy affects buying. Engagement rate also assumes followers are real, active accounts; inactive or purchased followers would lower a brand's rate regardless of its color strategy. Finally, normalizing by follower count addresses the most obvious difference in scale between accounts, but it does not account for other variables that plausibly affect engagement, such as posting frequency, paid promotion, caption quality, and how long each account has been active. Future research should collect and control for these variables to isolate any relationship between color and engagement.

7. Future Research

Future studies could build on this exploratory work in several ways. A larger sample of early-stage brands, with color measured from image data using software and engagement recorded at the post level, would allow quantitative testing such as multiple regression, including checks for interactions between saturation, cohesion, and gendered cues while controlling for posting frequency, paid promotion, and account age. Qualitative methods, such as a thematic analysis of the comments brands receive, could show whether users actually respond to color and identity in the way this study assumed; Zutshi's combination of survey and interview data offers one model for such a mixed-methods design. Studies that measure purchase intention or sales would test the link between engagement and buying that this study could not.

The intersection of color psychology and artificial intelligence is another promising direction. As AI tools become more capable of analyzing individual user behavior, engagement history, and inferred aesthetic preference, early-stage fashion brands may be able to personalize color presentation, tailoring which product images, palettes, or seasonal drops are shown to which viewers based on predicted color preference. This would extend the effect heuristic described by Finucane et al. from a brand-level strategy to an individually targeted one, and could allow startups to approximate, at a small scale, the personalized targeting that larger retailers achieve through far greater data and advertising resources.

8. Conclusion

This exploratory study examined how color strategy relates to Instagram engagement for eight early-stage fashion brands. The results did not match the original hypotheses. In this sample, color saturation did not track with like engagement, so Hypothesis 1 was not supported. Gender-coded palettes were not consistently associated with comment engagement, so Hypothesis 2 was not supported. When palette cohesion was measured by counting hue families, seven of the eight brands, including both the strongest and the weakest performers, had highly cohesive palettes, so Hypothesis 3 was not supported either. Separately, only one of the eight brands used scarcity language in its profile, suggesting that most of these startups are not applying the scarcity principle at all.

These results are still informative. They suggest that the color advice often given to brands, such as using bold colors or keeping a tight palette, is not enough on its own to explain why some young brands engage their audiences far better than others. The gap between By Viktorija and Kalpraag is striking, but both use cohesive palettes, so the difference likely reflects factors such as audience, posting behavior, or follower quality as well as, or instead of, color. With only eight brands, visual coding by a single researcher, and no statistical testing, these findings are a starting point rather than a settled answer. Larger studies that measure color precisely, collect post-level data, and control for other influences on engagement are needed to show whether, and how, color strategy matters for early-stage fashion brands on Instagram.

Works Cited

ABLANKPAGE [@ablankpage._official]. Instagram, www.instagram.com/ablankpage._official/. Accessed Aug. 2026.

Adeboye, Adebiyi Oladipupo, et al. "Factors Responsible for Influencing Colour Preference in Garment Selection among Undergraduates." East African Journal of Arts and Social Sciences, vol. 6, no. 1, 2023, pp. 31–38, https://doi.org/10.37284/eajass.6.1.1057.

Barton, Belinda, et al. "Scarcity Tactics in Marketing: A Meta-Analysis of Product Scarcity Effects on Consumer Purchase Intentions." Journal of Retailing, vol. 98, no. 4, 2022, pp. 741–58.

By Viktorija [@by.viktorija]. Instagram, www.instagram.com/by.viktorija/. Accessed Aug. 2026.

CACHÍ [@cachiofficielle]. Instagram, www.instagram.com/cachiofficielle/. Accessed Aug. 2026.

Casas, M. C., and J. Chinoperekweyi. "Color Psychology and Its Influence on Consumer Buying Behavior: A Case of Apparel Products." Saudi Journal of Business and Management Studies, vol. 4, no. 5, 2019, pp. 441–56.

Elliot, Andrew J. "Color and Psychological Functioning: A Review of Theoretical and Empirical Work." Frontiers in Psychology, vol. 6, 2015, article 368, https://doi.org/10.3389/fpsyg.2015.00368.

Finucane, Melissa L., et al. "The Affect Heuristic in Judgments of Risks and Benefits." Journal of Behavioral Decision Making, vol. 13, no. 1, 2000, pp. 1–17.

Fisher, Marshall L., et al. "Making Supply Meet Demand in an Uncertain World." Harvard Business Review, vol. 72, no. 3, 1994, pp. 83–93.

"Instagram Shopping Statistics [2026]: A Complete, Data-Driven Look at Social Commerce." Dataopedia, 11 Feb. 2026, dataopedia.com/instagram-shopping-statistics/. Accessed Aug. 2026.

Kalpraag [@kalpraag]. Instagram, www.instagram.com/kalpraag/. Accessed Aug. 2026.

Kanuri, Vamsi K., et al. "Standing Out from the Crowd: When and Why Color Complexity in Social Media Images Increases User Engagement." International Journal of Research in Marketing, vol. 41, no. 2, 2024, pp. 174–93.

Labrecque, Lauren I., and George R. Milne. "Exciting Red and Competent Blue: The Importance of Color in Marketing." Journal of the Academy of Marketing Science, vol. 40, no. 5, 2012, pp. 711–27.

Labrecque, Lauren I., et al. "The Marketers' Prismatic Palette: A Review of Color Research and Future Directions." Psychology & Marketing, vol. 30, no. 2, 2013, pp. 187–202.

Ladeira, Wagner Junior, et al. "A Meta-Analysis on the Effects of Product Scarcity." Psychology & Marketing, vol. 40, no. 7, 2023, pp. 1267–79.

LEONÍ Studio [@leoni_official]. Instagram, www.instagram.com/leoni_official/. Accessed Aug. 2026.

Luarn, Pin, et al. "Color Congruence on Instagram: The Role of Flow and Aesthetic Experience." Asia Pacific Journal of Marketing and Logistics, vol. 36, no. 10, 2024, pp. 2448–63.

Maghraby, Toqa M., et al. "The Psychology of Color and Its Effect on Branding." Journal of Textiles, Coloration and Polymer Science, vol. 21, no. 2, 2024, pp. 355–63.

Sabate, Ferran, et al. "Factors Influencing Popularity of Branded Content in Facebook Fan Pages." European Management Journal, vol. 32, no. 6, 2014, pp. 1001–11.
Singh, Satyendra. "Impact of Color on Marketing." Management Decision, vol. 44, no. 6, 2006, pp. 783–89.
Swani, Kunal, et al. "What Messages to Post? Evaluating the Popularity of Social Media Communications in Business versus Consumer Markets." Industrial Marketing Management, vol. 62, 2017, pp. 77–87.
The Trenza [@thetrenza]. Instagram, www.instagram.com/thetrenza/. Accessed Aug. 2026.
12LUNES [@12lunes]. Instagram, www.instagram.com/12lunes/. Accessed Aug. 2026.
Twenty Four [@tf24.twentyfour]. Instagram, www.instagram.com/tf24.twentyfour/. Accessed Aug. 2026.
"22.1% of New US Businesses Close Within a Year." LendingTree, 2025, www.lendingtree.com/business/small/failure-rate/. Accessed Aug. 2026.
Yu, C.-E., et al. "Coloring the Destination: The Role of Color Psychology on Instagram." Tourism Management, vol. 80, 2020, article 104110, https://doi.org/10.1016/j.tourman.2020.104110.
Zutshi, A. "Color, Emotion, and Perception: A Mixed-Methods Study on Social Media Design and Youth Well-Being." International Journal of Indian Psychology, vol. 13, no. 3, 2025, pp. 1698–1720, https://doi.org/10.25215/1303.157.

Appendix A: Instagram Profile Screenshots

Figures 1 through 8 show each brand's Instagram profile at the time of coding, alongside the color palette extracted from it.

Evaluating the Impact of Prompt Engineering Strategies on Large Language Model Financial Forecasting

By Hugo Shek (student)[1], Plínio Zanini (mentor)[2], Siddharth Krishnan (mentor)[2]

Abstract

Prompt engineering influences how large language models (LLMs) reason when producing financial forecasts. Leveraging recent advancements in the Artificial Intelligence (AI) industry and the growing ability to shape LLM performance through prompt engineering, this study empirically examines the different ways in which prompting strategies improve LLM generated stock forecast responses. Our case-study was conducted as a series of controlled experiments, in which twelve prompting techniques and one composite strategy of our own design (Fin-Pro-1) were applied to identical financial inputs for selected publicly traded companies. The results suggest that a combination of in-context examples, clear reasoning structure, customized persona constraints, along with reflective revision was associated with higher interpretability of LLM produced stock forecasts. Fin-Pro-1 also achieved lower MAPE and RMSE than zero-shot prompting across all three forecasting horizons (1-week, 4-week, and 2-month) under the evaluated conditions. This work highlights that prompt design is an important component in LLM-based financial analysis and motivates future research using larger samples and quantitative validation. All the forecasts generated in this study are used for research evaluation and should not be interpreted as investment advice.

Keywords

Prompt Engineering, Large Language Models (LLMs), Financial Forecasting, Stock Price Prediction, Forecast Interpretability

Section 1: Introduction

Artificial Intelligence (AI) algorithms are increasingly applied in today's financial institutions because of their capacity to identify data patterns within these numerous and complex datasets. In these high-stakes settings, DBS Bank reported a revenue uplift of more than $150 million from AI-supported operations and an additional $30 million in cost avoidance through AI enhanced monitoring, fraud prevention, and operational systems in 2022 alone [1]. As investors, businesses, and financial institutions rely on accurate forecasting daily, to mitigate risk and make informed financial decisions to strengthen their portfolios[2], the reliability of LLM reasoning becomes increasingly consequential. The task of being precise, however, is greatly challenging in the current financial market as it operates under high volatility and requires rapid analysis of multimodal data including market shifts, changes in monetary policy, economic status, and indicators of consumer sentiment [3].

Traditional forecasting approaches, including Random Forest, Autoregressive Integrated

Mills High School, 400 Murchison Drive, Millbrae, California, 94030, USA; thathugoboss@gmail.com

Moving Average (ARIMA), and Long Short-Term Memory (LSTM) networks, have been widely adopted in identifying historical patterns and predicting future movements. While these methods are effective in analyzing large amounts of quantitative data[4], they may struggle to account for shifts in investor sentiment and can be limited by the patterns and features present in their training data. AI models offer an advantage in fluctuating market environments by enabling adaptive analysis and supporting financial decision-making [2]. Amongst earlier AI driven methodologies, agent-based modeling (ABM) serves as an example of simulating the volatile economic environment through representing participants as "agents" that follow predefined decision rules[5]. While ABM can be useful for studying market dynamics, its heavy reliance on assumptions and parameter choices limits its ability to facilitate the development of a commercial-grade LLM agent with strong reasoning abilities[3,5].

Given these limitations, the reliability of LLM generated forecasts may depend not primarily on information provided, but rather how users instruct them. One practical approach for structuring LLM behavior is using prompt engineering: the design of input instructions that govern a model's response structure and evaluation process[6,7]. In financial applications, prompt design may impact whether a model produces a brief, generic response or an interpretable forecast that explains key strengths and weaknesses[8]. Recent work has proposed framework based prompting strategies including Chain-of-Thought, Graph-of-Thought, and similar architectures to improve the organization and clarity of model outputs[6,9]. Despite the growing momentum in financial forecasting, existing studies often lack comparative analysis on reasoning processes and predictive consistency under constrained experimental evaluation conditions[6].

Accordingly, this study examines the influence of different prompt engineering strategies on the structure and interpretability of LLM generated stock forecasting responses in case-study conditions. In these experiments, multiple prompting techniques were evaluated using consistent market data, and the resulting model outputs were analyzed using both quantitative forecasting metrics and qualitative assessment of response structure. This paper makes three main contributions. First, it provides a practical comparison of twelve established prompting strategies, examining not only their influences on financial forecasting accuracy but also a comparative evaluation on response interpretability under consistent conditions, an aspect often overlooked in previous studies. Second, it introduces and evaluates Fin-Pro-1, a composite framework combining multiple techniques to produce more structured financial analyses. The results suggest improvements over zero-shot prompting across all three forecasting horizons within the selected case studies, although these advantages were not consistent across all prompting strategies. Third, this study highlights an important distinction between forecast accuracy and interpretability, revealing that more structured reasoning does not necessarily lead to more robust forecasts. These findings provide insights into the potentials and limitations of prompt engineering in financial forecasting, while setting direction for further research through larger datasets, additional AI models, and more rigorous quantitative validation.

The remainder of the paper is organized as follows. Section 2 is a review of prompting strategies and their application towards the financial sector. In Section 3, we describe the

experimental design and both the quantitative and qualitative evaluation approaches. Section 4 presents our empirical results and analysis. Section 5 provides the limitations and future directions, and Section 6 concludes the paper.

Software disclosure: All tables and figures were created in Overleaf using LaTeX, except the MAPE and RMSE graphs, which were generated in Google Colab using Python.

Section 2: Prompt Engineering in Financial Forecasting

In recent years, prompt engineering, defined as the design of input instructions which guide large language model (LLM) behavior, has emerged as an important approach for shaping the quality of LLM responses[10]. A prompt's design contributes to the way in which an LLM presents its reasoning through a task and the extent to which it outputs information-dense responses. Prior research has introduced a range of prompting approaches, including few-shot learning[11], chain-of-thought prompting[12], and emotional prompting[13], which engineer the input given to an LLM without changing any model architecture. These prompting techniques intend to direct the model's assessment process in order that models produce more functional responses that are more structured for complex decisions that rely on consistent investment analysis.

This section reviews the progression of prompt strategies from zero-shot prompting to more recent prompting taxonomies. Figure 1 provides an overview of the prompting techniques examined in this study, grouped into four subcategories respective of their unique functionalities.

Figure 1. The web diagram illustrates the taxonomy of prompt engineering techniques reviewed in this study. The diagram lays out the larger prompt engineering methods into four groups including Example Learning, Self Reflection, Style & Persona, and Reasoning Structures. Each category then segments into the specific prompt techniques which are based upon the general prompting mechanism.

Zero-Shot Prompting

Zero-shot prompting[10] refers to prompts that are given to a language model without any examples they can point to. In common use, zero-shot prompts are the standard inputs without additional engineering [10] that often yield responses that lack depth and explanation. In an investment setting, an example for a standard prompt may be, "Project the documented cash flow of NVDA." In such cases, the model typically produces a response based primarily on its

pre-trained knowledge without explaining what led to the conclusion and whether its reasoning was based upon real statistical inferences. Standard prompting represents only one form of zero-shot prompting approaches, and several variations exist that introduce minimal structural guidance without adding examples. One approach is the "Take a Deep Breath" prompt[14], which allows space for the model to pause and synthesize all the relevant information before responding. As prompt structuralization strengthens, model responses often become more consistent, and in some cases, structured zero-shot prompts have been shown to outperform few-shot approaches[15]. However, zero-shot prompting has limitations. Because the method relies heavily upon its pre-trained data and lacks contextual awareness relative to its task, it is more vulnerable to hallucinations where the model generates entirely inaccurate and unsupported responses[3]. In this empirical project, zero-shot prompting serves as the baseline condition for the evaluation of LLM generated financial forecasts. By comparing zero-shot outputs to the more structured prompting strategies, the study can examine whether the additional prompt design improves reasoning depth, statistical consistency, and the reduction of unsupported assumptions.

Few-Shot Prompting

Few-shot prompting[11] builds on zero-shot prompting by embedding one or more examples within the prompt. For investment analysis, these examples can guide the model to a structural reasoning pattern, such as evaluating risk and historical trends in order to strengthen the validity of outputs. For example, a few-shot prompt may consist of a question and answer pair explaining a value investor's trading strategy specifying the length, structure, and evaluation criteria before asking the model to walk through another trading approach. In larger investment settings, exemplars allow LLMs to leverage in-context learning that support consistent responses across a high volume of organizational tasks and decisions[11]. Since investment queries tend to involve multi-step judgment rather than one direct answer, few-shot prompting can succeed by making its output more consistent amongst different companies, strategies, or market conditions. This context-rich input is then able to help the language model apply its instruction-following and pattern recognition capabilities within provided examples in the prompt. Even when examples are provided, few-shot prompting does not guarantee improved performance, since weak or irrelevant exemplars may fail to guide the model effectively. Therefore, the impact of few-shot prompting heavily depends on example quality as appropriate exemplars can elicit more systematic and explainable outputs[7]. Prior work on in-context learning suggests that model performance can improve as the number of examples increases, especially when these models are better at pattern identification from demonstrations[11]. However, in an investment analysis setting, the value produced by few-shot prompting should not be understood as a general performance gain, but rather as a way of making the model's financial reasoning more explicit. As shown in Figure 2, few-shot exemplars can help the model answer an investment query by providing the reusable architectural breakdown, in which the model can apply financial evidence, valuation logic, risk-return tradeoffs, market context, to form a final statement[8].

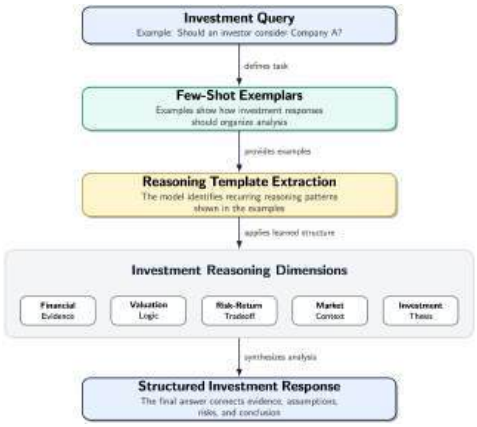

Figure 2. The flowchart illustrates how the few-shot prompting method is applied. The process begins with the investment query, followed by few-shot exemplars that demonstrate how the financial analysis structure is to be applied. From these exemplars, the model is able to identify reasoning patterns present, apply them towards the investment specific dimensions to focus upon before producing the investment analysis. Rather than performance measurement, this figure shows the proposed structure to help support more organized and explainable investment reasoning.

However, this relationship should not be generalized to every environment, as performance gains remain sensitive towards the task being addressed. In this empirical project, few-shot prompting is used as a comparison condition to evaluate whether added examples improve forecast consistency, reasoning structure, and evidence use compared with zero-shot prompting.

Personality Prompting

LLMs do not think like humans, but they respond differently when assigned to a specific persona or role, each with its appropriate expertise. Personality prompting, which is also referred to as strategy-specific prompting, instructs the models to integrate the role defined perspective when formulating responses[16]. Role-based prompting can steer the model towards greater response precision and reasoning depth by directing the model to consider different viewpoints in which a generic prompt may overlook. In financial forecasting tasks, this workflow can lead LLMs toward specific analytic properties, such as volatility, liquidity, trading strategy, or sector conditions such as technology, energy, and consumer markets[16]. Thus by specifying the scope, the prompt can help reduce irrelevant parameters which weaken the forecast. For example, a prompt may instruct the model to respond as a professional stock analyst focused upon the health and Medicare sector, while incorporating momentum based strategies and a medium risk assessment[16]. Rather than allowing the model to generate a broad market summary, persona-driven constraints force the output to remain within the intended analytical framework. For this empirical project, this makes it easier to examine whether role-based prompting improves forecast relevance, reduces unsupported assumptions, and limits hallucinated or off-sector reasoning. However, this strategy also has its limitations. When role definitions are too

vague or mixed with surrounding prompt context, a vague or contradictory role definition may weaken the model's reasoning and lead to less reliable forecasts[7].

Emotional Prompting

Emotional prompting involves integrating emotional cues into prompts such as an addition of politeness, compliments, or expressions of urgency in order to shape model responses[13]. The wording of a prompt does more than just request information as it has the capability to shape the way a model reasons. Emotional framing can affect how strongly the model elaborates on its forecast, the extent to which it weighs uncertainty, supports its assumptions, and justifies its conclusion. EmotionPrompt [13] formalizes the use of emotional cues within its prompts and reports improvements in several performance related measures. These results suggest that emotional prompting may help with answer clarity and transparency when compared to standard prompting in controlled evaluations[7]. In investment forecasting, this is especially relevant because the use of emotional prompting may encourage the model to take larger precautions as it deals with high-stakes decision making[3]. While these emotional cues can support more comprehensive responses, each one has its own restrictions. When emotional cues are overused, this may cause the response reliability to decrease and could even introduce unintended bias to the model[13]. Therefore, especially in financial forecasting, cautious consideration should be taken as biased reasoning can misrepresent investment interpretation[3]. In this evaluation, emotional prompting is not only examined by its effect upon explanation depth, but additionally for whether it helps preserve forecast reliability.

Self-Reflection Prompting: Take a Deep Breath & Take a Step Back Prompting

Large language models are trained to generate responses based on their learned patterns from large datasets. As a result, initial outputs may depend heavily on these learned patterns and produce responses with limited depth. However, prompting techniques that support revision help models reason more thoroughly and assess financial indicators with greater care[7]. The Take-a-Deep-Breath[14] and Take-a-Step-Back[17] prompting strategies were developed to encourage models to reason more carefully about the problem before generating their responses. These prompts instruct the model to reassess the problem before answering, hoping to minimize likelihoods of biases and unsupported model outputs. Through the incorporation of a reflective step, these methods aim to redirect the model to develop a more thoughtful analytical judgment instead of relying solely on the relevant pattern conclusions from its training data. By adding a reflective step before the final output, these methods can foster attributes such as evaluating upside potential, downside risk, market conditions, to support a more complete evaluation of the query[8]. This approach is especially relevant in financial forecasting, where reliable predictions require careful review of price movements, company fundamentals, sector trends, and the broader economic context[3]. In this empirical project, this approach is tested by examining whether prompts that suggest the model to take time to reason are more accurate and better justified than standard prompts.

Reasoning Structure Prompting

Reasoning structure prompting guides large language models to break down their responses using logical steps in an effort to improve response reliability. In financial forecasting, reasoning structure helps separate evidence, assumptions, and final predictions making the model's reasoning easier to evaluate[6]. This approach is useful in complex tasks including financial analysis, portfolio construction, risk modeling, all of which are highly dependent on clear reasoning and transparent explanations[8].

Chain-of-Thought (CoT) Prompting

Chain-of-Thought (CoT) prompting[12] was introduced in order to improve the LLM reasoning process by providing models with a way to construct out their reasoning steps rather than simply providing a final answer. CoT can also be prompted using short instructions such as "let's think step by step," which guides models into producing a structured solution outline, while including the final answer. In financial forecasting, this structure allows models to break down the finance problem into smaller pieces such as the various market factors and company trends, while producing the finalized thesis. Kojima et al. reported that adding the instruction "Let's think step by step" substantially improved GPT-3 performance on the MultiArith benchmark[18], increasing accuracy from 17.7% to 78.7% in their evaluation setting[19]. CoT can also be integrated into a few-shot format, where an example question is joined with a step by step answer breakdown[12]. This allows the model to mirror the analytical framework used in the example when evaluating a new financial situation so that the conclusion is lined up with a clearer chain of evidence. Figure 3 illustrates this prompt structure difference through a Snowflake case study in which the standard prompt includes a company customer description in its analysis, while the CoT few-shot prompt guides the model toward a more forecast-oriented investment analysis.

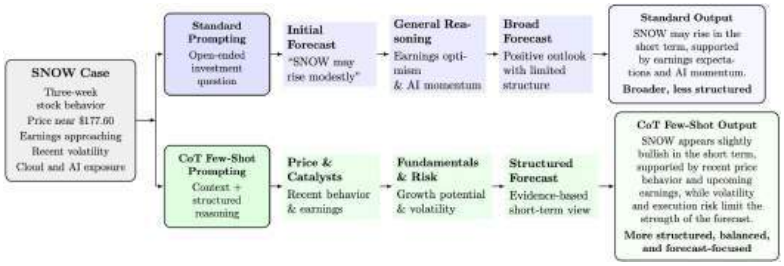

Figure 3. The standard prompt and chain-of-thought few-shot prompting comparison was used to show the different structural outlines to answer an investment query. The standard prompt generated a broad company description that shifted focus on general Snowflake use cases. However, the chain-of-thought few-shot prompt was able to organize the response around recent price behavior, earnings catalysts, fundamentals, valuation risk, and forecast implications. This comparison shows how a structured prompting strategy can help the model remain focused on the investment question rather than drifting into unrelated company backgrounds.

Even though CoT can improve the structure and interpretability in model outputs, its effectiveness still depends on the task and prompt quality. Additionally, the model's ability to

connect reasoning steps does not guarantee accurate reasoning[12]. Nonetheless, CoT can still play an important role in financial applications where forecasts require deep justification and weighing the various uncertainties before coming up with a final investment conclusion.

Tree-of-Thought (ToT) Prompting

Tree-of-Thought (ToT) prompting [20] extends CoT by exploring multiple candidate reasoning paths at each step and backtracking irrelevant steps rather than committing solely to one chain. A ToT prompt usually directs the model to first come up with several logical proposals for each step. Then each stage gets evaluated before proceeding to use the most effective options[20]. This structure helps models reduce the risk that comes from selecting one incomplete or irrelevant reasoning path, which is especially beneficial in financial decision making settings[6]. As a result, ToT has been shown to outperform CoT on the benchmark task "Game of 24" where the goal is to reach 24 from a set of numbers. Under one study's evaluation conditions, ToT had a 74% success rate compared to 4% for CoT[20]. In a financial aspect, ToT has been applied to facilitate a structured analysis of investment options, accounting treatments, and valuation strategies under uncertainty[6]. In investment analysis, ToT is especially useful because financial decisions often compare multiple interpretations rather than fixating upon one pathway. However, ToT also has limitations. Since the model generates several potential paths, it may produce contrasting but relevant reasoning steps that are difficult to merge into a final conclusion[9]. Therefore, ToT can support investment analysis by adding multiple layers of financial reasoning; however, its usefulness largely depends on whether those paths are synthesized into a unified financial assessment[9].

Graph-of-Thought (GoT) Prompting

Graph-of-Thought (GoT) prompting [9] improves upon prior reasoning methods by allowing relationships among each reasoning step. Unlike CoT which is based upon one reasoning pathway, and ToT which explores multiple ideas, GoT allows reasoning across steps to be revised, connected, and integrated. This is especially relevant to financial forecasting because it allows for more flexibility, and financial judgment often relies on several investment factors in order to formulate a conclusion. In practice, a GoT prompt tells a model to generate several answer solutions, review each step's relevance, improve upon weaker paths, and merge the strongest reasoning segments to form the final answer[6]. GoT may be more effective in the organization and combination of multi-step reasoning but its outputs still require additional review before forming a conclusion. In finance-related tasks, GoT-style prompting has been used to reason about connected factors such as market conditions, relationships between assets, and industry-specific drivers[6].

Financial-Chain-of-Thought (Fin-CoT) Prompting

The Financial Chain-of-Thought (Fin-CoT or FinCoT), a more recent prompting method[8], builds on standard Chain-of-Thought prompting by incorporating domain-specific

reasoning. Fin-CoT includes the step-by-step breakdown approach, while introducing explicit reasoning tags and domain specified financial reasoning outlines[8]. Fin-CoT also embeds in-depth financial reasoning through a mermaid diagram, which provides more structural guidance for solving more complex problems in the finance domain, as human input is incorporated to further support reasoning. As shown in Figure 4, the AMD investment forecasting process begins with a user defined problem, after which the model gathers relevant information such as AI-demand, profitability, and risk signals[8]. Next, the human selects the most relevant components of the output. Finally, the model refines the response through a summarization and self-reflection stage.

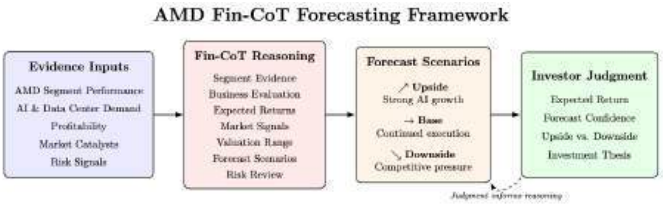

Figure 4. The AMD Fin-CoT forecasting framework illustrates the interaction that happens between evidence inputs, model reasoning, and investor judgment. The framework consists of seven reasoning stages including segment evidence, business evaluations, expected returns, market signals, valuation ranges, forecast scenarios, and risk review. Additionally, investor judgment and evidence inputs heavily influence the framework's reasoning allowing for a more effective investment thesis.

Prior work reported that Fin-CoT improved financial reasoning performance for several models on controlled settings, although the effects varied by model type[8]. However, the method requires considerable human-to-model interaction, which limits its practicality in time-sensitive forecasting scenarios[8]. Therefore, Fin-CoT may be beneficial when interpretability and structured decision-making are prioritized over speed, making it highly applicable to deeper investment forecasting research.

Section 3: Experimental Procedures

Materials

The experiment used OpenAI's original GPT-5 large language model, released on August 7th, 2025. GPT-5 was selected as the primary model for this experiment because it was an accessible model through ChatGPT and allowed all prompting strategies to be tested without additional model training or tuning. All initial trials were issued through the ChatGPT web interface using GPT-5 with default settings; no API parameters such as temperature were modified. The three accounts used for independent generations were on the same subscription tier, and the same model selection was used for every trial. Initial experimental prompts were administered from August 9 to 12, 2025. An expanded evaluation of Fin-Pro-1 and zero-shot prompting was conducted on October 8, 2026 using OpenAI API with the original GPT-5 model and original trial settings (model id: gpt-5-2025-08-07, reasoning effort: medium (Default), and no output token limit), with 30 trials per company to further assess forecasting consistency.

Financial market data were collected from Yahoo Finance for Advanced Micro Devices (AMD)[21], Snowflake (SNOW)[22], and Tesla (TSLA)[23]. Supplemental company financial data were collected from each company's investor relations materials, including AMD financial results[24], Snowflake quarterly results[25], and Tesla's Last Quarter 2021 Update[26]. Google Colab was used for metric calculations and visualizations. Overleaf LaTeX was used to format and generate the final presentation of tables and visualizations.

Prompting Methods & Our Prompting Method

We evaluated twelve established prompting methods: Zero-Shot, Few-Shot, Emotional-Compliment (shown as "Compliment" in tables and figures), Emotional-Urgency (shown as "Urgency"), Take-a-Step-Back, Take-a-Deep-Breath, Persona (Market Maker) (shown as "Market Maker"), Persona (Momentum) (shown as "Momentum"), Chain-of-Thought (CoT), Tree-of-Thought (ToT), Graph-of-Thought (GoT), and Financial Chain-of-Thought (Fin-CoT). Zero-shot prompting was used as the baseline condition for our experiment as it was the least structured prompt format evaluated within this study. All other prompting strategies were compared against zero-shot under the same model, company data, and forecast windows. In addition to the pre-existing prompting methods, this study also evaluated a pre-defined composite prompting strategy, referred to as Fin-Pro-1. Fin-Pro-1 integrates several prompting techniques: Emotional-Compliment, Emotional-Urgency, Take-a-Step-Back, Take-a-Deep-Breath, and Fin-CoT. Our hypothesis driven prompting strategy was created in order to support thorough financial forecasting evaluation including but not limited to response structure, careful assessment, and domain specific knowledge. We hypothesized that the combination of these techniques would produce more structured and interpretable forecasts than any strategy alone. Specifically, emotional cues were used to encourage the model to be more elaborate and task-focused, reflective prompts were used to elicit reassessment and caution, and Fin-CoT was used to provide an overarching, financially centered reasoning structure. Because Fin-Pro-1 was evaluated on only three companies in the case study, future research should test the same predefined prompt upon a larger set of companies.

Data Points

The data points included each company's historical stock performance across multiple time horizons: short-term performance over 1 week, medium-term performance over 4 weeks, and long-term performance over 2 months, 6 months, and 2 years. All-time performance was also included to give broader historical context. Supplemental financial context included non-GAAP net income, earnings per share (EPS), and company revenue for the most recent through the fifth most recent reporting periods. Analyst EPS estimates were also included. For each forecast trial, only information available on or before the forecast reference date was included in the prompt.

Tables 1 (Time Horizon) and 2 (Metric) present AMD's organized data points as an example of the data structure used in the experiment.

Time Horizon	Price Change
1-week (Short-Term)	+13.93%
4-week (Med-Term)	+32.71%
2-month (Long-Term)	+16.04%
6-month (Long-Term)	-13.14%
2-year (Long-Term)	+23.00%
All-Time	+666.83%

Metric	Value	Notes
Latest EPS	$0.96c	Beat estimate
Three-Month EPS	$1.09c	Beat estimate
Six-Month EPS	$0.92c	Beat estimate
Nine-Month EPS	$0.69c	Beat estimate
Year EPS	$0.62c	Beat estimate
Analyst Consensus EPS	$0.49c	–
Latest Quarterly Revenue	$7.44B	–
Three-Month Quarterly Revenue	$7.66B	–
Annual Revenue	$5.45B	–

Source: Yahoo Finance and AMD Investor Relations financial results, accessed August 1st, 2025.
Note: This same structure was applied to SNOW and TSLA before prompting the language model.

Trial Structure

For each company, a historical forecast reference date was randomly selected. AMD was evaluated as of May 16, 2025, SNOW as of November 25, 2022, and TSLA as of December 23, 2021. All financial information provided to the model was restricted to information available on or before the selected reference date. Each prompting strategy was initially evaluated through three independent generations per company, with company-level forecasts averaged for analysis. Fin-Pro-1 and zero-shot prompting were later expanded to 30 independent generations per company, yielding 90 responses per strategy per forecasting horizon. Accordingly, Tables 5–7 report different trial counts across strategies while retaining three distinct company-date scenarios. For eleven original comparison strategies, the reported MAPE, RMSE, bias, and directional accuracy values were calculated similarly, using an average of these three company forecasts. For Fin-Pro-1 and zero-shot, metrics were calculated directly from 90 individual forecasts per strategy per horizon, accounting for run-level variability. Consequently, comparisons between strategies with different aggregation methods should be interpreted cautiously.

In addition, company names and tickers were hidden from the model to reduce, but not eliminate, potential reliance on prior knowledge. Each prompting trial was also conducted on a new chat to prevent prior conversation context. For the three individual trials, which were later averaged, each trial was conducted through a different ChatGPT account while using the same model across trials. The prompting strategies were prompted in the same fixed order listed in the Prompting Methods section. Responses were also regenerated only when the model produced an empty forecast output. Specifically, regeneration happened when a safety related forecasting alert returned, or if it omitted a numerical forecast for one or more forecast horizons. In total, 12 responses were regenerated according to these predefined criteria.

Evaluation Metrics

The prompted LLMs were tasked with forecasting whether a stock's price would increase or decrease over predefined time horizons, including 1-week, 4-week, and 2-month periods from selected historical reference dates for which relevant data were available. For each time frame, the model was asked to provide percentage-based forecasts. To evaluate the effectiveness of each prompting strategy, this study examined both forecast performance and response interpretability.

As a quantitative comparison for evaluating LLM forecasting results, two non-LLM baselines were also incorporated: a no-change baseline and a historical-momentum baseline. The no-change baseline predicts a 0% future stock return across all three forecasting horizons. The historical-momentum baseline, associating projections with market trends, predicts that each stock's future return would equal its observed historical return from the prior period. To assess statistical reliability, 95% confidence intervals (CIs) for mean absolute error (MAE) were estimated using a t-based method, combining within-company variability across 30 independent trials per company for Fin-Pro-1 and zero-shot. These intervals reflect uncertainty from repeated generations within the three fixed company-date scenarios, rather than across broader market conditions.

Additionally, forecast performance was also measured using mean absolute percentage error (MAPE), root mean squared error (RMSE), forecast bias, and directional accuracy. Interpretability was separately evaluated using a six dimension scoring rubric. MAPE was used to compare forecast error relative to the size of the actual movement, with lower values indicating better forecast performance. RMSE measured how far predicted percentage movements were from the actual movements, with lower values indicating better accuracy. Forecast bias determined whether predictions consistently overestimated or underestimated actual stock movement, indicating the direction of any systematic skew. Directional accuracy was employed to determine whether the model identified the correct direction even when the predicted magnitude was inaccurate.

Interpretability score was used to measure the clarity and transparency of the model. The rubric scored six dimensions: evidence use, risk discussion, acknowledgement of uncertainty, use of financial indicators, specificity of forecast reasoning, and avoidance of unsupported claims. The score evaluates how clearly the model presents and supports its forecast reasoning. Additionally, this rubric was established before the final model outputs were analyzed, and the evaluator had limited awareness of the prompting condition associated with some responses, so the evaluation was not fully blinded. The initial model responses were evaluated and scored by the author using the predefined rubric. Furthermore, these assigned scores were subsequently reviewed by two research mentors. Each dimension was scored from 0 to 2, and scores were averaged across prompting strategies according to Table 3.

Dimension	0	1	2
Evidence Use	Does not use the provided data	Uses one to two relevant data point	Uses more than two relevant data points
Risk Discussion	Does not discuss risk	Mentions risks but no explanation	Explains multiple specific risk and explains
Acknowledgement of Uncertainty	Presents the prediction with little or no uncertainty	Uses general uncertainty language (e.g. maybe)	Explains possible conditions or alternative outcomes
Use of Financial Indicators	Does not use financial indicators	Uses one to two financial indicator	Uses more than two financial indicators
Specificity of Forecast Reasoning	Gives little explanation for the prediction	Provides a general reason for the prediction	Clearly explains step by step information
Avoidance of Unsupported Claims	Includes two or more unsupported claims	Includes one unsupported claim	Does not include unsupported claims

Table 3. Interpretability scoring rubric for LLM-generated financial forecasts.

Section 4: Results and Discussion

Non-LLM Quantitative Baselines Results

In the two non-LLM quantitative baselines, the no-change baseline achieved lower MAPE, RMSE, and absolute bias values across all three horizons, indicating that using past performance to predict future returns does not guarantee more accurate forecasts in these selected company-date scenarios. These findings also highlight the importance of comparing LLM-generated forecasts against non-LLM baselines rather than exclusively relying on comparisons among prompting strategies. However, the limited number of independent company-date scenarios restricts the broader applicability of these comparisons.

Horizon	Baseline	*n*	MAPE (%)	RMSE	Bias (%)	DirAcc (%)
1 week	No-change	3	100.00	3.97	1.11	0.00
1 week	Momentum	3	689.42	14.69	10.22	0.00
4 weeks	No-change	3	100.00	8.44	2.66	0.00
4 weeks	Momentum	3	242.93	17.47	7.01	100.00
2 months	No-change	3	100.00	25.31	-4.18	0.00
2 months	Momentum	3	26,735.86	64.58	24.09	33.33

Table 4. Forecasting performance of non-LLM quantitative baselines across three forecasting horizons, evaluated using three company-date scenarios (n = 3 per horizon).

95% Confidence Intervals Results

Across the 1-week, 4-week, and 2-month horizons respectively:
Zero-shot MAEs (95% CI; percentage points): 8.02 (7.08–8.96), 11.19 (9.29–13.08), 24.06 (20.64–27.48)
Fin-Pro-1 MAEs (95% CI; percentage points): 2.32 (1.93–2.72), 3.82 (3.07–4.56), 9.23 (7.80–10.66)

In the 95% confidence interval report, Fin-Pro-1 consistently achieved lower mean absolute errors than zero-shot across all three forecasting horizons, with the confidence intervals suggesting a potential forecasting accuracy increase under the evaluated conditions. However, these findings remain limited to three company-date scenarios.

Note: In Tables 5–7, n denotes the number of forecast generations per company (30 for Fin-Pro-1 and zero-shot; 3 for the other strategies). All strategies were evaluated across three company-date scenarios. DirAcc = directional accuracy (hit rate, %). The original eleven strategies used averaged company forecasts, while Fin-Pro-1 and zero-shot used individual forecast errors to measure variability across trials. Since the calculations differed, comparisons between the two groups should be interpreted with caution.

Short-Term Forecast Results

In the short-term 1-week forecast, Fin-Pro-1 achieved the lowest RMSE and the second lowest MAPE in this controlled case-study sample. Relative to the zero-shot baseline, Fin-Pro-1 reduced MAPE from 418.0% to 80.22% and RMSE from 9.28 to 3.29. These results suggest that the combined prompting technique may have helped model responses produce forecasts better aligned with the realized stock movements. Fin-Pro-1 also showed the smallest absolute bias at -0.48, suggesting that its forecasts were less directionally skewed within the sample. This pattern

may be consistent with the model remaining more focused in the financial context, although this bias metric does not alone establish this. However, since each strategy was evaluated on only three company-date scenarios, with different numbers of trials across prompting strategies, findings are not final and require further testing.

Prompting Strategy	n	MAPE (%)	RMSE	Bias (%)	DirAcc (%)
Zero-Shot	30	418.05	9.28	2.54	27.78
Few-Shot	3	171.97	11.46	8.74	33.33
Compliment	3	122.03	6.84	2.24	33.33
Urgent	3	197.03	11.08	7.08	0.00
Take-a-Deep-Breath	3	176.77	11.63	8.58	33.33
Take-a-Step-Back	3	145.47	9.10	6.48	33.33
Market	3	95.03	9.16	6.91	33.33
Momentum	3	220.03	11.91	7.08	0.00
CoT	3	137.80	9.23	6.74	33.33
ToT	3	159.33	9.83	7.11	33.33
GoT	3	158.30	10.41	7.94	33.33
Fin-CoT	3	58.10	6.30	3.68	66.67
Fin-Pro-1	30	80.22	3.29	-0.48	85.56

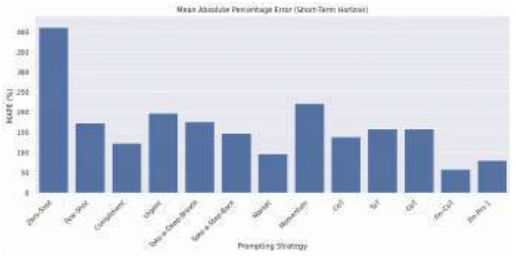

Table 5. Short-term 1-week performance metrics for each prompting strategy. Figure 5. Mean absolute percentage error (MAPE) for each prompting strategy in the short-term 1-week horizon forecast.

Medium-Term Forecast Results

In the medium-term 4-week forecast, Fin-Pro-1 achieved the second lowest MAPE at 43.48%, compared with 130.59% for the zero-shot baseline. Furthermore, Fin-Pro-1 reduced RMSE from 15.88 to 5.60 and achieved a higher directional accuracy of 91.11% compared with 63.33%. Although several other prompting strategies achieved 100.00% directional accuracy, they were evaluated using only three trials per company, limiting direct comparisons. In addition, Fin-CoT achieved the lowest RMSE at 2.29 and MAPE of 34.37%, although fewer trials make direct comparison with the 30 trial Fin-Pro-1 results less reliable. These findings suggest potential improvements over zero-shot prompting, although the limited company-date scenarios restrict broader conclusions.

Prompting Strategy	n	MAPE (%)	RMSE	Bias (%)	DirAcc (%)
Zero-Shot	30	130.59	15.88	6.14	63.33
Few-Shot	3	85.30	9.93	6.35	100.00
Compliment	3	51.97	3.18	−0.65	100.00
Urgent	3	49.17	3.78	0.35	100.00
Take-a-Deep-Breath	3	84.13	7.56	5.02	100.00
Take-a-Step-Back	3	109.17	7.33	3.69	100.00
Market	3	45.80	4.86	3.52	100.00
Momentum	3	110.23	8.40	5.19	100.00
CoT	3	84.13	6.48	2.19	100.00
ToT	3	117.70	7.74	7.52	66.67
GoT	3	69.23	6.17	4.79	100.00
Fin-CoT	3	34.37	2.29	−0.81	100.00
Fin-Pro-1	30	43.48	5.60	0.86	91.11

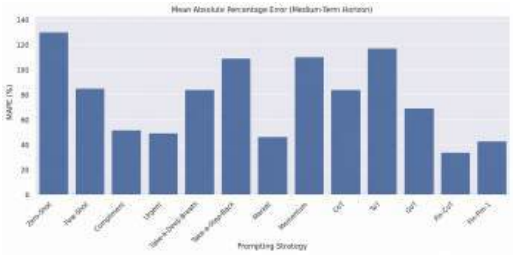

Table 6. Medium-term 4-week performance metrics for each prompting strategy. Figure 6. Mean absolute percentage error (MAPE) for each prompting strategy in the medium-term 4-week horizon forecast.

Long-Term Forecast Results

In the long-term 2-month forecast, Fin-Pro-1 achieved the lowest MAPE at 7,675.82%, compared with 12,943.24% for the zero-shot baseline and reduced RMSE from 31.40 to 11.85. Take-a-Deep-Breath also achieved the lowest RMSE at 8.48 and the highest directional accuracy at 100.00%, although fewer trials make direct comparison with the 30 trial Fin-Pro-1 results less reliable. However, the extremely large MAPE values should be interpreted with caution given

that SNOW's realized 2-month return was only +0.02%. Because MAPE divides forecast error by the actual return magnitude, SNOW's near-zero return (+0.02%) greatly inflates the percentage error. Therefore, MAPE is not a suitable comparison for this horizon, and RMSE provides a more informative forecast result comparison. Although Fin-Pro-1 outperformed zero-shot on both error metrics, other prompting strategies achieved lower RMSE, suggesting that its performance advantages were not consistent across all methods.

Prompting Strategy	n	MAPE (%)	RMSE	Bias (%)	DirAcc (%)
Zero-Shot	30	12 943.24	31.40	1.32	41.11
Few-Shot	3	12 590.90	31.49	−25.07	33.33
Compliment	3	19 997.37	18.20	−16.07	66.67
Urgent	3	16 755.73	33.03	−21.82	66.67
Take-a-Deep-Breath	3	41 654.47	8.48	−4.32	100.00
Take-a-Step-Back	3	13 333.00	20.44	−16.32	66.67
Market	3	39 984.20	10.57	−10.32	66.67
Momentum	3	16 747.80	33.39	−27.07	33.33
CoT	3	20 004.20	19.37	−18.32	66.67
ToT	3	24 992.37	15.87	−13.82	66.67
GoT	3	34 487.20	12.30	−11.47	66.67
Fin-CoT	3	14 995.57	19.61	−15.07	66.67
Fin-Pro-1	30	7 675.82	11.85	−1.32	77.78

Table 7. Long-term 2-month performance metrics for each prompting strategy. Figure 7. Root Mean Squared Error (RMSE) for each prompting strategy in the long-term 2-month horizon forecast.

Heatmap of LLM Response Interpretability Across Prompting Strategies

Figure 8 shows that structured reasoning strategies consistently scored higher on interpretability than simpler prompts. Fin-Pro-1 performed strongest across evidence use, financial indicators, and reasoning specificity, while zero-shot prompting scored lowest overall. Accordingly, these scores reveal the structure and transparency of the responses rather than the accuracy of their forecasts. Additionally, scores were based on the initial three responses per company for each prompting strategy, separate from the expanded forecasting accuracy evaluation.

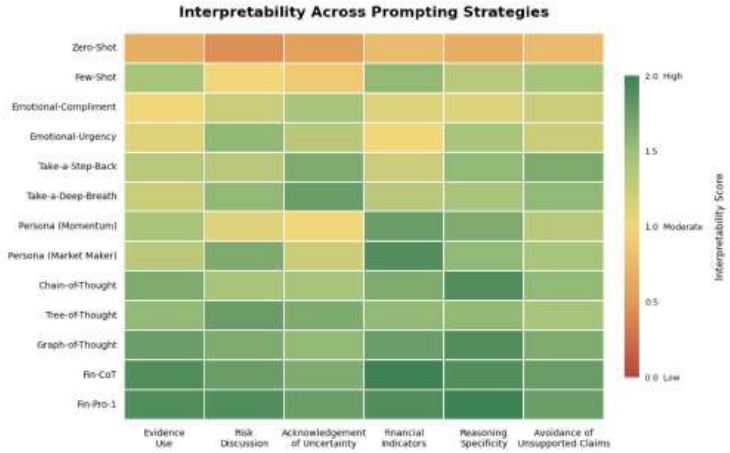

Figure 8. Heatmap comparing interpretability scores across prompting strategies and evaluation criteria.

Inference of Our Findings

Overall, the results suggest that prompt structure was associated with differences in forecast accuracy and the response interpretability for the model tested (GPT-5) under the

controlled conditions of this study. Fin-Pro-1 achieved the second-lowest MAPE in both the 1-week and 4-week forecasting windows, along with the lowest RMSE in the 1-week window. Fin-Pro-1 achieved lower RMSE than the no-change baseline at all three horizons, with values of 3.29 compared to 3.97 for 1-week, 5.60 compared to 8.44 for 4-week, and 11.85 compared to 25.31 for 2-month. Compared with zero-shot prompting, Fin-Pro-1 achieved lower MAPE and RMSE across all three forecasting horizons. However, other prompting strategies achieved lower RMSE at the 4-week and 2-month horizons, suggesting that Fin-Pro-1's performance improvements were not consistent across all methods. The results also show that financially structured prompting methods such as Fin-Pro-1 and Fin-CoT were associated with more organized and interpretable forecast responses than less structured prompts. Rather than giving the model a query to only rely on, these methods guided the model with a framework that helped it explain relevant risks, price movement assumptions, and reasoning behind the forecast. Additionally, these findings suggest that LLM-generated forecasts may help investors better understand financial information and potential risks, but those detailed explanations do not guarantee robust predictions. Therefore, LLMs should serve as supplementary tools beneficial for financial analysis rather than the primary basis for investment decision-making. Although Fin-Pro-1 and zero-shot prompting were evaluated using 30 independent trials per company, the remaining strategies were evaluated using only three trials per company. Furthermore, the study remained limited to three company-date scenarios, requiring further evaluation across larger datasets, additional market conditions, and AI models before drawing broader conclusions.

Section 5: Limitations and Future Directions

Limitations

This study is subject to several limitations. Although no-change and historical-momentum benchmarks were incorporated alongside zero-shot prompting, more advanced quantitative forecasting models, such as ARIMA and random forest, were not evaluated. In addition, the analysis included only three publicly traded companies, Advanced Micro Devices (AMD), Tesla (TSLA), and Snowflake (SNOW). Including a larger variety of models for analysis, finance specific models such as Fin-GPT, Fin-LLaMA, or Fin-BERT, would make the results apply to broader finance cases and support clearer comparisons. Second, the experiments were conducted using a restricted historical financial dataset. In order to reduce the potential reliance on pre-existing information, real company tickers and concurrent news content were not disclosed to the model. However, anonymization cannot fully guarantee elimination of historical data leakage because underlying LLMs may have encountered information about these companies during training. Therefore, the reported forecasting performance may partially reflect prior training knowledge rather than authentic analytical analysis. Third, although Fin-Pro-1 and zero-shot prompting were evaluated using 30 independent trials per company, the remaining prompting strategies were evaluated using only three trials per company. 95% confidence intervals were incorporated to assess forecasting variability for Fin-Pro-1 and zero-shot, but

these measures reflect uncertainty across repeated generations within the three selected company-date scenarios rather than broader market conditions. Future research should expand the number of independent company-date scenarios and apply consistent trial counts across all prompting strategies. Another limitation concerns the accessibility of prompt engineering for investors without technical expertise. One possible approach is to develop reusable frameworks, such as Fin-Pro-1, that simplify user interactions by requiring only relevant financial information, as prompting instructions are predefined. This approach reduces the need for users to construct specialized prompts themselves, but further research is necessary to determine the reliability of forecasts generated by non-expert users.

Difficulties

LLMs may produce hallucinations: fluent responses that are factually incorrect or entirely made up[3]. These hallucinations occur when LLMs are not properly pre-trained or lack high-quality data. Though not always frequent, these errors pose significant risks to the finance sector. For LLMs to be trusted as investment analysts, they must demonstrate ethical behavior and must be unbiased. Potential negative outcomes include skewed graphs, biased interpretations, and misleading summaries which can cause huge damage[2]. Such failures could result in bad investment decisions and even legal consequences[6]. However, these limitations are not distinctive to LLMs, as traditional forecasting models also rely on historical training data and market condition influences that can lead to inaccuracies and biases. Although LLMs can introduce biases through their training data and interpretations of market conditions, these limitations do not remove the value they offer in financial analysis compared with traditional forecasting methods.

Future Directions

Future progress in this area is likely to benefit from an addition of financial-based data retrieval systems. While prompt engineering can help with the structural output from a model, information retrieval techniques are necessary to ensure that models receive the most relevant and recent sources[27]. This is especially important in financial forecasting because finance related articles, analyst commentary may contain biases and selective framing that often affect LLM decision making. In order to help reduce risks, future systems could integrate Retrieval-Augmented Generation (RAG) to retrieve finance specific data sources such as earnings reports, SEC filings such as 10-K and 10-Q reports, and market data[27]. As a result, future LLM based forecasting systems should combine retrieval methods with structured prompt engineering in order to produce transparent and thorough financial forecasts.

Section 6: Conclusion

This study examined how prompt engineering strategies shape the structure and interpretability of LLM generated stock forecasting responses under case-study conditions. Through the application of prompting strategies to time-series stock data, the analysis focused on

understanding how prompt design plays a role in affecting the model reasoning and response architecture. The results suggest that, under the conditions that were tested, prompt design was associated with differences in LLM generated financial forecasts performance and interpretability. Under the three evaluated company-date scenarios, Fin-Pro-1 achieved lower MAPE and RMSE than zero-shot prompting across all three forecasting horizons, although its advantages were not consistent across all prompting strategies. Future research could compare these methods with soft prompting, where continuous prompt representations are learned automatically. Analysis on whether they improve financial forecasting performance and interpretability could also provide further insight into different variants of prompting methods in financial applications.

AI Acknowledgement

During the preparation of this work, the author used ChatGPT only as the large language model being evaluated in the experiment. Specifically, it was tasked to generate stock-forecasting responses in order to examine the different prompting strategies. ChatGPT was not used to write, draft, or revise the paper. The author also used Grammarly for spelling, grammar, and punctuation checks. Google Search was used to find better vocabulary replacements and find source material. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

Acknowledgements

I would like to thank my mentors Professor Plínio Zanini and Siddharth Krishnan, Ph.D., University of North Carolina - Charlotte, for their guidance in my finalized paper. I would also like to thank Kiranjyot Gill, a LIGO-affiliated astrophysicist, for her valuable feedback during the final stages of this manuscript.

Authors

Hugo Shek is a student at Mills High School in Millbrae, California. He is a data scientist, entrepreneur, and researcher interested in large language models, financial analysis, and technology for real-world problem solving. He also develops Refueler, his behavioral research translated into a digital wellness and motivational coach app.

Works Cited

(1)Remolina, N. Generative AI in Finance: Risks and Potential Solutions. Law Ethics Technol. **2024**. https://doi.org/10.55092/let20240002

(2)Chen, S.; Green, T. C.; Gulen, H.; Zhou, D. What Does ChatGPT Make of Historical Stock Returns? Extrapolation and Miscalibration in LLM Stock Return Forecasts. Elsevier BV 2024. https://doi.org/10.2139/ssrn.4941906

(3)Nie, Y.; Kong, Y.; Dong, X.; Mulvey, J. M.; Poor, H. V.; Wen, Q.; Zohren, S. A Survey of Large Language Models for Financial Applications: Progress, Prospects and Challenges. arXiv June 15, 2024. https://doi.org/10.48550/arXiv.2406.11903

(4)Darwish, M.; Hassanien, E. E.; Eissa, A. H. B. Stock Market Forecasting: From Traditional Predictive Models to Large Language Models. Comput. Econ. **2026**, 67 (6), 4553–4597. https://doi.org/10.1007/s10614-025-11024-w

(5)Chen, T. T.; Zheng, B.; Li, Y.; Jiang, X. F. New Approaches in Agent-Based Modeling of Complex Financial Systems. Front. Phys. **2017**, 12 (6), 128905. https://doi.org/10.1007/s11467-017-0661-2

(6)Joshi, S. Review of Prompt Engineering Techniques in Finance: An Evaluation of Chain-of-Thought, Tree-of-Thought, and Graph-of-Thought Approaches. IJIRCST **2025**, 13 (4), 51–67. https://doi.org/10.55524/ijircst.2025.13.4.6.

(7)Sahoo, P.; Singh, A. K.; Saha, S.; Jain, V.; Mondal, S.; Chadha, A. A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications. arXiv March 16, 2025. https://doi.org/10.48550/arXiv.2402.07927

(8)Nitarach, N.; Sirichotedumrong, W.; Pitchayarthorn, P.; Taveekitworachai, P.; Manakul, P.; Pipatanakul, K. FinCoT: Grounding Chain-of-Thought in Expert Financial Reasoning. arXiv June 19, 2025. https://doi.org/10.48550/arXiv.2506.16123

(9)Besta, M.; Blach, N.; Kubicek, A.; Gerstenberger, R.; Podstawski, M.; Gianinazzi, L.; Gajda, J.; Lehmann, T.; Niewiadomski, H.; Nyczyk, P.; Hoefler, T. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. Proc. AAAI Conf. Artif. Intell. **2024**, 38 (16), 17682–17690. https://doi.org/10.1609/aaai.v38i16.29720

(10)Vatsal, S.; Dubey, H. A Survey of Prompt Engineering Methods in Large Language Models for Different NLP Tasks. arXiv July 24, 2024. https://doi.org/10.48550/arXiv.2407.12994

(11)Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; Amodei, D. Language Models Are Few-Shot Learners. arXiv July 22, 2020. https://doi.org/10.48550/arXiv.2005.14165

(12)Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q.; Zhou, D. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv January 10, 2023. https://doi.org/10.48550/arXiv.2201.11903

(13)Li, C.; Wang, J.; Zhang, Y.; Zhu, K.; Hou, W.; Lian, J.; Luo, F.; Yang, Q.; Xie, X. Large Language Models Understand and Can Be Enhanced by Emotional Stimuli. arXiv November 12, 2023. https://doi.org/10.48550/arXiv.2307.11760
(14)Yang, C.; Wang, X.; Lu, Y.; Liu, H.; Le, Q. V.; Zhou, D.; Chen, X. Large Language Models as Optimizers. arXiv April 15, 2024. https://doi.org/10.48550/arXiv.2309.03409
(15)Sadr, N. G.; Madhusudan, S.; Emami, A. Think or Step-by-Step? UnZIPping the Black Box in Zero-Shot Prompts. arXiv February 5, 2025. https://doi.org/10.48550/arXiv.2502.03418
(16)Lopez-Lira, A. Can Large Language Models Trade? Simulating Financial Theories and Experiments Using LLM Agents. **2025**, 47. https://doi.org/10.2139/ssrn.5217340
(17)Zheng, H. S.; Mishra, S.; Chen, X.; Cheng, H.-T.; Chi, E. H.; Le, Q. V. Take a Step Back: Evoking Reasoning via Abstraction in Large Language Models. arXiv March 12, 2024. https://doi.org/10.48550/arXiv.2310.06117
(18)Roy, S.; Roth, D. Solving General Arithmetic Word Problems. arXiv August 20, 2016. https://doi.org/10.48550/arXiv.1608.01413
(19)Chen, J.; Chen, L.; Huang, H.; Zhou, T. When Do You Need Chain-of-Thought Prompting for ChatGPT? arXiv April 18, 2023. https://doi.org/10.48550/arXiv.2304.03262
(20)Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; Narasimhan, K. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. In Advances in Neural Information Processing Systems 36; Neural Information Processing Systems Foundation, Inc., 2023; pp 11809–11822. https://doi.org/10.52202/075280-0517
(21)Advanced Micro Devices, Inc. (AMD) Stock Price, News, Quote & History - Yahoo Finance. https://finance.yahoo.com/quote/AMD/ (accessed 2025-08-01)
(22)Snowflake Inc. (SNOW) Stock Price, News, Quote & History - Yahoo Finance. https://finance.yahoo.com/quote/SNOW/ (accessed 2025-08-01)
(23)Tesla, Inc. (TSLA) Stock Historical Prices & Data - Yahoo Finance. https://finance.yahoo.com/quote/TSLA/history/ (accessed 2025-08-01)
(24)AMD Investor Relations. Advanced Micro Devices, Inc. https://ir.amd.com/financial-information/financial-results (accessed 2025-08-01)
(25)Snowflake - Financials - Quarterly Results. https://investors.snowflake.com/financials/quarterly-results/default.aspx (accessed 2025-08-01)
(26)Tesla Investor Relations. https://ir.tesla.com/#quarterly-disclosure (accessed 2025-08-01)
(27)Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.; Rocktäschel, T.; Riedel, S.; Kiela, D. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv April 12, 2021. https://doi.org/10.48550/arXiv.2005.11401

Appendix

Future Return Window (AMD): Short-Term: -5.85%, Medium-Term: +7.87%, Long-Term: +36.62%
Future Return Window (SNOW): Short-Term: +3.47%, Medium-Term: -4.30%, Long-Term: +0.02%
Future Return Window (TSLA): Short-Term: -0.96%, Medium-Term: -11.54%, Long-Term: -24.10%
https://github.com/Hgos253/research_stockforecasting/blob/main/AI%20Prompting%20Experiment_%20Supplementary%20Materials%20(1).pdf

Traditional and Emerging Technologies for Glucose Management in Type 2 Diabetes
By Siddharth Tirumalai

Abstract

This paper compares traditional and emerging diabetes management technologies to identify how each affects glucose management for people with type 2 diabetes. The traditional approach for managing blood glucose is through blood glucose monitoring (BGM) along with multiple daily injections (MDI) throughout the day. Diabetes management has progressed over the years, creating newer technologies, including continuous glucose monitoring (CGM), insulin pumps, and patch pumps. They were developed to address the limitations of the traditional diabetes management techniques. Based on evidence from randomized trials, meta-analyses, and clinical reviews covering glycemic control, hypoglycemia risk, device related complications, and patient reported experience, this review finds that CGMs and insulin pumps produce modest but meaningful improvements in A1C compared with fingerstick monitoring and multiple daily injections.

Introduction

More than 40 million Americans currently live with diabetes, and type 2 diabetes accounts for 90 to 95 percent of those cases (American Diabetes Association). For nearly a century managing type 2 diabetes has come down to one thing. A patient will check their blood sugar and need to calculate their own insulin dose and give it to themselves and do this over and over with no outside check if something goes wrong. Insulin has been given by syringe or pen since 1922, and injection is still how most people with type 2 diabetes take insulin once they need it (American Diabetes Association). For most of that time blood sugar monitoring and insulin delivery worked the same basic way. A patient checks a single number with a glucose meter and then decides on their own what to do with it whether that means eating differently or moving more or adjusting their insulin (American Diabetes Association). This model became the default way diabetes was managed long before anything more advanced came along and it is still what many patients rely on today. That started to change over the last two decades. Continuous glucose monitors (CGM) let patients see where their blood sugar is heading instead of only where it stands at one moment while insulin pumps can deliver insulin more continuously than a pen or syringe (Berget et al.; Klonoff et al.). Some advanced systems combine a CGM with an insulin pump and an algorithm that adjusts insulin delivery on its own. These systems are generally called closed-loop or hybrid closed-loop systems (Berget et al.). This has changed care substantially for type 1 diabetes. For type 2 the evidence is thinner partly because the disease is more heterogeneous and partly because it has drawn less research attention (Freckmann et al.).

At the same time the research does not suggest that newer technology is automatically better for every patient. Studies of insulin pump therapy in type 2 diabetes have produced mixed results with some showing significant improvements and others finding little difference compared with MDI (Freckmann et al.). CGM has shown more consistent benefits in several studies but the size of those benefits can depend on the patient's baseline glucose control and

insulin treatment and ability to use the information the device provides (Poolsup et al.; Irace et al.). Technology can reduce some of the work involved in diabetes management but it also introduces new challenges such as cost and device maintenance and skin reactions and the need for patient and provider training (Freckmann et al.). This review works through the evidence in that order, looking at BGM and MDI first and then CGM, insulin pumps, and patch pumps, and it pulls from trials, meta-analyses, and clinical reviews. How do traditional self-managed tools such as BGM and MDI compare with newer technologies such as CGM, insulin pumps, patch pumps, and automated systems for people with type 2 diabetes? More specifically, how do these approaches compare in terms of glucose control, treatment burden, hypoglycemia, and quality of life?

Discussion

This question matters because poor glucose control can contribute to complications affecting multiple organ systems including the cardiovascular system and kidneys and eyes and nervous system. (Li et al.) Diabetes-related vascular complications are interconnected rather than isolated problems, so long-term metabolic control can influence health across several organs at once (Li et al.). Because of this the technology used to monitor and control glucose can influence how consistently patients manage their disease and how effectively they respond to changes in their glucose levels.

Blood Glucose Monitor

The tools available to manage that risk have changed considerably since insulin was first introduced, though for most of the history of diabetes care those tools demanded nearly everything from the patient. The blood glucose meter was never treated as the treatment and instead showed how patients figured out where they stood, and it was on them to decide what to do with that number, whether that meant eating differently, moving more, or adjusting their insulin. So for a long time, manual monitoring and injections were just how diabetes care worked, forming the traditional treatment approach against which newer technologies are compared.

BGM gives patients a fast, cheap way to check their blood sugar at any given moment, and it's still a core part of care for anyone on insulin (American Diabetes Association). But the problem with BGM comes down to timing. A fingerstick measurement captures glucose at a single point in time and provides no information about the direction or rate of change. Most people using BGM are only checking a handful of times a day, while CGM can pull in close to 300 readings over the same 24 hours (Klonoff et al.). That gap matters most exactly when people aren't testing, like overnight, when blood sugar can drop low enough to be dangerous without any symptoms to wake someone up (American Diabetes Association).

On top of that, meter accuracy is shakier than most people assume. There are actually two different accuracy standards for glucose meters, one from an international standards group and one from the FDA, and meters approved under an older version of the FDA standard aren't required to meet the newer one (American Diabetes Association). In one investigation of

eighteen popular glucose meters, only six actually met the accuracy standard they were tested against (American Diabetes Association). Everyday things like high dose vitamin C, low or high oxygen levels, or extreme temperatures can throw a reading off too, often without giving any warning that something's wrong (American Diabetes Association). So just because a meter is FDA cleared doesn't mean it's accurate in every situation, and patients may not be aware of these potential sources of measurement error.

Even when the meter is accurate, plenty of people just don't test as often as they're supposed to, and pain is a big reason why. A cross-sectional survey of 44,181 adults with type 1 or type 2 diabetes, drawn from the Kaiser Permanente Northern California population, found that only about a third of participants with type 2 diabetes met the recommended testing frequency. The same researchers noted that most diabetes educators believe patients find insulin injections less painful than finger pricks, which goes against the usual assumption that injections are the scarier part. Some companies have tried fixing this with alternative site testing, letting people test on their forearm or thigh instead of their fingertip since those spots have fewer nerve endings, but at least one randomized study found that didn't actually improve how often people tested.

All of this helps explain why different BGM studies end up with such different results. One meta-analysis of regular, unstructured BGM use in people with type 2 diabetes found it only lowered A1C by about 0.25 to 0.3 percentage points after six months, and that small improvement had mostly worn off by the one year mark (American Diabetes Association). Another trial combined daily BGM with text reminders and still found no real change in A1C after a full year (American Diabetes Association). By contrast,(Polonsky et al.) followed 483 insulin-naive patients with type 2 diabetes, all with a baseline A1C of at least 7.5%, over 12 months. Patients assigned to structured BGM, meaning they logged seven point glucose profiles at least quarterly and reviewed them with a provider on a set schedule, saw their A1C drop by 1.2 percentage points, compared to a 0.9 percentage point drop in the group receiving enhanced usual care alone. That 0.3 percentage point gap between the two groups was statistically significant ($p = 0.04$). The studies reviewed here suggest that the clinical benefit of BGM may depend substantially on how glucose information is incorporated into treatment decisions, rather than on the frequency of testing alone.

Multiple Daily Injections

MDI, delivered through insulin pens or syringes, lets patients administer insulin without requiring an external insulin-delivery system such as a pump, and it's still how most people with type 2 diabetes take insulin once they need it (American Diabetes Association). Pens do offer some flexibility, with dosing that can go as fine as half unit increments depending on the device (American Diabetes Association). But that flexibility comes at a cost. With MDI, the patient calculates and administers every single dose themselves, and traditional MDI lacks any mechanism for automated dose adjustment (American Diabetes Association).

Traditional MDI requires repeated patient decision-making and execution, and this burden carries a documented association with psychological strain in addition to its physical demands. Poor adherence to MDI has been associated with depression, diabetes-related distress,

and other markers of self-management burden, though it's worth noting that the studies establishing this relationship are observational, meaning they can show that poor adherence and psychological distress tend to occur together without proving that one causes the other (American Diabetes Association). A narrative review by Brixner et al. examining MDI burden in type 2 diabetes found that patients consistently identified several specific barriers: fear of needles, dosing that feels too complicated, physical difficulty administering the injection, fear of hypoglycemia, weight gain, disruption to daily routines, and embarrassment about injecting around others. These reported barriers may also have measurable clinical consequences. Peyrot et al. surveyed 1,530 insulin-treated adults across eight countries and identified injection-related pain, interference with daily activities, and regimen complexity as significant, independent factors associated with insulin omission. Separately, Edelman et al. followed a retrospective cohort of patients with type 2 diabetes on basal-bolus insulin therapy and found that patients who were non-persistent with their regimen had significantly higher A1C than those who remained persistent (8.84% vs. 8.38%, $p<0.005$) and lower rates of treatment success (39% vs. 55%, $p<0.009$). Because this was an observational cohort study rather than a controlled trial, it establishes an association between nonadherence and worse outcomes rather than direct proof that missed doses alone caused the difference.

This barrier is not confined to patients. A cross-sectional survey of 76 primary care providers found that 88% reported discomfort initiating insulin pump therapy for their patients, and 89% reported discomfort adjusting it once treatment had begun (O'Donovan et al.), suggesting that provider training gaps, not only patient hesitation, contribute to why so many patients remain on MDI longer than their glycemic control alone would justify. Even newer connected insulin pens, which log the timing and size of each dose, don't resolve this underlying limitation, since the patient still has to determine the dose independently at the moment of injection (Yang et al.).

This burden shows up in comparative outcome data as well, not only in patient reported reasons for nonadherence. Ghazanfar et al. compared 100 insulin pump users to 350 patients managing type 2 diabetes without a pump and found that patients using pens or syringes reported markedly lower treatment satisfaction than pump users did. Among the pump users in that same study, only a single participant reported dissatisfaction, tied to reservoir capacity rather than to device performance. This comparison is useful because it isolates treatment satisfaction as its own clinical outcome, distinct from A1C, and confirms that the burden described above extends into how patients actually experience their treatment.

The research covered in this section has its own gaps, separate from the barriers patients deal with using these devices directly. A lot of the BGM studies aren't fully comparable to each other, since some tested regular, unstructured checking, while others tested a full structured program with provider review built in, which makes it difficult to say exactly how much benefit BGM offers on its own outside of that context. The MDI research has a related issue. Several of the sources used here rely on retrospective cohort data or cross-sectional surveys rather than randomized controlled trials, which means the relationships they describe, between adherence

barriers and clinical outcomes, are associations rather than proven cause and effect. More controlled research directly linking specific MDI barriers, such as injection-related fear or pain, to measurable glycemic outcomes in type 2 diabetes specifically would strengthen this section considerably. Overall, the evidence supports the conclusion that both BGM and MDI depend heavily on the patient's ability to follow through consistently, but the research base itself isn't uniform enough to say precisely how large that dependency is, which is the exact problem the technology covered in the next section is trying to address.

Continuous Glucose Monitoring and Insulin and Patch Pump Therapy

CGM and insulin pump technology were developed to address the limitations discussed in the previous section. Rather than providing a single reading at one point in time or requiring a fully manual injection, these tools generate continuous data, and in some configurations, can adjust insulin delivery automatically (Berget et al.). A CGM is a small sensor placed under the skin that measures glucose every five to fifteen minutes and transmits the reading to a phone or receiver, allowing the patient to observe a continuous trend rather than an isolated value (Klonoff et al.). An insulin pump delivers rapid acting insulin through a small subcutaneous tube rather than through discrete injections, which allows for more precise adjustment of both background and mealtime insulin than a pen can provide (Berget et al.). When a CGM and pump are combined, the pump can use the CGM's data to adjust insulin delivery directly rather than simply displaying it, a configuration referred to as a closed loop or hybrid closed loop system (Berget et al.). Although these technologies were originally developed for type 1 diabetes, approximately ten percent of insulin pump users have type 2 diabetes, and that proportion has continued to grow, which is directly relevant to the scope of this paper (Berget et al.).

CGM addresses one of BGM's core limitations directly, replacing single point readings with continuous data. A meta-analysis pooling four randomized trials in adults with type 2 diabetes found that CGM use lowered A1C by 0.31 percentage points compared to fingerstick monitoring, a statistically significant result (Poolsup et al.). The trials included in that analysis were relatively consistent with one another, in contrast to the more mixed findings reported in the same review's pediatric type 1 diabetes data (Poolsup et al.). It is worth noting that participants in these type 2 trials had an average baseline A1C above eight percent, which suggests that the observed benefit may be strongest among patients who begin treatment poorly controlled, rather than representing a uniform effect across the broader type 2 population.

CGM's benefit is not limited to glycemic outcomes. One study included in the pooled analysis used CGM as a counseling tool paired with a physical activity program, and found that reviewing personal glucose data with patients helped them adhere to exercise goals, indicating that CGM can support behavior change independent of its role in insulin adjustment (Poolsup et al.). CGM devices equipped with low glucose suspend features, which automatically pause insulin delivery when glucose is predicted to fall below a set threshold, have been shown in a large multicenter trial to significantly reduce nocturnal hypoglycemia without increasing A1C (American Diabetes Association).

These benefits, however, come with documented barriers. CGM sensors can cause skin

reactions, including allergic contact dermatitis linked to adhesive components, and this issue has been reported consistently across age groups, meaning it represents a real physical barrier to sustained use even when the device performs well clinically (von Kobyletzki et al.). A further limitation involves the physiological lag between interstitial glucose, which CGM measures, and blood glucose. This lag can range from five to fifteen minutes and becomes most pronounced when glucose levels are changing rapidly, meaning a CGM reading may not fully reflect a patient's current blood glucose during periods of rapid rise or fall (Klonoff et al.).
Comparing CGM subtypes directly also reveals that the technology is not uniform in its effectiveness. Real time CGM, which displays data continuously as it is collected, produced significantly better outcomes than retrospective CGM, which requires downloading and reviewing data after the fact, because real time devices allow patients to act on a trend as it occurs rather than analyzing it after the opportunity to intervene has passed (Poolsup et al.). This distinction indicates that CGM's clinical benefit depends substantially on the specific device configuration a patient has access to, rather than on continuous monitoring alone.

Insulin Pump and Patch Pump Therapy

Insulin pump therapy addresses the manual dosing burden associated with MDI by allowing continuous basal insulin delivery alongside more precisely calculated bolus doses. One large scale randomized trial specifically evaluating pump therapy in type 2 diabetes, OpT2mise, found that pump therapy produced significantly greater improvement in A1C than MDI among patients with poorly controlled type 2 diabetes, and this improvement was sustained over a full year of follow up (Freckmann et al.). Patch pumps, including the V-Go, PAQ, and Simplicity, were designed specifically for type 2 diabetes and use preset basal rates with simple button activated bolus dosing rather than the carbohydrate counting and multiple basal settings typically required for type 1 diabetes management (Freckmann et al.). Studies evaluating these simplified devices have reported improved glycemic control, high patient satisfaction, and a reduction in several of the barriers that otherwise discourage patients from initiating pump therapy (Freckmann et al.). In a study comparing 100 insulin pump users with 350 patients managing type 2 diabetes without a pump, only a single pump user reported dissatisfaction, and that dissatisfaction was attributed to reservoir capacity rather than to device performance (Ghazanfar et al., 2016).

Despite these findings, the evidence for pump therapy in type 2 diabetes remains considerably less consistent than for type 1 diabetes. Because type 2 diabetes is a heterogeneous condition that varies substantially across patients in duration, insulin resistance, and baseline control, pump studies have produced mixed results overall, with some showing no advantage over MDI, others showing comparable outcomes, and others showing meaningful benefit, primarily among patients with long standing, poorly controlled disease (Freckmann et al.). This heterogeneity likely explains why, despite the positive result from OpT2mise, clinical guidelines continue to classify pump therapy as a secondary option for type 2 diabetes, may be considered for patients who do not achieve their targets on MDI rather than recommended as a first line treatment (Freckmann et al.) (American Diabetes Association). This represents a barrier related

to patient eligibility and guideline structure rather than a limitation of the device itself.

Practical barriers to pump therapy extend beyond guideline placement. Patch pumps lack a bolus calculator or connectivity to other devices, which limits their flexibility for patients whose insulin requirements fluctuate significantly (Freckmann et al.). Catheter occlusion, in which the infusion tubing becomes blocked and insulin delivery is silently interrupted, is a documented risk across both patch pumps and conventional pumps, and has been described in the infusion set literature as a persistent weakness of pump therapy generally ((Freckmann et al.). As with CGM, skin reactions related to adhesive use are a documented concern with pump therapy as well (von Kobyletzki et al.). Cost and insurance coverage present an additional barrier specific to this population, since some insurance plans do not cover pump therapy for patients with type 2 diabetes regardless of individual clinical need (Freckmann et al.). This barrier is compounded by a provider-side gap. A cross-sectional survey of 76 primary care providers found that 88% reported discomfort initiating insulin pump therapy for their patients, and 89% reported discomfort adjusting it once treatment had begun (O'Donovan et al.), a finding that parallels the provider training gap discussed in relation to MDI and suggests that clinician preparedness, not patient hesitation alone, contributes to underutilization of pump therapy in this population.

The gap between OpT2mise's clearly positive result and the more inconsistent findings reported elsewhere in the pump literature is itself worth examining directly. This inconsistency may reflect differences in the patients studied, though differences in study design and treatment protocols could also play a role (Freckmann et al.).

Conclusion

The evidence reviewed here does not point to one device as the best option for every person with type 2 diabetes. The research therefore suggests that the most important question is not whether new technology is better than old technology in general. The more useful question is which technology is most appropriate for a particular patient. Someone with well-controlled diabetes may not need a pump while someone who has struggled on MDI for years might benefit considerably from one. CGM may matter most for patients who need more insight into their patterns or who deal with hypoglycemia while others may prefer BGM's simplicity and lower cost. The overall picture is a slow shift from fully manual management toward tools that give more continuous information and automate certain parts of treatment. CGM has fairly strong evidence behind it for improving glucose control in selected adults with type 2 diabetes and pumps and patch pumps look most useful for patients who stay poorly controlled on MDI. Much of that evidence comes from trials with relatively short follow-up and patient populations that started out poorly controlled, so it remains unclear how well the same benefits hold for patients with better baseline control or over many years of continuous use. Even the most automated systems still require patients to wear and maintain devices, understand alerts, and stay engaged with realistic expectations and real training, so the technology reduces the decisions patients have to make without making their participation optional. Cost and access and usability and skin reactions and device problems and provider training still limit how far any of this reaches. Future research should keep looking past A1C at quality of life, treatment burden, hypoglycemia,

adherence, cost, and patient satisfaction. The goal of diabetes technology is to make managing the disease more accurate and sustainable and realistic for the people who have to do it every day.

Acknowledgement

I sincerely thank Indigo Research and my mentor, Dr. Mandy Rounds, for her invaluable guidance and support throughout the course of this research paper. Her expertise allowed for the successful completion of this paper.

Work Cited

American Diabetes Association. “7. Diabetes Technology: Standards of Medical Care in Diabetes—2021.” Diabetes Care, vol. 44, no. Supplement_1, Dec. 2020, pp. S85–99. Silverchair, https://doi.org/10.2337/dc21-S007.

Berget, Cari, et al. “A Clinical Overview of Insulin Pump Therapy for the Management of Diabetes: Past, Present, and Future of Intensive Therapy.” Diabetes Spectrum : A Publication of the American Diabetes Association, vol. 32, no. 3, Aug. 2019, pp. 194–204. PubMed Central, https://doi.org/10.2337/ds18-0091.

Brixner, Diana, et al. “Addressing the Burden of Multiple Daily Insulin Injections in Type 2 Diabetes with Insulin Pump Technology: A Narrative Review.” Diabetes Therapy, vol. 15, no. 7, July 2024, pp. 1525–34. Springer Link, https://doi.org/10.1007/s13300-024-01598-3.

Edelman, Steven V., et al. “Persistence with Basal-Bolus Insulin Therapy in Patients with Type 2 Diabetes Mellitus and Effect on Clinical and Economic Outcomes: A Retrospective Claims Database Study.” Journal of Managed Care & Specialty Pharmacy, vol. 25, no. 12, Dec. 2019, pp. 1420–31. jmcp.org (Atypon), https://doi.org/10.18553/jmcp.2019.19097.

Freckmann, Guido, et al. “Insulin Pump Therapy for Patients With Type 2 Diabetes Mellitus: Evidence, Current Barriers, and New Technologies.” Journal of Diabetes Science and Technology, vol. 15, no. 4, July 2021, pp. 901–15. SAGE Journals, https://doi.org/10.1177/1932296820928100.

Ghazanfar, Haider, et al. “Impact of Insulin Pump on Quality of Life of Diabetic Patients.” Indian Journal of Endocrinology and Metabolism, vol. 20, no. 4, Aug. 2016, p. 506. journals.lww.com, https://doi.org/10.4103/2230-8210.183472.

Irace, Concetta, et al. “Enhancing Type 2 Diabetes Care With CGM Integration: Insights From an Italian Expert Group.” Diabetes/Metabolism Research and Reviews, vol. 41, no. 5, 2025, p. e70059. Wiley Online Library, https://doi.org/10.1002/dmrr.70059.

Klonoff, David C., et al. “Continuous Glucose Monitoring: A Review of the Technology and Clinical Use.” Diabetes Research and Clinical Practice, vol. 133, Nov. 2017, pp. 178–92. ScienceDirect, https://doi.org/10.1016/j.diabres.2017.08.005.

Li, Yiwen, et al. “Diabetic Vascular Diseases: Molecular Mechanisms and Therapeutic Strategies.” Signal Transduction and Targeted Therapy, vol. 8, no. 1, Apr. 2023, p. 152. www.nature.com, https://doi.org/10.1038/s41392-023-01400-z.

O’Donovan, Alexander, et al. “Determining the Perception and Willingness of Primary Care Providers to Prescribe Advanced Diabetes Technologies.” Journal of Patient-Centered Research and Reviews, vol. 8, no. 3, July 2021, pp. 272–76. PubMed Central, https://doi.org/10.17294/2330-0698.1819.

Peyrot, M., et al. “Insulin Adherence Behaviours and Barriers in the Multinational Global Attitudes of Patients and Physicians in Insulin Therapy Study.” Diabetic Medicine, vol.

29, no. 5, 2012, pp. 682–89. Wiley Online Library, https://doi.org/10.1111/j.1464-5491.2012.03605.x.

Polonsky, William H., et al. “Structured Self-Monitoring of Blood Glucose Significantly Reduces A1C Levels in Poorly Controlled, Noninsulin-Treated Type 2 Diabetes: Results from the Structured Testing Program Study.” Diabetes Care, vol. 34, no. 2, Jan. 2011, pp. 262–67. Silverchair, https://doi.org/10.2337/dc10-1732.

Poolsup, Nalinee, et al. “Systematic Review and Meta-Analysis of the Effectiveness of Continuous Glucose Monitoring (CGM) on Glucose Control in Diabetes.” Diabetology & Metabolic Syndrome, vol. 5, no. 1, July 2013, p. 39. Springer Link, https://doi.org/10.1186/1758-5996-5-39.

von Kobyletzki, Laura Beate, et al. “Insulin Pump Therapy and Adverse Skin Reactions With Focus on Allergic Contact Dermatitis in Individuals Living With Diabetes Mellitus: A Systematic Review and Clinical-Based Update.” Journal of Diabetes Science and Technology, vol. 18, no. 6, Nov. 2024, pp. 1300–12. SAGE Journals, https://doi.org/10.1177/19322968241252613.

Yang, Runhuang, et al. “Insulin Delivery Devices in Diabetes Management: Applications and Advancements.” Intelligent Pharmacy, vol. 3, no. 3, June 2025, pp. 235–42. ScienceDirect, https://doi.org/10.1016/j.ipha.2025.02.002.

How Effective Is the Introduction of Regulatory Sandboxes In Helping UK Businesses Grow in the Fintech Market

By Veer S. Kooner (Student)[1] , Heitor G. C. Santos (mentor)[2]

Abstract

This paper investigates the impact of regulatory sandboxes on the growth of fintech businesses by assessing their relative gains while accounting for limitations. This was done by collecting data from the FCA website, a FOI request, KPMG pulse of fintech reports, Innovate Finances investment landscape reports, and IbisWorld. The data extracted was on the number of Financial Conduct Authority sandbox applications and those accepted, total UK Fintech investment and total number of employment within the UK Fintech sector. The data collected was used to analyse the effect of the introduction of regulatory sandboxes on FinTech businesses. To measure the effectiveness of regulatory sandboxes, the following criteria was applied: increase in investment of fintechs, market size through employment , and examples of approved firms. Results found that the FCA sandbox has had an impact on investment and employment both of which have experienced growth after the introduction of the regulatory sandbox. However, implications lie in measuring the extent of impact of the sandbox in the growth of these elements as the sandbox can provide multiple purposes: acting as a generative space by encouraging innovation or for regulatory compliance. Therefore, the ability to measure sandboxes' isolated impact is more nuanced.

Key Words

Fintech, Regulatory Sandbox, FCA, Financial Technology Application, UK Businesses

Introduction

With the growth of the fintech space, new ideas to allow for free innovation space have arrived, aiming to reduce barriers to entry within the market and increase the likelihood of success for smaller businesses. This is referring to regulatory sandboxes. "Regulatory sandboxes are experimental legal regimes which allow regulatees to conduct experiments involving innovative products or services. Existing laws and regulations often prevent firms from engaging in experiments." (Chen and Taeihagh 2026 page 1)[1]

Reasons for introducing regulatory sandboxes are that in the "financial sector regulations have been tightened to protect consumers and investors since the global financial crisis, making it difficult to determine whether to ease off these regulations. Moreover, it is not easy to know the effects of new technologies such as fintech, and whether they would have a positive effect on society or not". (Goo and Heo 2020 page 1)[2] Regarding startups, they can provide a free space for them to innovate products that could potentially facilitate the survival and growth of the company. This is because, with normal regulation, the environment is non-arable and so suits bigger corporations that have the means via purchasing power to work around it.

Dubai British School Emirates Hills, The Meadows, Dubai, 124904, United Arab Emirates
*Veerkooner@hotmail.com

The reason for the interest in the effect of sandboxes is that, if they can provide a safe space for businesses to test products in an ever competitive and growing industry, it can inspire more firms to enter the fintech market. This will lead to more competition, allowing for increased quality of products provided by the fintech market and so maximizing consumer utility. Also, it will allow for a more even spread of wealth in the fintech market, ensuring the threat of market oligopolies is mitigated.

However, sandbox regulation is not without challenges. "The need to establish a strong supervisory mechanism is important to maintain the balance between innovation and consumer protection and financial system stability. Therefore, the sustainability of sandbox effectiveness depends on close collaboration between regulators, industry players, and other stakeholders to create regulations that are adaptive and harmonious with technological developments". (Widjaja 2025 page 1)[3]. Despite the importance of understanding the impact of sandboxes, there isn't a readily available supply of studies dedicated to this. We still do not know much about whether sandboxes truly accomplish the innovation they intend to create due to their fairly recent implementation compared to traditional regulation, as well as difficulties in measuring the sandboxes' isolated impact on the improvement of innovation.

This paper investigates the impact of regulatory sandboxes on the growth of FinTech businesses within the UK. We start with an overview of the concept of sandboxes, and the impact studies have found that the Financial Conduct Authority's sandbox can result in. In the data and methods section, we explain the key factors to assess the extent of impact the FCA's sandbox has had on UK FinTech businesses, compiling data from an extensive timeline pre and post open model for employment, total investment compared against number of sandbox applications and those accepted throughout this time period. In the findings section, we analyze the correlation between growth in employment and number of sandbox accepted firms, as well as with investment. Also, we mention other potential factors that may have caused change in these indicators of growth, which helps to explain the anomaly in a huge spike of UK FinTech investment despite sandbox applications increasing at a rather steady rate. This is not to say the FCA did not play its part, as the rise in sandbox applications was substantial in proportion for the start of this anomaly.

Background

A sandbox is defined as a concrete set of frameworks which, by providing a structured context for experimentation, enable, where appropriate in a real-world environment, the testing of innovative technologies, products, services or approaches for a limited time and in a limited part of a sector or area under regulatory supervision, ensuring that appropriate safeguards are in place (Council of the European Union, 2020)[4].

In 2015, the Financial Conduct Authority launched the world's first sandbox. The reason for its introduction was to facilitate the commercialization of emerging technologies in "fintech". (Arner, Barberis and Buckley 2015)[5]. What can be inferred from this is that the fintech space revolves around innovation, and so without a regulatory sandbox, it would not be able to grow

commercially within the normal regulatory parameters. The FCA, who still run the UK regulatory sandboxes, enabled companies to perform tests every six months in the form of recruitment to 'cohorts' before the open model, which allows more flexibility in application windows. (Butor-Keler, Agnieszka, and Polasik 2020)[6], The aim of sandboxes from the FCA was to "engage with the ecosystem and encourage firms to embrace new ways of doing things in the interests of consumers". (FCA 2019)[7]

In regard to firms' perspective, sandboxes provide a way to work with the FCA to better understand whether and how a novel service could be compliant with regulation. For some firms, sandboxes are simply a testing ground for discussing potential regulatory issues during product development. For others, sandboxes are effectively an authorization process (Fahy 2022 page 145)[8]. Despite previously operating in cohorts, from 2021 onwards, the FCA follows an open model format, where firms can apply at any time to encourage more financial innovation. What has already previously been established for the effectiveness of sandboxes is, "evidence of positive externalities at the industry level. Participation in the sandbox by one start-up is followed by increased entry and more money being raised by other start-ups in the same industry. There is also some evidence that the sandbox screens out better companies" (Hellmann, Montag, Vulkan 2022 page i)[9]. This would consequently increase the quality of service provided within the fintech industry.

The reason for the focus on the UK is that it was the first country to implement a regulatory sandbox, acting as a pioneer for the use of sandboxes. Moreover, the UK is among the most prominent in regard to growth and investment as well as successful implementation of financial technology. Also, in comparison to a country like the US, despite having less growth and investment size in the fintech market, the US faces different jurisdictions with regards to sandboxes across the different states. As a result, the UK provides an easily identifiable sandbox process allowing us to measure the impact of the sandboxes on the UK fintech market, as the sandboxes are more readily available, such as through the FCA's official website.

One set of studies is more supportive of the impact of regulatory sandboxes. Kálmán and János (2025) found that evidence has shown FCA's sandbox does in fact have a positive impact on UK businesses that participate. For example, a comprehensive study found that "firms admitted to the sandbox exhibited approximately a 15% higher likelihood of securing external funding compared to their non-participating counterparts". The same study also found the sandbox increased the chances of the firms surviving and helped increase the rate of product development, "suggesting that sandbox participation can serve as a quality signal for both investors and consumers." (Kálmán and János 2025 page 3)[10] In another study (Zetsche, Buckley, Barberis, Arner 2017)[11], while reiterating the need for the Sandboxes to evolve with different challenges it may face, the positive impact sandboxes have is supported in that they allow a knowledge exchange across both parties in more detail than just information, and allow regulators to evaluate the impact on society over just personal interests, allowing for maximised utility to clients within the FinTech market.

Another set of studies is more critical. Chen (2020 page 28)[12] found that there is no evidence that they helped firms raise more finance. They also have the sense that "the regulators are awarding sandboxes for the sake of having sandboxes rather than serving real policy goals in some situations." In another study, Akwu (2024 page 8)[13] mentions that, with the disruptive nature of FinTech innovation, “not all efforts towards promoting innovation ensure a safe environment for the application of these innovative products, and by their very nature, it is extremely important that these products are kept in check”.

For these sets of studies, they often lack a longer-term perspective as they were primarily published pre 2025 and so measuring the impact now requires a more recent inference as the initial surprise element of fintech has settled since it is more widely recognized. Another use of the paper is that compiling data pre and post open model sandbox regarding market size and investment in the UK fintech proved to be challenging due to the lack of official source. As a result, the paper provides this in one place.

This study on the impact of sandboxes in the UK specifically provides utility in the sense that the FCA created the first regulatory sandbox and this led to global adoption, subsequently the UK can be viewed as the pioneer for regulatory sandboxes. As a result, the FCA regulatory sandbox has been adopted for the longest period of time which allows for a greater long-term perspective when measuring its impact and also being a key player in the Fintech industry and financial industry as a whole, with the UK being the largest net exporter of financial services signifying its global influence.

Data & Methods

To assess the extent of growth of FinTech businesses in the UK, the following criterion must be applied: the amount of investment for UK FinTech’s (Private Equity, Mergers & Acquisitions, Venture Capital), employment (jobs), and empirical evidence of participants in FCA Sandboxes. As a result, the paper focuses on analyzing the consequent impact of the FCA Sandbox on these indicators of growth for UK businesses and the FinTech market, reasons for this, while covering a broad range of perspectives from a long- and short-term lens.

The first variable the paper looks at is the number of regulatory sandboxes. Using the FCA website, I collected data and information on the number of regulatory sandbox applications and those accepted per year from pre and post open models. It was far more challenging due to the introduction of the open model. As a result, we had to find press releases for each year as the table on the website did not mention which cohort specifically, certain companies were in as well as verifying and completing a full timeline of data from 2016 to 2026 via a freedom of information request to the FCA.

Since the paper is looking at measuring the effectiveness of regulatory sandboxes on FinTech businesses in the UK, the second variable we had to collect was data on investment in UK Fintechs from 2014-2025 including Private Equity, Venture Capital, and Mergers and Acquisitions. To do so for the period 2014-2019, the KPMG Pulse of FinTech H2 report for 2019 was used where we acquired data on total UK investment. (KPMG 2020)[14] The period 2022 -

June 2025 data for FinTech investment was collected from the most recent KPMG pulse of FinTech report H1. (KPMG 2025a)[15]. We then took the KPMG 2025 H2 report for the final 2025 total investment figure (KPMG 2025b))[16]. However, to complete the timeline we needed data from 2020 and 2021. For 2020 this was done through the KPMG Pulse of Fintech 2020 report (KPMG 2021)[17]. For 2021, this was collected from Innovate Finances 2021 investment landscape report. (2021 Innovate Finance page 1)[18]

Another way we could measure the impact of regulatory sandboxes on FinTech businesses in the UK is market size. When attempting to collect figures on either revenue or total number of fintechs they were often concealed behind websites that required payments as this is not information required to be published by firms. As a result, the next best alternative was the number of fintech employment (jobs) in the UK from the period 2014-2025. We collected this from the IBISWorld bar chart on employment in UK FinTech. (IBISWorld 2025)[19].

Results

Figure 1: Number of Sandbox applications and percent of accepted applications over time. Created via google sheets.

Figure 1 shows relatively stable changes in the number of sandbox applications to the FCA until 2019, when we see a peak. After that, the number of applications returns to a baseline level until the period between April 2025 to March 2026 when we see another peak. At first, this could be interpreted as the success of the regulatory sandbox model as it becomes more familiar within the fintech space (especially if some firms become aware that other firms have tried it and seen positive results, which contributes to greater business confidence in the sandbox). However, even after the 2019 peak, although it could be interpreted that the demand to participate in the sandbox has fallen, the fact that it remained stable in terms of changes may reflect the consistency in the FCA sandbox to ensure that participating firms will benefit from the process and hence there is no drastic fall in applications while other external factors such as the emergence of Claude may have also played a part in the recent rise in applications.

When it comes to the percentage of accepted sandboxes, we see a lower percentage of accepted applications over time. Initially, in cohort 1, roughly 35% of applications were accepted. In cohort 5, when the number of pre-open model applications received is at its greatest, the FCA consequently becomes more selective, as the proportion of applications accepted is now 29%. If we then move ahead to the end of our timeline, where the greatest number of applications received is 140, the percentage of applications accepted is roughly 10%, which is 3.5 times more selective than cohort 1 and 2.9 times more selective than cohort 5. It is possible that the FCA decreased the barriers to entry at first and ended up accepting more applications. However, later, as the cumulative number of applications increased, the FCA became more selective, perhaps to show that they can uphold the standards and benefits of the sandbox to fintech firms. This directly contradicts the view of Chen (2020 page 28)[12] that, “regulators are awarding sandboxes for the sake of having sandboxes rather than serving real policy goals in

some situations." With a longer-term perspective, it is evident that the FCA have accepted applications more cautiously post 2020 and not just for the sake of awarding them, which highlights how recent evidence of the FCA regulatory sandbox can dispute more critical claims and details the focus on quality of service provided.

Figure 2: Investment in Private Equity, Mergers and Acquisitions, and Venture Capital in the UK, 2014-2025. Created via google sheets.

In regard to investment (Private Equity, M&A, Venture Capital) in the UK fintech market, Figure 2 shows that, before the Sandbox was first introduced in 2016, investment had reached the lowest point in the timeline. It is interesting to see that, post 2016, when the sandbox was introduced, investment in the UK fintech space was rejuvenated and beginning to climb towards a 2019 peak. This peak seems to be an anomaly when looking at Sandboxes because the climax was so substantial. However, this was mainly caused by the external factor of the "$42.5 billion acquisition of payments processing company Worldpay by FIS in July. Flourishing innovation ecosystems in numerous European countries, the broadening scope of fintech offerings, and growing interest from global investors also helped spark investment." (KPMG, 2019, page 54)[20]. It is important to note that in the years 2016-2018, when the sandbox was newly implemented, investment was rising at a greater rate than in the period immediately before its introduction. This shows that the sandbox likely played a part in the momentum that was building in investment until the 2019 peak. After 2019, there was a large decline due to the lack of other external factors pushing investment. However, when we look at the sandboxes' impact in 2020-2025, it remains stable and still greater than the period prior to the introduction of sandboxes.

Figure 3: Total UK Fintech Employment (jobs) per year from 2013-2025. Created on google sheets.

When it comes to employment, we also see a steady rise in employment in the entire period. When compared to investment, since the 2019 Worldpay deal involved the acquisition of one company, it is likely not to have played a significant impact on the impact in investment at a large scale. This could potentially suggest the steady rise was contributed by the greater involvement of the FCA sandbox. However, with employment, there is a significantly greater number of external factors, and while this may individually not impact figures, cumulatively, it can also pose an impact.

DISCUSSION

Sandboxes as sources of additional investment and employment boosters

Findings show that sandboxes do have a relatively pronounced impact on the two factors we aimed to measure: investment and employment. While the Worldpay deal might have been the main cause for the rise in investment, it is important to notice that the period when we see the

largest rise in accepted sandbox applications occurred from late 2017-2018 in cohorts 3-4 (11 more sandbox applications accepted relative to the prior cohort). This also marks the beginning of a rising trend in investment, which peaks in 2019. The year 2019 also happens to be the year when most applications were submitted, which could suggest firms became more willing to apply after seeing its positive benefits in the 2017-2018 period.

However, we must consider that sandboxes are not the only factor that impacts fintech investment and employment, as external factors likely play a major part such as firm confidence in the condition of the market, which impacts investment. This is not to say there is definitively not a correlation between sandbox applications and investment. Disregarding the 2018-2019 spike in investment, we find that sandboxes do have a clear and impactful purpose in encouraging longer-term steady yet sustainable growth in the fintech market. This is evident in the fact that the level of investment after the introduction of sandboxes was greater than prior. When it comes to employment, there is a stable trend in the figure rising year on year, and while the regulatory sandbox may have contributed to this for the firms that participated due to the need for more labor when introducing new product ranges, external factors are also significant when it comes to employment figures, such as mergers and takeovers. It is necessary that one is aware that the main purpose of a regulatory sandbox is not in fact in driving employment, with the measure included to potentially provide information on how the size of the market has changed due to factors like the sandbox.

Sandboxes as confidence-builders

Finally, sandboxes' impacts are not only limited to potentially boosting employment in the firms they deal with and increasing overall investment. FCA sandboxes can also act as a safety net for UK firms worrying about regulation, helping to encourage compliance. They can contribute to business confidence. For instance, the 2019 peak in investment may also have risen due to increased confidence because of the number of sandboxes accepted. This increase in confidence might have been because firms investing in the market recognized the positive impact of sandboxes in influencing innovation in the UK FinTech space, which in turn increased their confidence.

With more ability for innovation in a market heavily reliant on adapting and providing new game-changing products, the quality and variety of services received by UK clients will increase. For example, Chynge participated in Cohort 3 and the company has expanded its activities while also being supervised by the Monetary Authority of Singapore. The aim of Chynge is to make international payments instant, free and safe. Chynge uses tokens to digitalize payments on their blockchain to transfer funds immediately, create liquidity pools to eliminate bank credit lines and use AI or predictive analysis for transaction monitoring to ensure they can block suspicious patterns of money flows.

Meanwhile, NorthRow, which participated in Cohort 4, is now a technology partner for many supervised entities. (Butor-Keler and Polasik 2020)[21]. “NorthRow offers a simple solution to regulatory compliance by unifying financial crime compliance requirements into a single

environment, managing the full customer lifecycle and accelerating growth safely" (NorthRow, 2026 page 3)[22]. They do this through different products: "Know Your Customer, which verifies identity and does screening against sanction risks; Know Your Business, by validating them against public and commercial sources, AML compliance within KYB and KYC processes; Client Suitability which uses governance tools to check suitability, Monitoring, Analytics, Remote ID & V using liveness video checks and biometric facial recognition; and finally Right to Work, combining automatic MRZ validation with NFC technology". (NorthRow, 2026, page 2)[22]. NorthRow was acquired by ID-Pal in December 2025 to "deliver a unified compliance intelligence solution for Know Your Customer (KYC), KYB, and Anti-Money Laundering (AML)". (ID Pal 2025 page 1)[23].

Sandboxes as generative spaces

In regard to increasing variety of services, the effect of the FCA sandbox can be conceived as a generative space which can be defined as" providing distinct and varied social spaces separated from normal regulatory demands (protection); enabling participants to collabo-ratively examine innovations and iteratively develop new knowledge and regulation (adaptation); and supporting focused and immersive examination of innovations from a diverse range of perspectives (participation)." (Macrae and Ansell 2024 page 2)[24]. This is because, "the sandbox has accelerated the development of financial innovation and deepened the cooperation of large entities with start-ups by approx. 50%". (Butor-Keler and Polasik 2020 page 629)[21]. Those participating in the sandboxes found new innovative solutions meaning that the FCA's sandbox can act as a generative space in developing new products/services to the FinTech market. For example, from the FCA sandbox these new solutions include, "intelligent sensors in flood-prone areas for insurance purposes, real-time financial market monitoring, or guaranteed loans when selling a flat". (Butor Keler and Polasik 2020 page 630)[21] Furthermore, those tested in the sandbox provided the aforementioned generative space to test products that otherwise would have been risky to introduce directly into the market. Furthermore, "the authors carried out an analysis of individual participants of the regulatory sandbox, and it showed that no two solutions were considerably similar to each other". (Butor-Keler and Polasik 2020 page 630)[21]. This reveals the level of uniqueness brought by innovation within the FCA sandbox which will only benefit the UK FinTech market as successful technological based markets thrive on new products being introduced to maintain high levels of demand.

CONCLUSION

This paper had the goal of investigating the impact of regulatory sandboxes in the U.K, looking at how they impacted UK Fintech businesses through two key factors: investment and employment. Collecting data from FCA website for the sandbox figures, KPMG pulse of fintech, FsTech and Innovate Finance reports for total FinTech Investment and the IBISWorld chart for employment, we found that it is more complicated to definitively reach a judgement on its

impact due to the nature of the FinTech market and many other external factors at play such as the $42.5 billion acquisition of WorldPay by FIS which did contribute significantly to the highest peak in investment that otherwise would have been relatively more stable and consistent. However, this is not to say stability is a negative thing in the investment levels of FinTech within the UK, as a regulatory sandbox may have contributed to consistent levels of interest in developing the FinTech market through injecting capital where entrepreneurship and innovation are incentivized with growth. Overall, the paper's contribution is in finding that the FCA's Sandboxes have positively influenced firms and the UK FinTech market. The positive direct impact FCA Sandboxes have had is seen with their ability to create a generative space, encouraging innovation and creativity that results in a level of uniqueness, improving variety in the market. Furthermore, the direct impact sandboxes have had is seen with participants like NorthRow testing their products before they go to market at a larger scale, such as KYC which has led to the company being acquired by ID-pal as of 2025.

Furthermore, the paper's contribution is also in compiling data into one space for a complete timeline of 2013-2025 for employment in the UK FinTech space and 2014-2025 for total investment, acting as a complete timeline from when the FCA regulatory sandbox was introduced pre and post open model for regulatory sandbox applications and accepted applicants covering a decade-long timeline.

However, the challenge is measuring the extent of impact of FCA sandboxes on these direct and indirect impacts when looking at the market rather than specific FinTech firms. This is because other factors also contribute to diluting the sole influence of just sandboxes. As a result, this is a gap that other scholars could investigate to take the investigation and assessment forward. Although overall using the criteria established to assess the impact of regulatory sandboxes, the FCA sandboxes have had a positive impact on UK FinTech businesses as the following criterion: the amount of investment for UK FinTech's (PE, M&A, VC), employment (jobs), and empirical evidence of participants in FCA Sandboxes have all seen positive results from the introduction of the FCA sandbox in 2016-2025 because there has been a sustainable rise in UK FinTech employment and investment.

Finally, in regard to what future research should entail and investigate, it is necessary to delve deeper into the extent of impact of the regulatory sandbox on the criteria established regarding specific examples, total UK Fintech investment and total UK Fintech employment. This is because the regulatory sandbox has evidently impacted investment and employment with both figures being higher than before the sandbox introduction, as well as a range of qualitative factors such as ensuring regulatory compliance and acting as a generative space fostering innovation subsequently highlighting how the sandbox has contributed to the growth of UK Fintech businesses. However, through greater research into other factors and comparing their impact, it could allow for a more concrete framework to define the FCA regulatory sandbox as a resounding success with the resources allocated for the FCA regulatory sandbox potentially still relatively being subject to an opportunity cost if the impact is less pronounced. Those that

contributed to the largest rise in investment particularly could be questioned as to whether the regulatory sandbox did in fact indirectly or directly shape their decision to invest.

AI Acknowledgement

During the process of this work the author used Google Gemini and ChatGPT to try to pinpoint sources/where to locate specific data. After using this tool the author reviewed and edited the resulting content as necessary and took full responsibility for the content of the publication.

Author

Veer Kooner is a year 12 student interested in finance and its applications with modern technology. Taking Maths, Economics and History for A levels he partook in a 4 course fintech series by Upenn where he found his interest in Fintech developed in activities like a stripe internship.

Works Cited

Chen, G. K. H. and Taeihagh, A. (2026) 'Designing regulatory sandboxes: a comprehensive framework for aligning functionalities and objectives', Policy Design and Practice, 9(1), pp. 1–15. doi: 10.1080/25741292.2025.2570954.

Goo, J.J. and Heo, J.-Y. (2020). The Impact of the Regulatory Sandbox on the Fintech Industry, with a Discussion on the Relation between Regulatory Sandboxes and Open Innovation. Journal of Open Innovation: Technology, Market, and Complexity, 6(2), p.43.

Gunawan, Widjaja (2025) ANALYSING THE IMPACT OF SANDBOX REGULATION ON FINTECH INNOVATION AND FINANCIAL INCLUSION. Injoser, 3 (5).

Council of the European Union (2020) Council Conclusions on Regulatory Sandboxes and Experimentation Clauses as Tools for an Innovation-Friendly, Future-Proof and Resilient Regulatory Framework that Masters Disruptive Challenges in the Digital Age. Brussels, 16 November. Document 13026/20. Available at: data.consilium.europa.eu/doc/document/ST-13026-2020-INIT/en/pdf (Accessed: 28 March 2026).

Arner, D.W., Barberis, J. and Buckley, R.P., 2015. The evolution of Fintech: A new post-crisis paradigm. Geo. J. Int'l L., 47, p.1271.

Butor-Keler, A. and Polasik, M. (2020) "The role of regulatory sandboxes in the development of innovations on the financial services market: the case of the United Kingdom", Ekonomia i Prawo. Economics and Law, 19(4), pp. 621–638. doi: 10.12775/EiP.2020.041.

"From Innovation Hub to Innovation Culture." 2019. FCA. June 4, 2019. https://www.fca.org.uk/news/speeches/innovation-hub-innovation-culture.

FAHY, L. (2022) 'Regulator Reputation and Stakeholder Participation: A Case Study of the UK's Regulatory Sandbox for Fintech', European Journal of Risk Regulation, 13(1), pp. 138–157. doi:10.1017/err.2021.44.

Hellmann, T.F., Montag, A. and Vulkan, N. (2022). The Impact of the Regulatory Sandbox on the FinTech Industry. SSRN Electronic Journal. doi:https://doi.org/10.2139/ssrn.4187295.

Kálmán, János. 2025. "The Role of Regulatory Sandboxes in FinTech Innovation: A Comparative Case Study of the UK, Singapore, and Hungary" FinTech 4, no. 2: 26. https://doi.org/10.3390/fintech4020026

Zetzsche, D.A.; Buckley, R.P.; Arner, D.W.; Barberis, J.N. Regulating a Revolution: From Regulatory Sandboxes to Smart Regulation. Fordham J. Corp. Finance. Law 2017, 23, 31–103. [Google Scholar] [CrossRef]

CHEN, Christopher. Rethinking the regulatory sandbox for financial innovation: An assessment of the UK and Singapore. (2020). Regulating FinTech in Asia: Global context, local perspectives. 11-30.

Akwu, Patrick. "Sandbox or Quick Sand? Evaluating the Double-Edged Sword of Fintech Regulation." Evaluating the Double-Edged Sword of Fintech Regulation (May 01, 2024) (2024).

“Pul Se of Fi Ntech H2 2019.” 2020. https://assets.kpmg.com/content/dam/kpmg/xx/pdf/2020/02/pulse-of-fintech-h2-2019.pdf.

“Pulse of Fintech – UK Perspective.” 2025. KPMG. 2025a. https://kpmg.com/uk/en/insights/fintech/pulse-of-fintech-uk-perspective.html.

“Pulse of Fintech H2’2025: EMEA.” 2025. KPMG. 2025.b https://kpmg.com/xx/en/what-we-do/industries/financial-services/pulse-of-fintech/emea.html.

“Pul Se of Fi Ntech H2’ 20.” 2021. https://assets.kpmg.com/content/dam/kpmg/xx/pdf/2021/02/pulse-of-fintech-h2-2020.pdf

“FinTech Investment Landscape 2021.” n.d. Innovate Finance – the Voice of Global FinTech. https://www.innovatefinance.com/capital/a-record-breaking-year-in-fintech/.

IBISWorld. 2026. “Financial Technology in the UK.” Ibisworld.com. 2026. https://www.ibisworld.com/united-kingdom/employment/financial-technology/14450/.

NorthRow, Accelerate Growth, Safely, corporate brochure (London: NorthRow, n.d.), PDF, https://25597772.fs1.hubspotusercontent-eu1.net/hubfs/25597772/Brochures/Corporate%20Brochure.pdf.

“ID-Pal Announces Strategic Acquisition of KYB Specialists, NorthRow - ID-Pal.” 2025. ID-Pal. December 2025. https://www.id-pal.com/newsroom/northrow-acquisition/.

Macrae, Carl, and Christopher K. Ansell. "Generative spaces: Collaboration, learning and innovation in a regulatory sandbox." Learning and Innovation in a Regulatory Sandbox (May 13, 2024) (2024).

A PLS-SEM Based Analysis of Determinants Influencing Portfolio Construction and Risk Minimization Strategies By Atiksh Suneja

Abstract

This study investigates the determinants of portfolio construction efficiency and risk minimization as perceived by investment practitioners. It examines how behavioral, macroeconomic, and diversification determinants relate to perceived portfolio construction efficiency and perceived risk minimization outcomes. Using a positivist, deductive approach and Partial Least Squares Structural Equation Modelling (PLS-SEM), a distribution-free estimation technique, the study collected primary data from structured questionnaires completed by 101 financial professionals and investment bankers selected for their direct experience in constructing portfolios and developing risk strategies. The questionnaire was confirmed to be reliable, with Cronbach's alpha coefficients between 0.736 and 0.806 and average variance extracted values ranging from 0.484 to 0.540. Heterotrait-Monotrait (HTMT) ratios for all construct pairs were below 0.85, indicating good discriminant validity. Portfolio determinants explained 43% of the variance in perceived portfolio construction efficiency ($R^2 = 0.430$) and 33% of the variance in perceived risk minimization outcomes ($R^2 = 0.330$), with holistic risk governance and risk ownership clarity emerging as the strongest indicators of risk minimization. Because all constructs are self-reported and overall model fit is weak, the findings are exploratory and describe practitioner perceptions rather than realized portfolio performance.

Introduction

Since the crystallization of modern portfolio theory, diversification has remained a key aspect in managing the risk-return trade-off in investment practice (Markowitz). Investment funds have become a central mechanism for global capital allocation, with an estimated 140,000 open-end funds operating worldwide as of 2023, enabling investors with varying risk preferences to access diversified financial markets without specialized expertise ("Topic: Investment Funds"). However, varying levels of market volatility, tail risks, and recurring systemic shocks have been crucial indicators in exposing the limitations of traditional frameworks, such as mean-variance frameworks. As a result, factor-based portfolio construction has evolved as an ex-ante effective, more transparent, and cost-effective approach to risk management in comparison to conventional security-level allocation. Although individual risk factors are able to provide positive premia over a given time period, they remain vulnerable to significant drawdowns, especially during periods of economic stress. Additionally, cross-regional factor diversification exploits low correlation structures in order to mitigate overall volatility and extreme downside risk, with long-only global minimum variance strategies being identified as robust approaches to risk minimization (Kremer et al.).

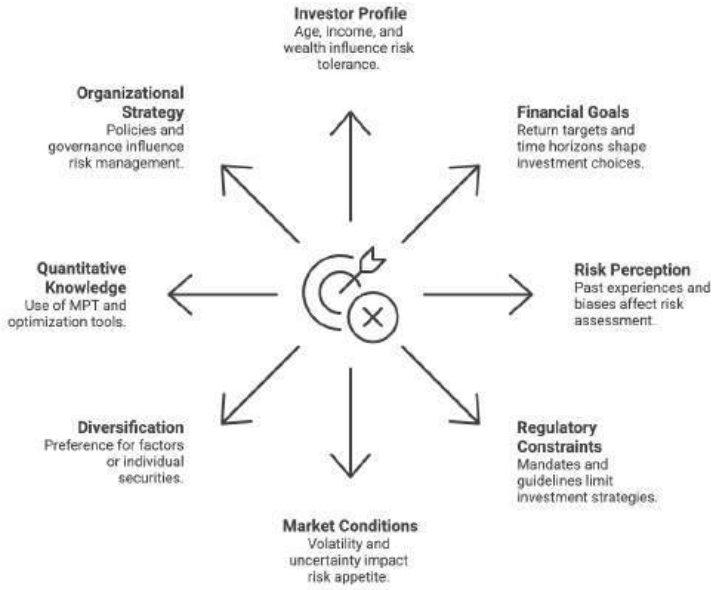

Fig 1: Determinants of Risk Appetite

Since Markowitz's foundational 1952 work, portfolio construction, coupled with risk management, has evolved considerably through novel quantitative methods and innovative evaluation frameworks. Rom and Ferguson, Roman and Mitra, and Araújo and Montini (qtd. in Lee and Eid) argue that the use of either variance or standard deviation as a measurement of risk is insufficient, particularly in the presence of asymmetric returns and tail risk. Similarly, Chan et al. and Santos and Tessari (qtd. in Lee and Eid) question the reliability of the sample covariance matrix due to its sensitivity to estimation error. Moreover, even portfolio optimization techniques with a heavy quantitative aspect are being criticized, with DeMiguel et al. portraying their limited robustness in practice. Despite extensive academic advances, high levels of uncertainty remain, especially regarding the extent to which these methods are adopted by market practitioners. Empirical evidence from Brazilian asset managers indicates that traditional mean-variance approaches are still preferred over post-modern portfolio theory, and that simplified rules such as concentration limits and normal distribution assumptions for risk estimation are preferred over more sophisticated models incorporating skewness and kurtosis (Lee and Eid).

The importance of portfolio construction and risk minimization has become a major concern for boards of directors and regulators since the financial crisis. Portfolio investment strategy has been influenced by enhanced prudential regulations and risk-based capital requirements (Basel Committee on Banking Supervision, *Basel III*). Hedge funds have typically been viewed as alternative investments that can provide absolute returns regardless of market conditions; however, that view was fundamentally challenged during the 2008 global financial crisis, which exposed many of the limitations of hedge funds (Roumpis and Syriopoulos). For example, Hedge Fund Research recorded a record 1,471 hedge fund liquidations in 2008, followed by a further 668 in the first half of 2009 (Parmee).

These events raised serious questions as to the effectiveness of traditional portfolio theories, particularly those based on static measures of risk and return such as the mean-variance framework and models based on capital asset pricing. In addition, prior research has demonstrated that the returns of hedge funds are dynamic and exhibit non-linear characteristics, making them difficult to analyze using traditional asset classes (Fung and Hsieh). The concern for portfolio risk management is especially pronounced in the U.S., where almost 50% of the assets in all regulated open-end funds are domiciled and where more than 50% of U.S. households owned mutual funds in 2023; therefore, firm-level risk arising from errors in investor behavior could significantly impact the stability of financial markets ("Topic: Investment Funds"). As such, while the pursuit of effective ways to construct portfolios and reduce risk has remained an important issue since 2008, there remains a lack of a strong theoretical or empirical basis for these efforts (Roumpis and Syriopoulos).

Since the 2008 financial crisis, factor investing has become increasingly attractive as a cost-efficient alternative to traditional active management. Factor-based portfolios offer effective diversification while also enabling the capture of systematic risk premia, leading to widespread adoption of smart beta and multifactor strategies within equity markets. Recent empirical evidence illustrates the relevance of risk-based asset allocation within multifactor investing. In particular, factor-risk parity approaches have been highly effective across major equity markets by allocating risk evenly across both cyclical and defensive factors. These findings position factor-based portfolio construction as an effective method of construction and as a robust framework for disciplined, risk-aware investment (Shimizu and Shiohama).

Risk management has evolved significantly over the years, becoming especially critical in periods of heightened uncertainty and enabling organizations and individuals alike to remain resilient and avoid suboptimal decision-making. Global disruptions such as the COVID-19 pandemic have exposed the limitations of existing risk frameworks and intensified calls for more sustainable economic and organizational models. This context has heightened awareness of the need to integrate risk considerations with the three pillars of sustainability (environmental, social, and economic) and to ensure long-term objectives are met. However, the theoretical relationship between risk management systems and sustainability management remains only partially understood, as it is multifaceted. Recent research therefore emphasizes the need for integrated frameworks that link multidimensional risk assessment to sustainable performance outcomes (Settembre-Blundo et al.).

The remainder of the paper is organized as follows. The Literature Review surveys research on risk factors, portfolio construction, and risk minimization strategies. The Research Methodology section sets out the conceptual framework, sampling strategy, and analytical procedure. The Results section reports the measurement and structural model assessments and answers the two research questions. The Discussion interprets these findings, and the paper closes with its contributions and implications, limitations and directions for future research, and a conclusion.

Literature Review

Portfolio construction and realized risk outcomes are influenced by a complex set of interacting determinants, including but not limited to investor characteristics, market microstructure, information frictions, behavioral biases, regulatory constraints, and analytical capabilities. Existing research evaluates portfolio selection models with distinct objectives (risk minimization, capital diversification, and risk diversification), highlighting the growing importance of risk parity or equal risk contribution approaches. While traditional risk parity frameworks rely mainly on volatility as the primary risk measure, recent extensions of risk parity frameworks aim to incorporate Conditional Value-at-Risk to better capture downside risk. Comparative evidence across equity, bond, and mixed-asset universes shows that each portfolio strategy exhibits specific strengths and limitations depending on the source of risk. These insights encourage a broader examination of how structural and behavioral determinants jointly influence portfolio decisions and risk minimization strategies (Cesarone and Colucci).

Project portfolios are widely recognized as strategic instruments for implementing corporate strategy, yet they are exposed to multiple sources of risk. Existing studies in project portfolio optimization have primarily addressed risk by maximizing expected returns (max E[R] or max E[NPV]) or by imposing upper-bound constraints on aggregate portfolio risk. However, these approaches largely neglect the explicit minimization of extreme downside outcomes. Recent research addresses this gap by adopting risk-averse decision criteria based on Conditional Value-at-Risk (CVaR). CVaR-based portfolio models demonstrate greater robustness under adverse scenarios while supporting diversification and enabling portfolio selection aligned with decision-makers' risk appetite (Dixit and Tiwari).

Related work examines gap risk in dynamic portfolio insurance strategies that extend the Constant Proportion Portfolio Insurance (CPPI) framework by allowing the leverage multiple to vary over time. That analysis focuses on conditional CPPI strategies applied to hedge fund portfolios, with return dynamics modelled using Johnson distributions and an EGARCH specification with Johnson innovations. Downside risk is explicitly controlled using Value-at-Risk (VaR) and Expected Shortfall (ES) measures, and upper bounds on the portfolio multiple are derived to limit gap risk under adverse market conditions, with empirical illustrations drawn from the Credit Suisse Hedge Fund Index (Naguez).

Mean-variance portfolio optimization has been widely criticized for poor out-of-sample performance and the generation of extreme and unstable asset weights, particularly in high-dimensional settings. These shortcomings largely stem from estimation errors in key inputs, notably the expected return vector and the covariance matrix. Recent studies demonstrate that imposing an ℓ_1-norm penalty on portfolio weights (ℓ_1-regularization) can stabilize solutions and improve out-of-sample performance by inducing sparsity. However, ℓ_1-regularization may also introduce biased portfolio allocations. To address this limitation, alternative penalty structures that incorporate financial information have been proposed, yielding portfolios that outperform standard benchmarks in empirical tests using U.S. equity data (Fastrich et al.).

Other work examines portfolio selection and risk management in a distribution-free setting based on Cover's universal portfolio framework (Cover). Rather than relying on explicit market modelling or known return distributions, this approach adopts an online learning perspective in which portfolio decisions are updated sequentially. Under a minimax regret criterion, risk-adjusted algorithms are proposed to track experts exhibiting the lowest maximum drawdown. Theoretical analysis establishes worst-case performance bounds consistent with those derived in Cover's seminal work, and empirical evidence using long-horizon NYSE data demonstrates the superior performance of strategies that explicitly incorporate drawdown-based risk control (Mohr and Dochow).

Conditional Value-at-Risk (CVaR) is a widely used portfolio risk measure due to its desirable properties relative to other risk measures, such as convexity and sub-additivity. Although problems of minimizing CVaR can be expressed as linear programming problems, they face a major hurdle when dealing with large asset universes and large sets of scenarios due to computational complexity. In response to this drawback, advanced techniques have been proposed for developing a multi-phase optimization framework that combines differentiable optimization approaches with convergent non-differentiable ones to produce nearly optimal results in an efficient manner while ensuring convergence to optimality. Empirical benchmarks demonstrate significant performance improvements over standard commercial solvers, particularly in high-dimensional settings (Lim et al.).

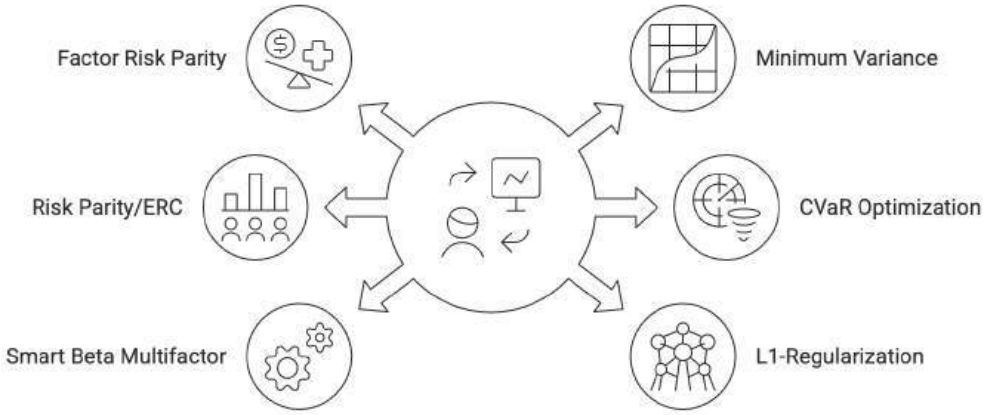

Fig 2: Portfolio Construction Techniques

Portfolio construction is the process of building a portfolio that produces an optimal combination of risk and return by allocating or balancing assets, and it requires the use of different principles to make investment decisions. One of the main principles of asset allocation is diversification; however, issues such as transaction costs, estimation errors, and high-dimensional asset spaces led to the creation of sparse, regularized, and robust portfolio selection frameworks. Additionally, the complexity of capital markets has prompted the development of learning-based stochastic control methods for adaptive portfolio decision-making. As portfolio optimization has developed through various methods, from

traditional mean-variance techniques to advanced downside risk measures and heuristics, portfolio management has become highly complex and requires a more integrative, data-driven approach to managing risk (Gunjan and Bhattacharyya).

Research Methodology

Using a deductive epistemological framework, this study takes a positivist, quantitative research approach. The investigation implements Partial Least Squares Structural Equation Modelling (PLS-SEM), following the recommendations of Hair et al. and Ringle et al., as a distribution-free estimation approach to analyze latent construct relationships. Three reasons for selecting PLS-SEM over covariance-based SEM (CB-SEM) are that (a) it can be used with small to medium-sized samples; (b) it can model complex structural relationships without making distributional assumptions; and (c) it produces accurate predictive results in exploratory settings that use concurrently operationalized theoretical constructs. A cross-sectional survey design was used to collect primary data from a purposively sampled group of respondents with direct experience constructing portfolios, developing investment strategies, and managing risks. This design was used to test the structural pathways between the determinants of the portfolio construction process and the minimization of risk as defined in the conceptual framework.

Conceptual Framework and Construct Operationalization

The structural model comprises three principal latent constructs derived from theoretical synthesis across behavioral finance, macroeconomic portfolio theory, and risk governance literature:

R1 (Portfolio Determinants): This reflective construct measures multiple factors that influence investor portfolio choices, such as behavioral biases, including overconfidence (Barber and Odean), loss aversion (Kahneman and Tversky), and anchoring (Tversky and Kahneman), sensitivity to broader economic conditions (interest rates, inflation expectations, geopolitical events), and willingness to diversify holdings. It was operationalized using five Likert items (Items 1–5), rated from "Strongly Disagree" (1) to "Strongly Agree" (5).

R2 (Perceived Portfolio Construction Efficiency, PCE): This endogenous construct captures respondents' assessments of how well their portfolios are constructed, looking primarily at the asset weightings applied within the portfolio (i.e., whether the investor successfully applied a performance-risk trade-off), and assessing whether the constructed portfolio meets GMV-type and universal portfolio designs while being optimally allocated on risk-based principles. It was measured using Items 6–10.

R3 (Perceived Risk Minimization Outcomes, RMO): This endogenous construct captures respondents' assessments of how well their risk management framework reduces risk, covering overall risk governance and the identification of risk owners, as well as how well downside risk was contained under stress, how variability was reduced during stressful periods, and how robust performance was against a distribution-free model. It was operationalized using Items 11–15.

All three constructs are measured through respondents' self-reported perceptions. The study therefore examines how practitioners assess portfolio construction efficiency and risk minimization, rather than realized portfolio weights, returns, or drawdowns.

The model specifies two structural paths, corresponding to two hypotheses:

H1: Portfolio determinants are positively associated with perceived portfolio construction efficiency (R1 → R2).

H2: Portfolio determinants are positively associated with perceived risk minimization outcomes (R1 → R3).

No mediation or moderation paths were specified in this iteration of the model.

Population, Sampling Strategy, and Data Collection

The target group consisted of investment professionals, financial analysts, portfolio managers, academic researchers in finance, and informed individual investors. To ensure that respondents had relevant experience with the topics included in the survey, a structured questionnaire was created and distributed using purposive and snowball sampling methods. One hundred and one usable responses were retained after the data were screened for completeness and response consistency. This sample meets the minimum requirement for PLS-SEM (Hair et al.) under the 10-times rule, which estimates that the sample should be at least ten times the largest number of formative indicators or structural paths leading to any endogenous construct. The questionnaire used a 5-point Likert scale (1 = Strongly Disagree to 5 = Strongly Agree) and consisted of 15 reflective indicator items spread over three constructs. Content validity was established through expert review, followed by a pilot test with 10 individuals before large-scale data collection.

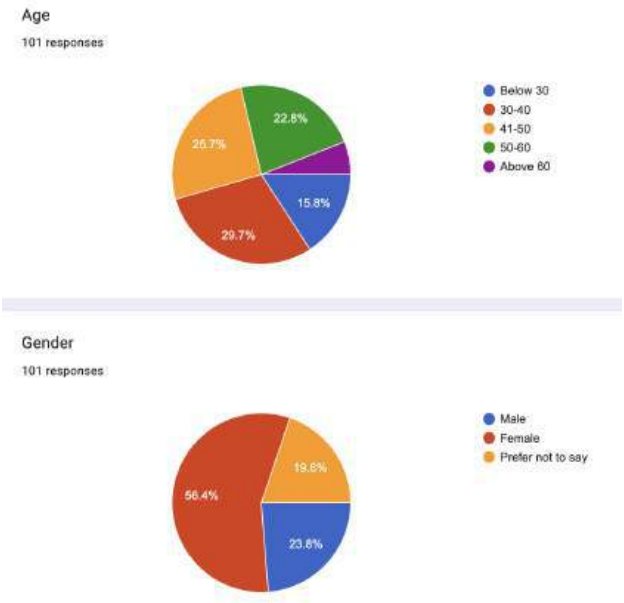

Fig 3: Respondent Demographics

Analytical Procedure

SmartPLS 4 (Ringle et al.) was used to execute all analyses. The analytical strategy adopted the two-stage procedure proposed by Anderson and Gerbing, which was implemented in PLS-SEM by Hair et al.

Step 1, Measurement Model Evaluation: The reliability of the reflective construct was determined on the basis of Cronbach's alpha (α), composite reliability (ρ), and convergent validity through Average Variance Extracted (AVE). Discriminant validity was evaluated using Heterotrait-Monotrait (HTMT) ratios and the Fornell-Larcker criterion. Variance Inflation Factor (VIF) statistics were checked to ensure there were no collinearity issues among the indicators.

Step 2, Structural Model Evaluation: To assess the explanatory and predictive capabilities of the structural model, path coefficients (β), the coefficient of determination (R^2), and Cohen's effect size (f^2) were calculated.

Results

This section presents the findings of the PLS-SEM analysis in two stages: the measurement model assessment, which confirms construct validity and reliability, and the structural model assessment, which evaluates the hypothesized relationships and provides empirical answers to both research questions.

Indicator Reliability and Outer Loadings

Indicator reliability was assessed by reviewing the outer loadings of all 15 reflective indicators across the three constructs. Loadings of ≥ 0.60 are considered acceptable for exploratory research (Hair et al.), whereas a loading of ≥ 0.70 is the norm for confirmatory studies. The full list of outer loadings is shown in Table 1. Across all three constructs, outer loadings ranged from Item 3 ($\lambda = 0.621$) to Item 12 ($\lambda = 0.791$), with most at or above 0.65. The lowest loading (Item 3, $\lambda = 0.621$) is above the minimum acceptable value of 0.60 and is retained because of its theoretical relevance to the construct domain; all other loadings also exceed this minimum, indicating that the indicators meet the required levels of reliability.

Table 1: Outer Loadings, Cronbach's Alpha, and AVE by Construct

Construct	Indicator	Outer Loading	Cronbach's α	AVE
R1: Portfolio Determinants	Item 1: Behavioral bias recognition in asset selection	0.729	0.736	0.484
	Item 2: Macroeconomic sensitivity in allocation	0.742		
	Item 3: Diversification awareness	0.621		
	Item 4: Risk tolerance framing	0.718		
	Item 5: Portfolio rebalancing intent	0.661		

R2: Perceived Portfolio Construction Efficiency	Item 6: Efficiency of asset weighting	0.732	0.769	0.490
	Item 7: Performance-risk trade-off clarity	0.681		
	Item 8: Global minimum variance alignment	0.650		
	Item 9: Universal portfolio compatibility	0.689		
	Item 10: Risk-based allocation optimality	0.742		
R3: Perceived Risk Minimization Outcomes	Item 11: Holistic risk governance effectiveness	0.788	0.806	0.540
	Item 12: Risk ownership clarity	0.791		
	Item 13: Downside risk containment	0.722		
	Item 14: Variance reduction under market stress	0.678		
	Item 15: Distribution-free risk model robustness	0.685		

Source: SmartPLS 4 output (Ringle et al.).

Construct Reliability and Convergent Validity

Construct reliability, calculated using Cronbach's alpha (α), composite reliability (rho_a), and composite reliability (rho_c), exceeded 0.70 for all constructs (Hair et al.), as shown in Table 2.

For R1 (Portfolio Determinants), α = 0.736, CR (rho_a) = 0.738, and CR (rho_c) = 0.824, indicating that the minimum reliability threshold is achieved and demonstrating acceptable internal consistency. According to Nunnally, Cronbach's alpha values between 0.70 and 0.80 are adequate for exploratory constructs.

For R2 (Portfolio Construction Efficiency), α = 0.769, CR (rho_a) = 0.778, and CR (rho_c) = 0.827; composite reliability values between 0.80 and 0.90 indicate good reliability, suggesting strong internal consistency for the five PCE indicators.

Lastly, with α = 0.806, CR (rho_a) = 0.837, and CR (rho_c) = 0.854, R3 (Risk Minimization Outcomes) has the strongest reliability profile of the three constructs, with a Cronbach's alpha above 0.80 confirming high internal consistency for the holistic risk governance and risk ownership indicators.

Convergent validity was assessed through AVE values. For R1, the AVE is 0.484, slightly below the common 0.50 threshold, but it can be considered acceptable when accompanied by strong composite reliability (Fornell and Larcker; Hair et al.). Each item loaded above 0.60. R2 has an AVE of 0.490, close to the benchmark, and R3 has an AVE of 0.540, which exceeds it, confirming full convergence.

Table 2: Construct Reliability and Convergent Validity Summary

Construct	Cronbach's α	CR (rho_a)	CR (rho_c)	AVE	√AVE
R1: Portfolio Determinants	0.736	0.738	0.824	0.484	0.696
R2: Perceived Portfolio Construction Efficiency	0.769	0.778	0.827	0.490	0.700
R3: Perceived Risk Minimization Outcomes	0.806	0.837	0.854	0.540	0.735
Threshold (acceptable)	> 0.70	> 0.70	> 0.70	> 0.50	—

Note: CR = Composite Reliability; AVE = Average Variance Extracted; √AVE = square root of AVE (used for the Fornell-Larcker criterion).

Discriminant Validity

The Heterotrait-Monotrait ratio (HTMT) was used because it has been established as a better way to assess discriminant validity, or how distinct constructs are from each other, than the traditional Fornell-Larcker criterion (Henseler et al., "New Criterion"). The conservative HTMT criterion of 0.85 was applied.

Table 3: Discriminant Validity (HTMT Ratio Matrix)

Construct Pair	HTMT	Threshold	Status
R1 ↔ R2 (Portfolio Determinants ↔ PCE)	0.739	< 0.85	Pass
R1 ↔ R3 (Portfolio Determinants ↔ Risk Min.)	0.669	< 0.85	Pass
R2 ↔ R3 (PCE ↔ Risk Minimization)	0.639	< 0.85	Pass

Note: All HTMT values are below 0.85, confirming discriminant validity for all construct pairs.

All pairwise HTMT values were below the 0.85 threshold (R1 ↔ R2 = 0.739; R1 ↔ R3 = 0.669; R2 ↔ R3 = 0.639), confirming that the three constructs are empirically distinct. This finding is strengthened by the conceptual closeness of portfolio construction efficiency (R2) and risk minimization outcomes (R3), indicating that these constructs measure two different dimensions of investment performance. In addition, the Fornell-Larcker criterion was verified, with the square root of each construct's AVE exceeding its correlations with the other constructs.

Structural Model Assessment and Hypothesis Testing

Following satisfactory measurement model assessment, the structural model was evaluated to test the two hypothesized relationships. Table 4 summarizes the path coefficients (β), R^2 values, and Cohen's f^2 effect sizes, along with hypothesis outcomes.

Table 4: Structural Path Coefficients, R^2, f^2, and Hypothesis Testing Summary

Hyp.	Path	β	R^2	f^2	Effect Size	Result
H1	R1 (Portfolio Determinants) → R2 (Perceived Portfolio Construction Efficiency)	0.656	0.430	0.755	Large	Supported
H2	R1 (Portfolio Determinants) → R3 (Perceived Risk Minimization Outcomes)	0.575	0.330	0.494	Large	Supported

Note: $f^2 > 0.35$ = large effect (Cohen). Both paths significant at $p < 0.01$ based on bootstrapping (5,000 subsamples).

RQ1: Portfolio Determinants and Perceived Construction Efficiency

RQ1: How are behavioral, macroeconomic, and diversification determinants associated with practitioners' perceived portfolio construction efficiency?

The structural path from R1 (Portfolio Determinants) to R2 (Perceived Portfolio Construction Efficiency) produced a path coefficient of $\beta = 0.656$ ($p < 0.01$), indicating a strong positive association. The coefficient of determination for R2 was $R^2 = 0.430$ (adjusted $R^2 = 0.424$), meaning that portfolio determinants explain 43% of the variance in perceived construction efficiency. Following Cohen and Hair et al., this represents a moderate-to-substantial level of explained variance in social science research, and the f^2 effect size of 0.755 exceeds the 0.35 threshold for a large effect.

These results indicate that practitioners who report greater awareness of behavioral biases, greater sensitivity to macroeconomic conditions, and a stronger diversification orientation also report constructing portfolios more efficiently across the GMV-type, risk-based, and universal portfolio considerations captured in Items 6–10. Because these items record practitioners' assessments rather than realized portfolio weights or returns, the finding describes how determinants relate to perceived efficiency, not measured performance-risk trade-offs. The remaining 57% of the variance may reflect factors outside the model, such as liquidity constraints, transaction costs, or institutional mandates.

H1 supported: Portfolio determinants are positively and significantly associated with perceived portfolio construction efficiency ($\beta = 0.656$, $R^2 = 0.430$, $f^2 = 0.755$).

RQ2: Portfolio Determinants and Perceived Risk Minimization

RQ2: To what extent do portfolio determinants explain practitioners' perceived risk minimization outcomes, and which aspects of risk minimization are most salient?

The path from R1 to R3 (Perceived Risk Minimization Outcomes) produced a path coefficient of $\beta = 0.575$ ($p < 0.01$) and $R^2 = 0.330$ (adjusted $R^2 = 0.324$), indicating that portfolio determinants explain 33% of the variance in perceived risk minimization outcomes. The f^2 value of 0.494 indicates a large effect.

Within R3, the two highest-loading indicators are Holistic Risk Governance Effectiveness (Item 11, $\lambda = 0.788$) and Risk Ownership Clarity (Item 12, $\lambda = 0.791$), ahead of Downside Risk

Containment (Item 13), Variance Reduction under Market Stress (Item 14), and Distribution-Free Risk Model Robustness (Item 15). Because governance and ownership are indicators of the risk minimization construct rather than separate predictors, this result does not show that governance improves risk outcomes or the model's predictive power. It shows instead that governance and clear ownership are the most central elements of how practitioners understand successful risk minimization.

H2 supported: Portfolio determinants are positively and significantly associated with perceived risk minimization outcomes ($\beta = 0.575$, $R^2 = 0.330$, $f^2 = 0.494$).

Discussion

The results offer three insights into how practitioners perceive portfolio construction and risk minimization. First, behavioral, macroeconomic, and diversification influences appear to operate together as a single set of underlying determinants (R1) that is strongly associated with both perceived construction efficiency (R2) and perceived risk minimization (R3). This offers a way of connecting the separation of behavioral and macroeconomic variables in traditional portfolio theory (Markowitz; Sharpe) with more recent behavioral theories of portfolio construction (Shefrin and Statman; Thaler and Sunstein).

Second, the large effect sizes of the structural paths ($f^2 = 0.755$ for R1 → R2; $f^2 = 0.494$ for R1 → R3) indicate that these associations are practically meaningful as well as statistically significant. PLS-SEM, as a distribution-free estimation technique, is well suited to this kind of exploratory analysis of non-normal survey data on investment practice (Hair et al.).

Third, holistic risk governance (Item 11, $\lambda = 0.788$) and risk ownership clarity (Item 12, $\lambda = 0.791$) are the two strongest-loading indicators in the entire model. This suggests that practitioners see governance architecture as central to what successful risk minimization means. This perspective is consistent with the principles of corporate governance according to the Basel Committee since it places ownership of risk with business lines, which is the first line of defense (Basel Committee on Banking Supervision, *Corporate Governance*). This is also in accordance with the Risk Management Framework for Mutual Funds published by the SEBI in 2021 which says that the managing companies of mutual funds should assign risk ownership to some specific CXO officials (Securities and Exchange Board of India). In addition, this finding suggests that in their investment management the institutional investment managers may gain by being transparent with respect to their risk ownership.

Contributions and Implications

PLS-SEM analysis of data collected from 101 practitioners working as financial professionals or investment bankers produced several statistically robust results. The Portfolio Determinants construct (R1), comprising behavioral, macroeconomic, and diversification dimensions, was strongly associated with Perceived Portfolio Construction Efficiency (R2), explaining 43% of its variance ($R^2 = 0.430$, H1), and with Perceived Risk Minimization Outcomes (R3), explaining 33% of its variance ($R^2 = 0.330$, H2). The reliability of Portfolio

Determinants (Cronbach's α = 0.736, AVE = 0.484) and Portfolio Construction Efficiency (Cronbach's α = 0.769, AVE = 0.490) meets exploratory psychometric criteria, and the risk minimization construct shows the highest reliability of the three (Cronbach's α = 0.806, AVE = 0.540). Discriminant validity for all three construct pairs was established, with HTMT ratios below the 0.85 threshold (R1 ↔ R2 = 0.739; R1 ↔ R3 = 0.669; R2 ↔ R3 = 0.639).

Previous studies have typically treated behavioral determinants, macroeconomic variables, and diversification approaches as distinct phenomena (Markowitz; Shefrin and Statman). The Integrative Determinant Architecture (IDA) proposed in this research brings these together in a single framework and provides initial empirical support for it from practitioner data. Rather than measuring realized trade-offs between global minimum variance, risk-based, and universal portfolio strategies, the IDA captures how practitioners assess their own construction efficiency against these paradigms (H1, $R^2 = 0.430$). This complements quantitative evidence that no single portfolio strategy consistently outperforms across the risk spectrum (DeMiguel et al.) by showing which determinants practitioners associate with building portfolios well.

Finally, this study introduces the Determinant-Driven Portfolio Risk Governance (DPRG) Framework as a conceptual model for operationalizing the IDA. It is built on three policy tiers: Determinant Identification, Construction Efficiency, and Risk Mitigation.

One major aspect that makes this framework useful is its application for three groups of audiences. The framework informs institutional fund managers that they should carry out behavioral bias audits prior to construction, as well as implement quantitative mapping of determinants in the pre-construction process. The framework has established grounds for regulatory bodies like SEBI, the FCA, and the SEC to propose mandatory disclosure of risk ownership. The framework allows people working in distribution-less environments to adopt PLS-SEM as a reliable empirical method to assess data gathered in exploratory studies on non-normal financial data.

Limitations and Future Research

Several limitations qualify these findings. First, all constructs are measured through respondents' self-reported perceptions. The study therefore shows how practitioners assess their own portfolio construction and risk outcomes, not how their portfolios actually performed. Validating the model against realized portfolio weights, returns, and drawdowns is the most important next step.

Second, overall model fit was weak. The SRMR was 0.160 for the saturated model and 0.164 for the estimated model, well above the 0.08 threshold recommended for PLS-SEM (Henseler et al., "Using PLS"), and the NFI of 0.425 is low. The measurement outcomes are credible and the estimates of the model are meaningful, however the model itself is not able to fully mirror the relationship that was observed, thus discoveries made from this research need to be regarded as exploratory rather than conclusive.

Moreover, the use of cross-sectional survey design makes claiming causality impossible; the mentioned paths represent correlation relationships and testing them timewise will be

necessary to determine the direction. Even though the chosen purposive sampling of 101 respondents is adequate to conduct PLS-SEM testing, it limits generalization regarding other groups and does not represent well the whole diversity of the global investor category. The AVE values of Portfolio Determinants (0.484) and Portfolio Construction Efficiency (0.490) fall below the commonly accepted threshold of 0.50 for sufficient convergent validity; while being acceptable for exploratory PLS studies, future applications of this model require further improvement of measurement tools. The framework has also not been validated in various regimes or asset types.

Future research may explore moderation effects, such as market volatility regimes as a moderator of the R1 → R3 path. More broadly, future research should build on the DPRG framework along five interrelated trajectories. First, empirical research should examine how behavioral and macroeconomic determinants are integrated through comparative determinant-based portfolio modeling, testing whether the identified constructs have differential predictive ability across bull, bear, and sideways market phases. State-dependent determinant modeling is seen as encouraging because it can lead to a dynamic optimization of governance frameworks rather than applying static assumptions.

Second, cross-industry validated results are salient. The current study involved the sample that was restricted only to finance specialists and bank investment experts. It would be useful to repeat the study using samples from REITs, sovereign funds or private equity funds, and retail investors, which would allow broadening the applications of the study. Such studies would also prove whether behavior determinants are equally predictive across all industries or only in specific ones.

Third, strategic portfolio frameworks based on structural equation modelling could use both PLS-SEM and covariance-based SEM to conduct more restrictive hypothesis tests under conditions of normality and to compare how measurement models perform across distributional environments. Structural modeling of diversification determinants in this context would clarify the boundary between diversification determinants and macroeconomic environments by no longer treating them as a composite factor.

Fourth, behavioral risk integration models that are based on Prospect Theory (Kahneman and Tversky) and taxonomies of cognitive bias may offer a method of putting into practice the behavioral determinant construct to a greater degree than what surveys on their own might achieve.

Lastly, developing agentic and algorithmic extensions of the DPRG framework with machine learning-assisted portfolio optimization systems represents the leading edge of applied research integrating governance theory with computational finance.

Conclusion

Using PLS-SEM, this research contributes to the portfolio management and financial governance literatures by proposing the Determinant-Driven Portfolio Risk Governance (DPRG) Framework and providing initial empirical support for it. Based on a purposive sample of 101

financial professionals and investment bankers, and supported by reliable measurement constructs (Cronbach's alpha coefficients between 0.736 and 0.806, and AVE values ranging from 0.484 to 0.540), the analysis indicates that behavioral, macroeconomic, and diversification determinants together explain 43% of the variance in perceived portfolio construction efficiency and 33% of the variance in perceived risk minimization outcomes.

The finding that holistic risk governance and ownership are the strongest indicators of perceived risk minimization is in line with the established rule of project management literature, according to which where one cannot indicate someone who is responsible for managing risks there exists not only an operational mistake but also a structural problem (*Guide to the Project Management Body*). The presented research accomplishes measuring this rule within a practitioner-oriented method of research.

Apart from its theoretical significance, DPRG Framework provides practitioners with the means of assessing factors they consider significant for constructing a successful portfolio and comparing global minimum variance with risk-based and universal portfolio strategies within a certain context. As this research is based on the self-reported data and the model fit overall is weak, these findings are only exploratory, and the next step is the testing of the presented framework against the real outcomes of the portfolio.

The importance of the DPRG Framework goes beyond the limits of theory. As the financial world becomes more complicated, distribution-free, behaviorally influenced, disrupted by the political forces, and more algorithmic by nature, regulators, institutional investors, and strategists are in need of a structured, tested, and adaptable framework to integrate determinants into portfolio construction.

Works Cited

Anderson, James C., and David W. Gerbing. "Structural Equation Modeling in Practice: A Review and Recommended Two-Step Approach." Psychological Bulletin, vol. 103, no. 3, 1988, pp. 411–23.

Barber, Brad M., and Terrance Odean. "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." The Quarterly Journal of Economics, vol. 116, no. 1, 2001, pp. 261–92.

Basel Committee on Banking Supervision. Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems. Bank for International Settlements, June 2011.

---. Corporate Governance Principles for Banks. Bank for International Settlements, July 2015.

Cesarone, Francesco, and Stefano Colucci. "Minimum Risk versus Capital and Risk Diversification Strategies for Portfolio Construction." Journal of the Operational Research Society, vol. 69, no. 2, 2018, pp. 183–200, https://doi.org/10.1057/s41274-017-0216-5.

Chan, Louis K. C., et al. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." The Review of Financial Studies, vol. 12, no. 5, 1999, pp. 937–74.

Cohen, Jacob. Statistical Power Analysis for the Behavioral Sciences. 2nd ed., Lawrence Erlbaum Associates, 1988.

Cover, Thomas M. "Universal Portfolios." Mathematical Finance, vol. 1, no. 1, 1991, pp. 1–29.

DeMiguel, Victor, et al. "Optimal versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?" The Review of Financial Studies, vol. 22, no. 5, 2009, pp. 1915–53.

Dixit, Vijaya, and Manoj Kumar Tiwari. "Project Portfolio Selection and Scheduling Optimization Based on Risk Measure: A Conditional Value at Risk Approach." Annals of Operations Research, vol. 285, nos. 1–2, 2020, pp. 9–33, https://doi.org/10.1007/s10479-019-03214-1.

Fastrich, Björn, et al. "Constructing Optimal Sparse Portfolios Using Regularization Methods." Computational Management Science, vol. 12, no. 3, 2015, pp. 417–34, https://doi.org/10.1007/s10287-014-0227-5.

Fornell, Claes, and David F. Larcker. "Evaluating Structural Equation Models with Unobservable Variables and Measurement Error." Journal of Marketing Research, vol. 18, no. 1, 1981, pp. 39–50.

Fung, William, and David A. Hsieh. "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds." The Review of Financial Studies, vol. 10, no. 2, 1997, pp. 275–302.

A Guide to the Project Management Body of Knowledge (PMBOK Guide). 6th ed., Project Management Institute, 2017.

Gunjan, A., and S. Bhattacharyya. "A Brief Review of Portfolio Optimization Techniques." Artificial Intelligence Review, vol. 56, no. 5, 2023, pp. 3847–86, https://doi.org/10.1007/s10462-022-10273-7.

Hair, Joseph F., et al. "When to Use and How to Report the Results of PLS-SEM." European Business Review, vol. 31, no. 1, 2019, pp. 2–24.

Henseler, Jörg, et al. "A New Criterion for Assessing Discriminant Validity in Variance-Based Structural Equation Modeling." Journal of the Academy of Marketing Science, vol. 43, no. 1, 2015, pp. 115–35.

Henseler, Jörg, et al. "Using PLS Path Modeling in New Technology Research: Updated Guidelines." Industrial Management & Data Systems, vol. 116, no. 1, 2016, pp. 2–20.

Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk." Econometrica, vol. 47, no. 2, 1979, pp. 263–91.

Kremer, Philipp J., et al. "Risk Minimization in Multi-Factor Portfolios: What Is the Best Strategy?" Annals of Operations Research, vol. 266, nos. 1–2, 2018, pp. 255–91, https://doi.org/10.1007/s10479-017-2467-6.

Lee, S. C., and William Eid Junior. "Portfolio Construction and Risk Management: Theory versus Practice." RAUSP Management Journal, vol. 53, no. 3, 2018, pp. 345–65, https://doi.org/10.1108/rausp-04-2018-009.

Lim, Churlzu, et al. "Portfolio Optimization by Minimizing Conditional Value-at-Risk via Nondifferentiable Optimization." Computational Optimization and Applications, vol. 46, no. 3, 2010, pp. 391–415, https://doi.org/10.1007/s10589-008-9196-3.

Markowitz, Harry. "Portfolio Selection." The Journal of Finance, vol. 7, no. 1, 1952, pp. 77–91.

Mohr, Esther, and Robert Dochow. "Risk Management Strategies for Finding Universal Portfolios." Annals of Operations Research, vol. 256, no. 1, 2017, pp. 129–47, https://doi.org/10.1007/s10479-016-2176-6.

Naguez, N. "Dynamic Portfolio Insurance Strategies: Risk Management under Johnson Distributions." Annals of Operations Research, vol. 262, no. 2, 2018, pp. 605–29, https://doi.org/10.1007/s10479-016-2121-8.

Nunnally, Jum C. Psychometric Theory. 2nd ed., McGraw-Hill, 1978.

Parmee, Nick. "Hedge Fund Liquidations Slow." WealthBriefing, 18 Sept. 2009, www.wealthbriefing.com/html/printarticle.php?id=25347.

Ringle, Christian M., et al. SmartPLS 4. SmartPLS GmbH, 2024, www.smartpls.com.

Rom, Brian M., and Kathleen W. Ferguson. "Post-Modern Portfolio Theory Comes of Age." The Journal of Investing, vol. 3, no. 3, 1994, pp. 11–17.

Roman, Diana, and Gautam Mitra. "Portfolio Selection Models: A Review and New Directions." Wilmott Journal, vol. 1, no. 2, 2009, pp. 69–85.

Roumpis, E., and T. Syriopoulos. "Dynamics and Risk Factors in Hedge Funds Returns: Implications for Portfolio Construction and Performance Evaluation." The Journal of Economic Asymmetries, vol. 11, 2014, pp. 58–77, https://doi.org/10.1016/j.jeca.2014.07.001.

Securities and Exchange Board of India. "Risk Management Framework (RMF) for Mutual Funds." Circular no. SEBI/HO/IMD/IMD-1 DOF2/P/CIR/2021/630, 27 Sept. 2021.

Settembre-Blundo, Davide, et al. "Flexibility and Resilience in Corporate Decision Making: A New Sustainability-Based Risk Management System in Uncertain Times." Global Journal of Flexible Systems Management, vol. 22, suppl. 2, 2021, pp. 107–32, https://doi.org/10.1007/s40171-021-00277-7.

Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." The Journal of Finance, vol. 19, no. 3, 1964, pp. 425–42.

Shefrin, Hersh, and Meir Statman. "Behavioral Portfolio Theory." Journal of Financial and Quantitative Analysis, vol. 35, no. 2, 2000, pp. 127–51.

Shimizu, H., and T. Shiohama. "Multifactor Portfolio Construction by Factor Risk Parity Strategies: An Empirical Comparison of Global Stock Markets." Asia-Pacific Financial Markets, vol. 26, no. 4, 2019, pp. 453–77, https://doi.org/10.1007/s10690-019-09274-4.

Thaler, Richard H., and Cass R. Sunstein. Nudge: Improving Decisions about Health, Wealth, and Happiness. Yale UP, 2008.

"Topic: Investment Funds." Statista, 17 Dec. 2025, www.statista.com/topics/5670/investment-funds/.

Tversky, Amos, and Daniel Kahneman. "Judgment under Uncertainty: Heuristics and Biases." Science, vol. 185, no. 4157, 1974, pp. 1124–31.

An Evolutionary Perspective on Human Susceptibility to Alzheimer's Disease
By Hannah Kim

Abstract

Alzheimer's disease (AD) is one of the most prevalent neurodegenerative diseases and the main cause of dementia worldwide. The two pathological hallmarks of AD are the aggregation of amyloid-beta (Aβ) peptides into senile plaques and the formation of tau protein neurofibrillary tangles (NFTs). Although those hallmarks have been observed in other primates, only humans have the full pathological features of AD. One fundamental evolutionary question is why the human brain is especially susceptible to AD. In this paper, we provide an evolutionary overview of AD pathology and discuss the possibility that human brain susceptibility to AD stems from trade-offs associated with evolving a large, complex brain. We first introduce the molecular mechanism of Aβ accumulation and Tau phosphorylation and discuss how normal processes may become harmful during aging. Following this, we discuss several evolutionary hypotheses, including reduced selection in late life, antagonistic pleiotropy, and human brain expansion. Finally, we discuss how comparative genomics, single-cell analysis, and human brain organoids may help identify human-specific factors that might contribute to AD pathology. This paper provides a complementary framework for understanding current evolutionary hypotheses and future directions for studying biological features important for AD susceptibility.

1. Introduction

Alzheimer's disease (AD) is a progressive neurodegenerative disorder and the most common cause of dementia(Arachchige et al.). Because of its high prevalence, it is a major public health crisis, especially as life expectancy has steadily increased. The number of patients with AD has dramatically increased, resulting in substantial healthcare burdens *("2023 Alzheimer's Disease Facts and Figures").* Although there has been numerous ongoing research on AD, there is still no cure for this disease. Current therapeutic approaches mainly focus on symptomatic relief; the molecular mechanisms of AD are still poorly understood (Arachchige et al.).

AD is particularly interesting from an evolutionary perspective because its full pathological phenotype has not been observed in any nonhuman primates (NHPs). However, other primates still have AD-related changes. For example, aged chimpanzees, gorillas, and other primates develop amyloid-beta (Aβ) deposits and certain forms of tau pathology (Edler et al.; Perez et al.). Therefore, humans and other NHPs share several features of brain aging, including neuroinflammation, synaptic loss, and protein accumulation. The important question is not simply why other primates do not develop AD, but rather why humans appear to be particularly vulnerable to severe AD-related neurodegeneration.

The evolution of the human brain may be one factor that contributes to our susceptibility to AD (Fox). During human evolution, the neocortex became larger and more complex, supporting cognitive functions such as learning, memory, and problem-solving. These changes

gave humans important cognitive advantages, but they also increased the demands placed on neurons. Maintaining a large and complex brain over a long lifespan requires considerable energy and may make neurons more vulnerable to age-related cellular stress. Several evolutionary hypotheses have been proposed to explain how these characteristics of the human brain might contribute to AD.

In this review, we examine the hypothesis that human vulnerability to AD may stem from evolutionary trade-offs related to human brain size, complexity, and longevity. We first review the molecular mechanisms of Aβ and tau pathology, and discuss several evolutionary hypotheses involving longevity, brain expansion, metabolism, and environmental mismatch. Finally, we discuss several biological approaches that may help identify the biological features associated with human susceptibility to AD.

2. Molecular Mechanisms of AD in an Evolutionary Context

2.1 The Amyloid Cascade Pathway in an evolutionary context

Since Glenner and Wong discovered that amyloid-beta protein (Aβ) is the central component of extracellular amyloid plaques in AD, the Aβ hypothesis suggests that an imbalance between Aβ42 and Aβ peptide production and clearance may initiate AD (Glenner and Wong). Aβ is generated from the amyloid precursor protein (APP) by sequential cleavage by β- and γ-secretases (Dawkins and Small). APP is cleaved by β-secretase (BACE1) and the γ-secretase complex, producing Aβ40 and Aβ42 peptides (Sheng et al.). Aβ42 then readily self-assembles into soluble oligomers, which are the most neurotoxic molecules at synapses (Qiu et al.). APP and its cleavage products are involved in many brain functions, including synaptic plasticity, neurite outgrowth, and cell adhesion (Sheng et al.).

APP processing and Aβ production have been detected in other nonhuman primates (Martin et al.). The conservation of Aβ-related processes across mammals suggests that Aβ alone is unlikely to explain why humans are particularly susceptible to AD. Evolutionary differences in Aβ production, aggregation, clearance, and cellular response to Aβ may contribute to species-specific vulnerability (Rosen et al.). From an evolutionary perspective, as humans age and live longer, Aβ clearance may become less efficient (August et al.). Because of longer lifespans, humans are more vulnerable to cellular damage, and impaired protein clearance can accumulate over several decades. This imbalance between Aβ production and clearance could lead to an accumulation of soluble Aβ oligomers.

Aβ oligomer accumulation may contribute to early synaptic dysfunction before large extracellular amyloid plaques form. From an evolutionary perspective, researchers may not suggest that humans produce more Aβ. Instead, the current theory suggests that biological systems in humans may become maladaptive because they are associated with longer lifespan and age-related decline. Although this evolutionary explanation could be a hypothesis, it provides a possible framework for understanding why normal Aβ biology may become harmful during aging.

2.2 Tau Hyperphosphorylation in an evolutionary context

Tau is a microtubule-associated protein that helps stabilize microtubules and supports the transport of molecules within neurons (Nelson et al.). In AD, tau can become abnormally hyperphosphorylated, causing it to detach from microtubules and form insoluble filaments that eventually accumulate as neurofibrillary tangles (NFTs)(Nelson et al.). As tau pathology progresses, it is associated with synaptic dysfunction, neuronal loss, and cognitive decline.

From an evolutionary perspective, changes in the human brain may have contributed to our vulnerability to tau pathology. The human neocortex expanded considerably during evolution, especially cortical regions involved in higher cognitive functions. Human cortical development also shows an extended period of synaptic maturation compared with that of other primates (Liu et al.). These changes supported more complex cognitive functions, but maintaining a large and highly connected brain over a long lifespan may also place greater demands on neurons. Such demands could make some neuronal populations more vulnerable to age-related stress and tau pathology.

Several brain regions that underwent significant evolutionary changes are also particularly vulnerable to AD. For example, the hippocampus and its related regions play important roles in memory and show substantial changes during human evolution (August et al.). The CA1 region and subiculum are especially vulnerable to neuronal loss and tau pathology in AD (Bouwman et al.). In addition, certain long-range projection neurons in the human neocortex appear to be particularly vulnerable to degeneration.

Tau is highly conserved among primates, and its amino acid sequence is very similar between humans and other primates (Holzer et al.). This suggests that differences in the tau protein itself may not explain why humans are particularly vulnerable to tau-related neurodegeneration. Other factors, such as differences in brain organization, neuronal connectivity, and the long human lifespan, may also be important. These features support human cognitive function, but they may also make the brain more vulnerable to tau pathology as it ages.

3. Evolutionary Hypotheses for Human AD Susceptibility

While AD is one of the most prevalent neurodegenerative diseases in the human species, other non-human primates usually do not develop the form of AD(Rapoport and Nelson). To explain this species-specific prevalence of AD, researchers consider several features of human evolution. These include our longer lifespan, the larger and more complex human brain structure, the brain's high energy use, genetic trade-offs, and changes in the modern environment. In this chapter, we will cover several evolutionary hypotheses.

3.1 Longevity and antagonistic pleiotropy

The idea of a mismatch between longevity and evolutionary selection would explain the human susceptibility to AD. Because survival and reproduction strongly affect the next generation, natural selection is stronger during early life stages (Williams). In older life stages, natural selection tends to have less effect (Williams). Although a specific biological trait would

be harmful in later stages, it can be more beneficial during early life. Therefore, this concept suggests that AD might derive from biological traits that benefit brain function in early and middle age but are less effective in late life.

This doesn't mean humans evolved to live only 40-50 years. Average life expectancy in the past was much lower partly because many people died during childhood or early adulthood. Individuals who survived these early risks could still live to older ages. Therefore, a more important evolutionary explanation is that natural selection generally weakens with age (Williams and Day). At the same time, the cellular costs of maintaining a highly complex and metabolically active brain may accumulate over time (Fox).

This idea is also related to antagonistic pleiotropy, a term describing a situation in which a genetic trait is beneficial earlier in life but harmful in old age (Williams). APOE E4 is one of the strongest genetic risk factors for sporadic AD (Belloy et al.). At the same time, APOE4 plays a critical role in lipid metabolism, immune responses, infection, and tissue repair in some ancestral environments (Belloy et al.).Therefore, the continued presence of some AD-related genetic variants in human populations may not mean that evolution "failed" to remove them. Instead, these variants may have provided benefits in some situations while creating risks later in life.

3.2 Brain expansion, development, and metabolism

Unusual evolution of the human brain may also explain human susceptibility to AD. Humans have a relatively larger brain size than other apes, especially in the neocortex, which is involved in complex thinking, learning, and social behavior(Rapoport and Nelson). Brain size alone doesn't explain the pathology of AD, since some animals have larger brains than humans but do not develop the same form of AD. Instead, the organization and function of specific brain regions may be critical to AD development.

Several brain areas affected during AD onset, including the frontal, temporal, and parietal cortex, have also changed substantially during evolution (Rapoport). Among them, the researchers focused on the hippocampus and entorhinal cortex because they play a critical role in learning and memory and are among the earliest regions affected by AD. Compared with the brain development of other animals, the human brain develops more slowly and maintains a high level of plasticity into adulthood, requiring more energy. Although the brain is a small portion of total body weight, it uses about 20% of the body's energy at rest.

Maintaining this level of activity for many decades may increase cellular stress, including oxidative stress, mitochondrial damage, and calcium regulation problems. The fact that humans developed a large, highly active brain and lived longer than other primates could support this idea. To maintain a complex brain structure for many decades, humans may accumulate cellular damage more readily over time.

3.3 Environmental mismatch

Another possible explanation focuses on the difference between the modern environment and the environment humans lived in when they evolved (Fox). The environmental mismatch hypothesis proposes that throughout human evolutionary history, humans were more physically active, couldn't predict food availability, and relied on natural light and dark cycles for daily activity (Fox). In modern environments, humans don't need to be as physically active and have constant access to high-calorie diets, so their life cycles no longer need to rely on natural light/dark cycles. Those changes might affect biological vulnerabilities that developed during human evolution.

Sleep would be especially critical. The glymphatic system, the brain's waste-clearance system, becomes active during sleep (Zare et al.). The glymphatic system helps remove toxic substances, including Aβ and Tau, from the brain, and long-term sleep disruption could reduce their removal (Zare et al.).

In addition to sleep, modern metabolic problems also contribute to AD risk (Terzo et al.). In modern lifestyles, eating patterns result in a greater tendency to have several eating-related problems, including obesity, insulin resistance, chronic inflammation, and cardiovascular disease, and these increase stress on neurons. Exposure to air pollution, heavy metals, pesticides, and other pollutants has been associated with biological processes such as oxidative stress and neuroinflammation (Yegambaram et al.). Therefore, environmental pollution could be another factor affecting AD pathology. These environmental factors may add additional stress to a brain that already has high metabolic demands, resulting in AD pathology (Ashleigh et al.).

4. Future Directions

To examine the evolutionary changes that make humans specifically susceptible to AD, we could use several biological methodologies. In this chapter, we will review future directions and methodological strategies to resolve the precise molecular mechanisms.

4.1 Multi-omics at single-cell resolution

Single-cell RNA sequencing (scRNA-seq) would be a useful tool for resolving the molecular mechanisms that contribute to selective cellular susceptibility to AD. This multi-omic approach improves understanding of the distinct transcriptomic and epigenetic landscapes of cell types susceptible to AD pathology (Morabito et al.). Researchers can analyze human post-mortem tissues at the single-cell level and isolate gene expression networks in specific cell types shown to be vulnerable to AD pathology. Integrating scRNA-seq with other single-cell approaches, such as single-cell ATAC-seq and spatial transcriptomics, could further reveal how gene regulatory elements and cellular interactions contribute to selective vulnerability(Avey et al.; Morabito et al.). This approach may also help identify molecular signatures that distinguish vulnerable human neurons from relatively resistant cell populations, providing potential targets for understanding and preventing AD pathology.

4.2 Human brain organoids

Rodent models carrying AD-associated mutations have been widely used to study AD pathology. However, there are important differences between rodent and human brains that can make it difficult to model some aspects of the disease. Human cerebral organoids generated from induced pluripotent stem cells (iPSCs) provide another way to study AD in a human cellular system (Sreenivasamurthy et al.). Unlike conventional 2D cell cultures, these 3D models can reproduce some features of human brain development and cellular organization, making them useful for studying disease mechanisms that may be difficult to examine in animal models. Using genome-editing tools such as the CRISPR/Cas9 system, researchers could further introduce human-specific evolutionary changes into brain organoids and test whether mutations relevant to evolutionary adaptation affect susceptibility to Ab-mediated pathology.

4.3 Comparative primate studies

Comparative studies of gene expression across humans and other primates, including chimpanzees, bonobos, and macaques, may help identify biological features that are unique to the human brain (Khrameeva et al.). Researchers can compare gene expression in brain regions such as the prefrontal cortex and hippocampus to identify differences in gene regulation and other molecular pathways across species. These studies may help explain why humans are more susceptible to AD, while other primates appear to be relatively resistant to severe AD-related neurodegeneration. Understanding these differences could also identify biological pathways that may be important for future AD research.

5. Conclusion

The current study suggests that AD may not be a simple disease that appears accidentally old, but could be a byproduct of human evolutionary history. Throughout millions of years, the human brain, especially the neocortex, rapidly expanded and became more complex, allowing human beings to have better cognitive capabilities, including learning and memory, language, problem-solving,g and complex social interaction. However, that rapid expansion also had a disadvantage. Human neurons have complex structures that require large amounts of energy. These features would make humans more susceptible to aging and neurodegenerative disorders. For example, because the brain is highly active and specialized, humans experience more metabolic stress over time. In addition, humans live much longer than non-human primates, and humans live in environments very different from those of earlier generations. Environmental factors including excess calories and disrupted sleep may further increase vulnerability to AD. From an evolutionary perspective, AD likely did not arise from a single evolutionary cause. Instead, an evolutionary perspective may help us better understand why certain human brain cells and networks are more vulnerable to neurodegeneration. To better understand species-specific susceptibility to AD, research will use cutting-edge technology, including single-cell RNA transcript analysis, brain organoids, and comparisons between humans and other primates. This may help identify the biological differences that contribute to this vulnerability. Understanding

AD as a possible consequence of evolutionary trade-offs may therefore provide new insights into both the origins of the disease and potential ways to protect the aging human brain.

Works Cited

"2023 Alzheimer's Disease Facts and Figures." Alzheimer's & Dementia, vol. 19, no. 4, Apr. 2023, pp. 1598–695. DOI.org (Crossref), https://doi.org/10.1002/alz.13016.

Arachchige, Arosh S. Perera Molligoda, et al. "Alzheimer's Disease: A Comprehensive Review of Epidemiology, Pathophysiology, Diagnosis, and Treatment." AIMS Neuroscience, vol. 13, no. 2, 2026, pp. 208–43. PubMed, https://doi.org/10.3934/Neuroscience.2026009.

Ashleigh, Theophania, et al. "The Role of Mitochondrial Dysfunction in Alzheimer's Disease Pathogenesis." Alzheimer's & Dementia, vol. 19, no. 1, Jan. 2023, pp. 333–42. DOI.org (Crossref), https://doi.org/10.1002/alz.12683.

August, Isabel, et al. "Evolution of Human Susceptibility to Alzheimer's Disease: A Review of Hypotheses and Comparative Evidence." Evolutionary Anthropology, vol. 34, no. 1, Mar. 2025, p. e22054. PubMed, https://doi.org/10.1002/evan.22054.

Avey, Denis R., et al. "Uncovering Plaque-Glia Niches in Human Alzheimer's Disease Brains Using Spatial Transcriptomics." Molecular Neurodegeneration Advances, vol. 1, no. 1, 2025, p. 2. PubMed, https://doi.org/10.1186/s44477-025-00002-z.

Belloy, Michaël E., et al. "A Quarter Century of APOE and Alzheimer's Disease: Progress to Date and the Path Forward." Neuron, vol. 101, no. 5, Mar. 2019, pp. 820–38. PubMed, https://doi.org/10.1016/j.neuron.2019.01.056.

Bouwman, Maud M. A., et al. "Hippocampal Subfields: Volume, Neuropathological Vulnerability and Cognitive Decline in Alzheimer's and Parkinson's Disease." Alzheimer's Research & Therapy, vol. 17, no. 1, May 2025, p. 121. PubMed, https://doi.org/10.1186/s13195-025-01768-w.

Dawkins, Edgar, and David H. Small. "Insights into the Physiological Function of the β-Amyloid Precursor Protein: Beyond Alzheimer's Disease." Journal of Neurochemistry, vol. 129, no. 5, Jun. 2014, pp. 756–69. PubMed, https://doi.org/10.1111/jnc.12675.

Edler, Melissa K., et al. "Aged Chimpanzees Exhibit Pathologic Hallmarks of Alzheimer's Disease." Neurobiology of Aging, vol. 59, Nov. 2017, pp. 107–20. PubMed, https://doi.org/10.1016/j.neurobiolaging.2017.07.006.

Fox, Molly. "'Evolutionary Medicine' Perspectives on Alzheimer's Disease: Review and New Directions." Ageing Research Reviews, vol. 47, Nov. 2018, pp. 140–48. DOI.org (Crossref), https://doi.org/10.1016/j.arr.2018.07.008.

Glenner, G. G., and C. W. Wong. "Alzheimer's Disease: Initial Report of the Purification and Characterization of a Novel Cerebrovascular Amyloid Protein." Biochemical and Biophysical Research Communications, vol. 120, no. 3, May 1984, pp. 885–90. PubMed, https://doi.org/10.1016/s0006-291x(84)80190-4.

Holzer, Max, et al. "Tau Gene (MAPT) Sequence Variation among Primates." Gene, vol. 341, Oct. 2004, pp. 313–22. PubMed, https://doi.org/10.1016/j.gene.2004.07.013.

Khrameeva, Ekaterina, et al. "Single-Cell-Resolution Transcriptome Map of Human, Chimpanzee, Bonobo, and Macaque Brains." Genome Research, vol. 30, no. 5, May 2020, pp. 776–89. PubMed, https://doi.org/10.1101/gr.256958.119.

Liu, Xiling, et al. “Extension of Cortical Synaptic Development Distinguishes Humans from Chimpanzees and Macaques.” Genome Research, vol. 22, no. 4, Apr. 2012, pp. 611–22. PubMed, https://doi.org/10.1101/gr.127324.111.

Martin, L. J., et al. “Amyloid Precursor Protein in Aged Nonhuman Primates.” Proceedings of the National Academy of Sciences of the United States of America, vol. 88, no. 4, Feb. 1991, pp. 1461–65. PubMed, https://doi.org/10.1073/pnas.88.4.1461.

Morabito, Samuel, et al. “Single-Nucleus Chromatin Accessibility and Transcriptomic Characterization of Alzheimer’s Disease.” Nature Genetics, vol. 53, no. 8, Aug. 2021, pp. 1143–55. PubMed, https://doi.org/10.1038/s41588-021-00894-z.

Nelson, P. T., et al. “Molecular Evolution of Tau Protein: Implications for Alzheimer’s Disease.” Journal of Neurochemistry, vol. 67, no. 4, Oct. 1996, pp. 1622–32. PubMed, https://doi.org/10.1046/j.1471-4159.1996.67041622.x.

Perez, Sylvia E., et al. “Early Alzheimer’s Disease-Type Pathology in the Frontal Cortex of Wild Mountain Gorillas (Gorilla Beringei Beringei).” Neurobiology of Aging, vol. 39, Mar. 2016, pp. 195–201. PubMed, https://doi.org/10.1016/j.neurobiolaging.2015.12.017.

Qiu, Tian, et al. “Aβ42 and Aβ40: Similarities and Differences.” Journal of Peptide Science: An Official Publication of the European Peptide Society, vol. 21, no. 7, Jul. 2015, pp. 522–29. PubMed, https://doi.org/10.1002/psc.2789.

Rapoport, S. I. “Hypothesis: Alzheimer’s Disease Is a Phylogenetic Disease.” Medical Hypotheses, vol. 29, no. 3, Jul. 1989, pp. 147–50. DOI.org (Crossref), https://doi.org/10.1016/0306-9877(89)90185-0.

Rapoport, Stanley I., and Peter T. Nelson. “Biomarkers and Evolution in Alzheimer Disease.” Progress in Neurobiology, vol. 95, no. 4, Dec. 2011, pp. 510–13. PubMed, https://doi.org/10.1016/j.pneurobio.2011.07.006.

Rosen, Rebecca F., et al. “Comparative Pathobiology of β-Amyloid and the Unique Susceptibility of Humans to Alzheimer’s Disease.” Neurobiology of Aging, vol. 44, Aug. 2016, pp. 185–96. PubMed, https://doi.org/10.1016/j.neurobiolaging.2016.04.019.

Sheng, Morgan, et al. “Synapses and Alzheimer’s Disease.” Cold Spring Harbor Perspectives in Biology, vol. 4, no. 5, May 2012, p. a005777. PubMed, https://doi.org/10.1101/cshperspect.a005777.

Sreenivasamurthy, Sai, et al. “Current Progress of Cerebral Organoids for Modeling Alzheimer’s Disease Origins and Mechanisms.” Bioengineering & Translational Medicine, vol. 8, no. 2, Mar. 2023, p. e10378. PubMed, https://doi.org/10.1002/btm2.10378.

Terzo, Simona, et al. “From Obesity to Alzheimer’s Disease through Insulin Resistance.” Journal of Diabetes and Its Complications, vol. 35, no. 11, Nov. 2021, p. 108026. PubMed, https://doi.org/10.1016/j.jdiacomp.2021.108026.

Williams, George C. “PLEIOTROPY, NATURAL SELECTION, AND THE EVOLUTION OF SENESCENCE.” Evolution, vol. 11, no. 4, Dec. 1957, pp. 398–411. DOI.org (Crossref), https://doi.org/10.1111/j.1558-5646.1957.tb02911.x.

Williams, Paul D., and Troy Day. “Antagonistic Pleiotropy, Mortality Source Interactions, and the Evolutionary Theory of Senescence.” Evolution; International Journal of Organic Evolution, vol. 57, no. 7, Jul. 2003, pp. 1478–88. PubMed, https://doi.org/10.1111/j.0014-3820.2003.tb00356.x.

Yegambaram, Manivannan, et al. “Role of Environmental Contaminants in the Etiology of Alzheimer’s Disease: A Review.” Current Alzheimer Research, vol. 12, no. 2, 2015, pp. 116–46. PubMed, https://doi.org/10.2174/1567205012666150204121719.

Zare, Farshad, et al. “Sleep-Dependent Clearance of Brain Metabolites via the Glymphatic System: Implications for Alzheimer’s Pathophysiology.” Brain and Behavior, vol. 16, no. 4, Apr. 2026, p. e71374. PubMed, https://doi.org/10.1002/brb3.71374.

Urban Design and Human Behavior: To what extent does urban design affect human behavior? By Ana-Maria Vulturu

International Computer High School of Bucharest

Abstract

Urban design plays a significant role in shaping the physical environment in which people live, while simultaneously influencing the ways individuals move, interact, and experience their surroundings. This research investigates the relationship between urban design and human behaviour, with particular focus on the extent to which the organization and quality of urban environments can influence quality of life, mental well-being, social interaction, physical activity, and the sense of community and belonging. The research explores urban design as a multidisciplinary process involving the planning of buildings, public spaces, mobility and transport systems, green areas, urban landscapes, and the regeneration of existing environments. It examines how these elements can contribute to the creation of sustainable, accessible, inclusive, and people-oriented cities.

The study first examines the historical development of urban design, beginning with the structured settlements of ancient civilizations and continuing through the Medieval, Renaissance, Baroque, and Industrial periods. This historical perspective demonstrates how changes in population, transportation, technology, political authority, and social conditions continuously influenced the physical structure of cities. The research then focuses on the contemporary concept of walkable cities, examining the relationship between pedestrian accessibility, transportation, green spaces, environmental sustainability, and human well-being. Walkable environments can provide greater access to essential facilities while encouraging physical activity and creating opportunities for social interaction. In contrast, car-oriented environments can contribute to traffic congestion, pollution, reduced pedestrian accessibility, and fewer opportunities for spontaneous social connections.

Furthermore, the research investigates the importance of public spaces in encouraging social interaction and strengthening a sense of community and belonging. Streets, squares, parks, plazas, playgrounds, and other shared environments can provide opportunities for individuals from different social groups to gather and interact. The physical appearance and accessibility of these spaces can influence whether people choose to engage with, move through, or remain in an environment. Green spaces are also examined as an important component of urban design because of their connection to mental well-being, environmental quality, and the reduction of the urban heat island effect.

The research additionally examines the decision-making processes behind urban development. It considers the roles of local governments, urban planners, technical experts, and residents, emphasizing the importance of cooperation, financial resources, public participation, planning

regulations, and land-use policies in transforming urban design proposals into implemented projects. Helsinki is considered as an example of how governmental regulations and sustainability requirements can influence buildings, transportation, green areas, and the overall structure of an urban environment.

Methodologically, the research is based on a literature review conducted using sources identified through searches of Google Scholar and Wikipedia. Approximately 120 articles were initially identified, with 52 selected for in-depth review and 15 ultimately meeting the established date, methodology, and accuracy criteria. The reviewed studies, published between 2011 and 2024, were analyzed according to variables including study type, publication year, sample size, statistical evidence, and quantitative and qualitative data. A comparative perspective involving Copenhagen and Bucharest further illustrates how differences in urban planning, infrastructure, mobility, green spaces, governance, and public spaces can influence urban experiences. Overall, the research demonstrates that urban design and human behaviour are closely connected, and that well-planned, sustainable, accessible, and people-oriented environments can contribute to improved quality of life, stronger social connections, and greater well-being.

Introduction

Urban Design is an approach concentrated on the planning and arrangement of buildings, public spaces and the interstitial areas between them. It is structured into several main branches, such as spatial design (urban structure), public space design, mobility and transport, urban landscape design, urban regeneration, socio-cultural dimension, and the strategic and economic component, each one contributing in shaping a sustainable urban environment. The complex design processes involved in restoring, replenishing, and modernizing specific areas within an urban design landscape can gradually unfold over a long-term duration, generally 10 to 20 years (OECD (2013). Compact City Policies: A Comparative Assessment). It is also defined as a “collaborative and multi-disciplinary process of shaping the physical setting for life – the art of making places” (1). Urban design covers a great range of locations such as towns, green areas, neighborhoods, city centers, suburbs, train or subway stations.

Furthermore, urban design has been proven to positively correlate with improved mental health and higher satisfaction in quality of life. According to scientific consensus, quality of life has “no standard definition or measure” “nor is there a standard model for investigating subjective QOL” (3). However it is often understood as a multidimensional construct encompassing physical well-being, psychological health, social relationships, and overall life satisfaction, and is associated with “various terms such as happiness, life satisfaction and well-being have been used in the context of investigating QOL” (3), where QOL refers to quality of life.

The environment can influence individuals, impacting their emotional and social state. As a result, studies suggest that well-designed cities can provide social cohesion,” A sense of community and belonging” (2) as well as offering urban areas for citizens to integrate. However,

when these environments are badly structured and poorly designed, they can play a very important role when dealing with anxiety, unease and lack of social interaction. By exploring and understanding the relationship between urban design and human behaviors, we can establish and alter cities which simultaneously increase quality of life and sustain mental health.

The aim of this research paper is to investigate and analyze the connections formed between urban design and human behavior. By exploring this topic the purpose is to demonstrate that there is a direct connection between urban design and human behaviors, and explain to what extent urban design influences human behavior, while explaining the correlation between both. Furthermore, the goal is to prove that a well-designed, aesthetically pleasing, sustainable, sociable environment can release stress, anxiety and negative emotions, to a certain extent.

Methodology

This chapter aims to summarize a comprehensive body of resources used to support the methodological framework necessary to identify, select, and analyze the urban design literature that depicts the strengths and limitations of how this discipline can positively influence human behaviour. The purpose of this review is to analyse the direct connection between urban design and human behaviour as well as investigating how well-designed surroundings can positively influence a mode of action or how poorly designed places can negatively influence a human's actions. The findings support the primary goal of the thesis: demonstrating that an aesthetically pleasing, green, and well-planned area can make a positive impact on human behaviour as well as increasing quality of life.

To identify the most relevant and up-to-date evidence, searches were conducted using databases such as Google Scholar, Wikipedia, ScienceDirect. The research period began January 2026 and extended to September 2026.

The keywords used during the process, as well as their various combinations, included “urban design”, “human behaviour”, “quality of life”, “green spaces”, "interstitial spaces”, “well-being”, “planning”, “governments”, “regulations”; Boolean operators (AND/OR) were used to narrow or expand search results as the need for more thorough search parameters arose.
In the context of this thesis, the studies spanned between 2016 and 2026 to ensure a comprehensive analysis of the topic and provide credible urban design and human behaviour proof about long-term outcomes. Moreover, only articles published in English were considered as it offered a broader library of resources to formulate compelling findings.
The initial search yielded a library of approximately 120. After a thorough review of their titles and abstracts, only 52 articles were selected for in-depth reviews, of which 15 met the date, methodology, and accuracy criteria established. The following variables were decisive in the process of data extraction from the final pool of studies: type of study, publication year, sample size, statistical evidence, quantitative and qualitative data.

The urban design literature reviewed revealed trends consistent with the findings of this thesis. The methodologies applied proved that long term outcomes of human behaviour can be ascertained through the strategic planning of urban environments.

The history of urban design

The history of this discipline has been around from ancient civilizations (early origins). The concept effectively began around 3000-2000 BCE, because before these years most human settlements were mainly villages that had no clear structure. By 3000 BCE major changes took place, population increased and society became more organized and hierarchical. Eventually, this made people reflect more about how cities must be planned. Initially, small adjustments were made, from 2500 BCE to 500 CE cities started using grid systems, particularly in Greek and Roman cultures, mainly for order and defence.

Thereafter, in the Medieval Era, unlike the ancient grid cities, the medieval towns were unplanned cities. Due to the central authority that was missing, cities expanded as people settled where they had space, therefore the planning of the cities had irregular paths and almost gave the impression of a maze.

In the late 15th century, when the medieval period ended, urban design transitioned into the Renaissance and Baroque eras, which lasted between the 1400s and 1750s. As opposed to the principles that governed urban design in the Medieval Era, Renaissance cities rediscovered the interest in geometry, symmetry and order inspired by the Greek and Roman ideas “Renaissance architecture adopted distinguishing features of classical Roman architecture. However, the forms and purposes of buildings had changed over time, as had the structure of cities, which is reflected in the fusion of classical and 16th century forms” (5)

The cities were designed by using centralized plans, sometimes forming “a bastion fort or trace italienne (a phrase derived from non-standard French, meaning 'Italian outline') [ ] a fortification in a style developed during the early modern period” (4). This pattern was used as a response for gunpowder weapons such as cannons.These designs emphasized the harmony and aesthetics making a transition into well-thought urban planning.

During the Baroque years, all urban intervention became more monumental and ambitious. As a result, street design widened, becoming boulevards, the streets were no longer crooked nor had an irregular layout; moreover, public spaces were often designed to symbolize control, dominance, power and visually express authority “One of the most significant programmes of urban planning in the Baroque Period was Rome. Inspired by the ideal of the Renaissance city, Pope Sixtus V's ambitious urban reform programme transformed the old environment to emulate the "long straight streets, wide regular spaces, uniformity and repetitiveness of structures, lavish use of commemorative and ornamental elements, and maximum visibility from both linear and circular perspective." “.

During the Industrial Revolution, between the 18th and 19th century, another major transformation occurred. Because rapid industrialization influenced the population, leading into a massive growth, the main problem became overcrowding and poor sanitation; cities became divided by stark contrasts between poverty and wealth “Giant sprawling cities developed during this era, exhibiting the luxuries of wealth and the meanness of poverty in sharp juxtaposition” (5). Due to this, environmental and public health problems appeared. However, implementing engineering upgrades to water supplies and sanitation systems was crucial for the expanding of cities. Moreover, during these centuries, factories played a huge role in triggering the mass migration from rural areas to the city, because of the huge numbers of jobs it allowed. As people moved into the city they were faced with overcrowding, poverty, average to low living conditions and health problems (as densely populated cities witnessed a more prolific rate of disease spread).

“At the same time, middle and upper class people chose to move toward the suburbs, seeking to escape the overcrowding, noise, pollution and desire to improve their living standards by having a quieter and cleaner life.” (4) All these changes took place due to the new transportation systems such as railways and trams, being the most important ones, followed up by streetcars (that were initially horse-drawn, and later on electric) and canals.

Furthermore, because of the overcrowding of cities, individual houses were no longer the go-to choice, so landowners and developers thought of a solution that aimed to maximize profit. The buildings were not designed based on preference, they were thought to fit as many people as possible within a limited space. This, later on, led to early tenement buildings, apartment blocks, and shared housing units. In the industrial era, during the 19th and 20th centuries, cities were not prepared to host such a big population that was, at the same time, rapidly growing. During the construction of these buildings the lack of urban planning was definitely missing, because the whole process led to more overcrowded buildings and overall low living standards.

Nowadays, urban design is more than organizing roads and buildings, it focuses on creating and managing sustainability, efficiency and inclusivity. Modern urban design aims to bring benefits to the quality of life while dealing with important matters such as climate change, population growth, rapid urbanization and new technological developments. Nevertheless, contemporary urban design works toward building “walkable cities”, as an approach to reducing traffic congestion, air and noise pollution, as well as health problems and environmental impacts (6).

The concept of walkable cities

Nowadays, people are seeking to speed everything up. They want to find an alternative, a faster way, for most tasks in their lives. A study shows that 7% of all car trips were short enough so they could have been walked (7). Therefore, due to the continuous need for new developments, technologies and the desire to always move quickly and faster, people rarely pick the alternative

of walking or cycling to their destination, if possible. It is very often ignored and replaced by the car, which results in more traffic, air pollution and overcrowding. Because of this, ultimately, the layout of a city becomes more "car-centred" (putting focus on roads), rather than "people-centred" allowing citizens to experience wider and better thought sidewalks (8). Nevertheless, the removal of cars from daily life is an impossible endeavor. What we can improve upon is our reliance on vehicles as a sole means of transportation.

A study shows that walking is the most powerful and accessible form of exercise. By walking, cardiovascular health improves, chronic diseases can be mitigated and "healthy aging" can be further promoted. In summation, walking helps the overall mental well-being, as it effectively reduces stress, anxiety, and offers "space for reflection or social connections" (9).

In ancient times, towns were built so that pedestrians could arrive anywhere on foot. However, in the 20th century, this concept morphed. Due to the accelerated rise of automobiles, the roads for cars widened, meanwhile the pavements, dedicated for pedestrians, dramatically shrunk. For example, in New York City, broad sideways were narrowed so additional space for automobiles was created (21). Furthermore, car-oriented cities impact many aspects of our daily life negatively, such as mental health, environment and climate change, social inclusion, and economic vitality (9).

Non-walkable cities negatively influence our well-being, social inclusion and environmental sustainability because the structure directly changes how individuals interact and move. When pedestrian infrastructure is limited and insufficient and destinations can be far, residents end up relying on their own vehicles. This, automatically, reduces the physical activity a person does, as well as, missing the opportunities to social interact and create new connections. These two aspects strongly support and reinforce our psychological well-being (23). Moreover, studies show that car dependency leads to an increase in air pollution, due to the fuel consumption (24). Furthermore, all these aspects demonstrate that if a city is car-orientated then this can contribute to health problems, as well as leading to negative effects on well-being and mental health.

Despite all these negative aspects, the aim of this paper is to demonstrate why citizens feel more comfortable in cities that are considered walkable. Researchers found that a lot of individuals opt for living in walkable neighborhoods because it provides a combination of health benefits and social interactions (22). People who live in such communities enjoy their time there because they can easily find everything they need, such as their workplace, shops, markets and public transport, which helps them travel long distances while reducing the usage of a private vehicle. Citizens often associate walkability with public spaces that attract them and help them create great social interactions, while building a stronger sense of community (10).

Other research showed that people who move into new walkable cities are highly motivated by the knowledge that green spaces, public transport and other essential facilities are located within walking distance (11). Moreover, the findings noted that even if housing affordability plays a

very important role, if a prospective buyer wants a larger home this can lead them to opt for less walkable cities, such as suburban areas. On average, people clearly expressed a preference for communities that are accessible and support walking as a part of their daily life (11). Altogether, these studies demonstrate that walkable cities are much more attractive than car-centered cities, because they give people new opportunities when it comes to having social interactions, as well as providing great access to the basic facilities.

As established earlier, a walkable city is designed for citizens, providing them access to daily necessities, such as supermarkets or groceries stores, schools, workplaces, and parks, all within a short walkable distance, usually no more than a 15 minutes walk or bike ride (12). A main feature and one of the most important when deciding if a city is people oriented or not is the availability of green spaces. Very often, people talk about how beneficial it is to have these green spaces, however, rarely, do people actually spread awareness and take this matter into their own hands. A study shows that on average Americans, 342 million people (in 2025), who weigh about 4.2% out of the entire population (in 2025) (13), spend 87% of their time indoors and 6% in vehicles (14).

Researchers found that not being connected with nature, and spending so much time indoors, surrounded with artificial light can be detrimental. As a result, this can directly affect our well-being and mental health. Just by going for a walk or taking in the sun, nature improves well-being, especially in mitigating the phenomenon called brain fog (a symptom that involves trouble remembering, focusing or managing stress (15) (16). The study showed that it reduces the level of stress and anxiety and promotes overall happiness (16).

Nevertheless, green spaces do not only support our mental health but also the health of our environment. Trees absorb carbon dioxide and help purify the air. However, green spaces play a much more important role in cities. They help make a major difference by reducing the urban heat island effect. The urban heat island effect is a phenomenon in which urban areas experience significantly higher temperatures than surrounding rural areas (17). This happens when buildings and other surfaces trap heat, making the area considerably hotter. As a result, this also reduces the air quality, which can later, during a heatwave, increase the health risks people can have, as well as leading to a higher level of energy used.

All these aspects directly contribute to the benefits that walkable cities offer citizens, such as sustainability, overall health, and social inclusive urban environments. By minimizing the number of private vehicles, the air quality increases, stronger social connections are born and our well-being and mental health improves.

Social interaction and the sense of community

A social interaction is when 2 individuals or more initiate and participate in a conversation, expressing different opinions and giving arguments (18). With the help of networking spaces,

people can engage in conversations, providing ideas and thoughts, due to the existent “social nets formation” (19). The aim of these urban public spaces is to promote social participation, communication skills improvement, self-esteem and a sense of belonging. Some of these quality spaces and pedestrian-orientated design can be found in streets and boulevards, plazas, squares, play grounds, malls, and many other forms where gatherings can take place.

A researcher proved that individuals feel a strong bond with a particular place and with their people when they interact with them. In addition, another crucial factor influencing human behaviour is the physical appearance an environment gives you. Studies demonstrated that the visual attractiveness of a place is one of the key features that encourages people to engage with, walk through or remain (19). As a result, urban public spaces provide individuals, as well as communities, with an overall good well-being. As cities start to grow and become more populated, the need for high-quality public spaces, where different social groups can gather and spontaneous encounters can take place, increases. Nevertheless, just as important is experiencing the sense of community and belonging. These public spaces play a significant role, especially in cities with a variety of ethnicities. By offering access to everybody, regarding their age, race or socioeconomic status, places like this provide a greater sense of belonging, as well as strengthening the community cohesion (20).

Studies also showed that well designed public spaces promote improving physical health. Having access to parks, plazas, squares and others encourages individuals to practice a certain form of exercise. Whether it is walking, jogging, running or even cycling, they all bring lots of benefits for physical health and mental well-being. And so, if cities provide pedestrians with walking paths and designated outdoor areas for exercise, they clearly show that they support a healthy lifestyle, regarding all ages (20).

In 2026, for the annual edition of Happy City Index, Copenhagen, Denmark, was ranked number one as the happiest and most livable city due to their overall outstanding quality of life, universal healthcare, exceptional work-life balance, and strong green infrastructure. Copenhagen received the highest score, especially for the environments and citizens categories (25). However, its reputation focuses not entirely on its physical environment, but also on its urban landscape. Copenhagen’s urban landscape is not a recent development; rather, it is a result of decades of decisions and long-term planning. A key factor is the municipal government. The city managed to transform because it was backed up by strong political decisions. Copenhagen flourished when different levels of government started cooperating, and together helped maintain continuity in urban development (25). One of Copenhagen’s strongest suits consists of their cycling lanes. They are clear evidence that infrastructure can be designed to be people-centred, rather than car oriented. Cycling was then incorporated in the city’s transport and became a very attractive and practical way of travelling around. Nonetheless, it is important to remember that these changes cannot take place without governmental support (26).

In contrast, one of the lowest scoring cities in the Happy Index 2026 edition was Bucharest. Bucharest was ranked 247 out of 251 cities, being considered as one of the lowest-performing urban centres in the world (29). The results were analyzed based on critical factors that contribute to citizens' well-being, QOL, sustainability, governance, health, economy, environment, and mobility (29). For Bucharest, however, the result is much more complex than it seems. Despite the significant cultural scene it gives and the increasing technology sector, as well as the growing international visibility it gets, Bucharest fights with many recurring and constant issues related to traffic congestion, air pollution, access to public spaces, and the overall badly designed infrastructure. As a result, all these aspects directly influence the well-being and QOL people have. These findings highlight a significant reminder that urban success is mainly measured by the performances a city serves (infrastructure, services, green spaces, and other urban policies), and by how well a city works for the individuals who actually live in it (29).

Differently from Copenhagen, a very important box that Bucharest does not tick is that the city's infrastructure is car-orientated, and not people centred. Nevertheless, due to the increased number of cars or other vehicles pedestrians also have to face other challenges. For instance, frequently sidewalks are occupied by parked cars, limiting pedestrian infrastructure as well as reducing access and comfort. Even though parking regulations and other policies exist and are ongoing, people voluntarily choose not to respect them, perpetuating the rights citizens have when walking on the sidewalks. Consequently, Bucharest's existing infrastructure reflects that it has a strong reliance on cars, not focusing on pedestrians and their overall comfort (30).

Ultimately, urban public places play a critical role in enhancing and improving an individual's QOL (quality of life), as well as their sense of belonging. Moreover, it increases community cohesion, while also improving physical and mental health. Places thought especially for citizens, that are well-designed and visually attractive can directly encourage people to gather and interact with each other. Copenhagen is a prime example that provides all these benefits, illustrating a higher QOL, as well as a strong green and cycling infrastructure.

Decision-making in urban development

As previously discussed, urban planning is a technique through which cities can express themselves and show their identity. Strategies and critical decisions are taken so that a whole city's system (public transport, public spaces, green spaces) is enhanced and people oriented. When thoughtfully implemented, urban design becomes a powerful tool for long-term resilience, facilitating not just environmental sustainability, but also stronger civic engagement and cultural identity.

However, before any project takes place and is implemented in the city, it must go through a process where it is analyzed by authorities. The local governments in a city play a significant role in urban development. Local governments work alongside urban planners to ensure that urban projects are being developed rightly and, moreover, promote the wishes that communities

have (28). Meanwhile, the collaboration between governments, local governments or municipal authorities and urban planners is crucial for the development of urban, people oriented cities, as well as for the infrastructure of city projects. Nevertheless, aspects of effective governance perfectly align to promote a greater quality of life for citizens, but also to boost the economic growth of cities (28).

Nonetheless, just having an idea for improving a city's overall urban design and environment is not sufficient so it can become reality. Whether it is proposed by an individual or an organisation, an urban development project requires significant financial resources, as well as coordination and approval (28).

However, any urban development decision involves the public authorities, urban planners, people with technical expertise, and residents (the community). The role of local governments is to establish the priorities a development has, while the professionals and the public help shape the decision. An effective urban policy needs cooperation between different levels of governments and the citizens community. When an urban problem is identified, the process begins, whether it is a challenge regarding public transport, public or green spaces, or other pedestrian conditions. Authorities and planners work together to find possible solutions by analyzing and considering residents' needs, so that comfort is established, as well as considering potential effects each solution may have. However, a crucial factor that must not be missed is that different solutions may benefit different categories of people, and as a result, decision-making is a much more complex project than initially thought (31).

In addition, citizens can directly contribute to these decisions through public meetings and consultations. Their participation is crucial because they help provide information and key aspects of their everyday life. Therefore, citizens' experiences help authorities to understand how certain proposals may affect communities (31).

Once a decision regarding an urban development has been taken, its implementation is considered a complex process of laws, regulations, and planning. Planning regulations consist of how different land can be used and the establishment requirements, meaning how tall buildings can be, their density, parking and green spaces, and other activities that are allowed in the area. These rules ensure that each project fits perfectly within the initial structure, while also validating the objectives of the city. Moreover, researchers found that land-use planning gives urban governance a better chance at balancing environment, social, and economic interests (31).

A practical example of a city that can be shaped by governmental regulations is Helsinki, Finland. In Helsinki it can be easily observed that it uses intricate land-use plans (to determine what different areas can be used for, and the amount and height of projects, streets, green areas, and location of parking spaces). Furthermore, this city's building regulations must meet specific and critical "green factor" targets, in order to ensure the health of the environment and incorporate other nature-based solutions into construction. As a result, these measures prove how

governmental planning decision can shape the harmony of the city, in our case Helsinki, and moreover demonstrate that sustainability objectives can be morphed into concrete requirements that influence the development of certain buildings, transportation, and green spaces, ensuring to integrate them into the urban environment (32).

Therefore, these regulations create connections between political decisions and the appearance of cities and, in addition, urban regulations are not just some simple technical rules, they can be instruments that authorities (mainly governments) use to shape the way cities are redesigned, how they function, and how residents experience them.

Discussion

Taken together, the findings of this paper demonstrate that urban design plays a significant role in shaping human behaviour, social interaction, and overall quality of life. The physical structure of a city influences not only how efficiently individuals move from one location to another, but also how they experience and interact with the urban environment. Elements such as street design, pedestrian infrastructure, green spaces, public areas, cycling networks, and transportation systems can influence everyday decisions regarding mobility, recreation, and social participation. Urban design should therefore not be understood simply as the physical organization of buildings and infrastructure. Instead, it represents an important environmental factor capable of creating conditions that either encourage or discourage behaviours associated with individual and collective well-being.

One of the clearest relationships identified throughout this research is the connection between walkability and opportunities for social interaction. When urban environments provide accessible, safe, and well-connected pedestrian infrastructure, residents are more likely to consider walking as a realistic alternative to private transportation. Increased pedestrian activity can consequently result in individuals spending more time within public spaces rather than experiencing the city primarily through private vehicles. This creates more opportunities for spontaneous encounters, informal conversations, and repeated interactions between members of the same community. Over time, such interactions may contribute to stronger feelings of familiarity, social connection, and belonging. Walkability can therefore be considered not only a transportation issue, but also a social characteristic of the urban environment.

However, the relationship between urban design and behaviour is not determined by the simple existence of infrastructure. A pedestrian pavement, for example, is unlikely to encourage walking if it is perceived as unsafe, poorly maintained, disconnected from important destinations, or difficult to access. Similarly, the presence of a park does not automatically result in increased recreation or social interaction. The findings discussed throughout this paper suggest that the quality, accessibility, safety, maintenance, and connectivity of these spaces are central to determining whether people actually use them. Effective urban planning therefore requires

consideration of how different elements of a city function together rather than evaluating each element independently.

Green and intentionally designed public spaces appear particularly relevant to this relationship. Parks, squares, pedestrian areas, and other shared environments can provide opportunities for physical activity, relaxation, recreation, and social interaction. They can also provide residents with spaces in which participation does not necessarily depend on commercial consumption. Nevertheless, their effectiveness depends considerably on their integration into everyday urban life. A well-maintained green space that is difficult to reach may have a more limited impact than a smaller but highly accessible public space located within walking distance of residential areas. This reinforces the idea that successful urban design depends not simply on the quantity of infrastructure available, but on its usability and relationship with the surrounding urban environment.

The case-study comparison between Copenhagen and Bucharest further illustrates how different approaches to urban planning can result in considerably different experiences of city life. Copenhagen provides an example of an urban environment in which cycling infrastructure, pedestrian accessibility, public transportation, green spaces, and social areas have been integrated into the broader functioning of the city. The city's high position in quality-of-life and happiness rankings reflects a broader model in which sustainable transportation and human mobility have become important components of urban planning. Rather than designing mobility primarily around private vehicles, Copenhagen has created infrastructure that gives residents multiple practical alternatives for navigating and experiencing the city.

Bucharest, by comparison, continues to face challenges associated with traffic congestion, pedestrian infrastructure, connectivity, and the availability and quality of public spaces. These challenges can affect residents' ability or willingness to walk, cycle, use shared spaces, or spend extended periods within the public environment. The comparison therefore demonstrates how planning priorities can influence everyday behaviour. When walking or cycling is inconvenient or unsafe, individuals may become more dependent on private vehicles. This dependence can, in turn, contribute to congestion and reduce the amount of everyday interaction occurring within public spaces.

At the same time, the Copenhagen–Bucharest comparison must be interpreted cautiously. The purpose of the case study is not to suggest that Bucharest should simply reproduce Copenhagen's urban model. The two cities have different historical developments, economic resources, geographical characteristics, institutional structures, and existing infrastructure. Policies that are successful in Copenhagen may therefore require significant adaptation before they could be effective in Bucharest. Instead, Copenhagen can function as a reference point for understanding how sustained investment in pedestrian infrastructure, public transportation, cycling networks, and accessible public spaces can contribute to a more people-oriented urban environment. The

broader lesson is therefore not that there is one universal model of successful urban design, but that planning decisions should respond to the specific characteristics and needs of the population they are intended to serve.

Furthermore, this research highlights that urban environments are shaped not only by architects and urban designers, but also by government institutions, regulation, funding, and political priorities. Even well-designed proposals depend on institutional support for their implementation and long-term maintenance. Decisions concerning transportation budgets, zoning, pedestrianization, green-space protection, infrastructure development, and public investment ultimately determine which forms of urban development become possible. Consequently, creating people-oriented cities requires coordination between design expertise and public policy. Urban design should therefore be understood as both a spatial and governance issue.

It is equally important not to overstate the findings. Urban design alone cannot solve every social, environmental, or well-being problem experienced by residents. An individual's quality of life is influenced by numerous interconnected variables, including income, employment, housing, health, education, personal circumstances, and social relationships. Similarly, constructing a park cannot guarantee that residents will exercise, and developing a cycle lane cannot ensure that individuals will stop using cars. Urban design does not directly determine human behaviour; rather, it can shape the opportunities, incentives, and barriers that influence behavioural choices.

This distinction is particularly important to the central argument of the paper. The value of effective urban design may lie less in forcing particular behaviours and more in expanding the range of realistic choices available to residents. A city in which walking, cycling, public transportation, recreation, and social interaction are safe and accessible gives individuals greater freedom to incorporate these behaviours into their everyday lives. Conversely, when urban environments make such activities difficult, unsafe, or inconvenient, residents' choices become more restricted.

Ultimately, the findings indicate that effective and carefully considered urban design can contribute to environments that support both individual and collective well-being. Walkable streets can encourage physical movement and increase opportunities for spontaneous interaction; accessible green spaces can provide environments for recreation and relaxation; well-designed public spaces can support community activity; and integrated transportation networks can increase mobility while reducing dependence on private vehicles. The effects of these interventions depend heavily on accessibility, quality, safety, connectivity, and the broader social and institutional context in which they are implemented.

Therefore, designing better cities ultimately means designing environments around the people who inhabit them. Urban planning cannot determine how individuals will behave or guarantee their well-being, but it can create the conditions in which healthier, more social, and more

sustainable behaviours become easier and more accessible. As cities continue to grow and respond to changing environmental and social demands, understanding this relationship between physical space and human behaviour will become increasingly important. Urban design should consequently be treated not merely as a question of how a city looks or functions, but as an important component of how people experience everyday life within it.

Works Cited

1. Urban Design Group (2019) What is Urban Design
2. Agnieszka Olszewska-Guizzo (2024) How Urban Design Can Impact Mental Health & Well-being
3. Susan S. Fainstein (2026) The era of industrialization
4. College Sidekick Renaissance Architecture
5. Rachel Cole, Gavin Turrell, Mohammad Javad Koohsari, Neville Owen & Takemi Sugiyama (2017) Prevalence and Correlates of Walkable Short Car Trips: A Cross-Sectional Multilevel Analysis
6. Katarzyna Turon, Piotr Czech, Michal Juzek (2017) The concept of a walkable city as an alternative form of urban mobility
7. Patricia Fierro-Newton (2025) Streets for People: The Case for Walkable Cities
8. Katherine Bookfield (2016) Residents' preferences for walkable neighbourhoods
9. Paul Whybrow, G. Bramley, Caroline Brown (2021) 'It's just one of those natural progressions': Stories of relocating to neighbourhoods of high and low walkability
10. M. Martinez Euklidiadas (2023) What Makes A Walkable City
11. USAFacts (2026) How many people live in the US?
12. Tristan Roberts We Spend 90% of Our Time Indoors. Says Who?
13. American Brain Foundation (2025) Brain Fog
14. Queensland Government Department of the Environment, Tourism, Science and Innovation (2025) Why are green spaces good for us?
15. My NASA data Urban Heat Island
16. UNESCO Social Interaction
17. Vahid Bigdeli Rad, Ibrahim Bin (2013) The Role of Public Spaces in Promoting Social Interactions
18. Roberto Gasparini (2025) Why Public Spaces Matter
19. NYC Public Design Commission (2022) Designing New York: Streetscapes for Wellness
20. Xumei Zhu, Zhipeng Lu, Chia-Yuan Yu, Chanam Lee, George Mann (2013) Walkable communities: Impacts on residents' physical and social health
21. Catherine Sundling, Marianne Jakobsson (2023) How Do Urban Walking Environments Impact Pedestrians' Experience and Psychological Health?
22. Abdulla Baobeid, Sami G. Al-Ghamdi, Muammer Koc (2021) Walkability and Its Relationships With Health, Sustainability, and Livability: Elements of Physical Environment and Evaluation Frameworks
23. Happy Index (2026) Healthy City Index
24. Til Koglin (2015) Organisation does matter - planning for cycling in Stockholm and Copenhagen
25. KXM Group Insights On Danish Working Culture
26. Behar Font (2025) How Urban Planning Drives Sustainable Growth In Modern Cities

27. BR Business Review (2026) Bucharest ranks among the world’s lowest-scoring cities in Happy City Index 2026
28. Raluca Popescu (2022) The culture of parking on the sidewalks
29. OECD (2025) Global State of National Urban Policy 2021
30. Sirkku Juhola (2018) Planning for a green city: The Green Factor tool

Non-Familial Transmission of Rice Cultivation Knowledge in Aizu-Wakamatsu, Fukushima By Rikako Umezawa

Abstract

Rice cultivation in Japan has been handed down within families, on the same land, across generations. In Aizu-Wakamatsu, Fukushima, that channel is closing. Children commonly leave the area for urban cities for schooling at eleven or twelve and do not return to live, and no member of this study's principal case community inherited a farm from a parent. This article asks how knowledge of rice cultivation is transmitted when the traditional channel, inheritance within a family across generations, is disappearing, and what this reveals about what "tradition" requires in order to survive.

Inheritance and transmission, once a single event, have separated. Land passes to people never taught to work it, and teaching passes to people who will inherit nothing. Where transmission continues, it runs through three non-familial channels: structured apprenticeship between unrelated adults, voluntary in-migration by people with no ancestral claim to the land, and community participation extending to seasonal returnees, part-time farmers, and consumers. On the second question, tradition here has never required descent. It has required a teaching relationship and the occasions on which one can form. What is endangered is not any particular practice but those occasions.

The findings draw on three months of ethnographic fieldwork conducted in 2026: three recorded semi-structured interviews, twenty-nine informal conversations, and participant observation in rice transplanting, community restoration work, and direct produce sales. The article identifies trust as the condition governing whether the substitute channels work. A fixed apprenticeship term functions as a demonstration of reliability rather than as a period of skill acquisition, and instruction without trust does not transfer tacit knowledge, which is why institutional training fails where community apprenticeship succeeds.

1. Introduction

On my first night in Aizu-Wakamatsu, I was sent to a paddy to help regulate its water level. I was carrying a water container with my arms until a farmer told me to use my hip and to find my balance before my grip. That instruction exists nowhere in writing. It existed because someone stood close enough to say it.

For generations, the judgment behind that instruction moved from parent to child, on the same ground. That route is closing. By 2022, 80.4 per cent of local farmland had been concentrated among core operators (Aizu-Wakamatsu City Agricultural Revitalization Council, 2024), and the schooling pathway removes children before succession is a question they could answer. The expectation is that the knowledge closes with the channel. It has not.

This article asks how knowledge of rice cultivation is transmitted in Aizu-Wakamatsu when the traditional channel for transmission, inheritance within a family across generations, is disappearing, and what this reveals about what "tradition" actually requires to survive.

It answers that inheritance and transmission, historically a single event, have come apart, and that what is endangered is neither a practice nor a population but the teaching relationship and the ordinary occasions on which one forms.

A farm can report a successor and transmit nothing. That is a problem anywhere farming has been taught at home rather than in a classroom, because the figures used to measure succession will not catch it. The case also puts a question to anthropology: whether knowledge like this needs a family to move through, or only needs one person willing to teach and another willing to learn.

Section 2 positions the study against three literatures, on rural demographic decline, on farm succession, and on tacit and traditional ecological knowledge, and identifies the gap where they meet. Section 3 sets out the method, the composition of the sample, and my own position in the field. Section 4 develops three themes from the material and answers the question. Section 5 states what follows for how agricultural decline is measured. The argument runs in one direction throughout. What succession scholarship treats as a single variable is, in this region, two variables, and the second one fails quietly enough that instruments built for the first will not register it.

2. Literature Review

2.1 Rural decline and demographic change in Japan

Peer-reviewed work frames Japanese rural shrinkage as a compound process, not a simple fall in population. Matanle and Sato (2010), working from Niigata Prefecture and Sado Island, show demographic contraction deepening through out-migration combined with natural decrease, and extending toward regional urban centres. Feldhoff (2013) locates the *genkai shūraku* (marginal settlement) phenomenon within the declining manageability of ageing non-metropolitan communities. The analytical units throughout are people and land: households remaining, average producer age, area of abandoned cultivation.

Decline has not produced movement in only one direction. Population-geographic research distinguishes return or U-turn migration from non-return movement out of metropolitan centres, discussed as *den'en kaiki* (Inoue et al., 2022), and Klein's (2020) multisited ethnography documents urban-origin migrants relocating to rural Japan in pursuit of alternative work and life. This literature establishes that outsiders arrive. It does not ask what happens to place-specific agricultural knowledge once they do.

This article accepts that the demographic decline is real and locally severe. It questions the assumption left unexamined: that the loss of agricultural knowledge is the same curve seen from another angle.

2.2 Farm succession and agricultural knowledge transmission

In this literature, identifying a successor is the variable determining whether a farm survives, and the family is where succession is assumed to happen. Research on rural families in Yamanashi finds farming households reproducing single-heir, *i.e.,* based on property succession even as co-residence and participation diversified (Tsutsumi, 2001).

The nearest prior work is Uchiyama et al. (2008), and the relationship requires stating precisely. They observe that succession research has concentrated on physical assets while neglecting intangible ones, and examine the transfer of managerial skill and farm-specific knowledge through patterns of delegated responsibility. Their unit of analysis is the succession process *after a successor has been identified*, and that successor is a family member. This article accepts their premise that intangible transfer is the interesting part of succession, and removes the condition their design depends on. Where they ask how knowledge is delegated within an identified family succession, this article asks what becomes of it when no family successor is identified at all. Japan-focused work on entrants from non-farming households documents public training, financial support, land access, and relationships with municipalities, cooperatives, and experienced farmers (Egawa, 2005; McGreevy et al., 2019). These describe how someone gets into farming, not how the judgment particular to a field is passed on.

The gap is therefore precise rather than absent. Succession scholarship asks *whether a successor exists*; new-entrant scholarship asks *how someone enters farming*; Uchiyama et al. ask *how knowledge moves once a family successor is in place*. None asks whether the successor was taught in cases where the family channel has failed. The questions were rarely distinguished because under family succession they collapsed into one event: the child who inherited the land was the child who had spent eighteen years beside it. Once transmission moves outside the family, they separate. A farm may satisfy the first question and fail the second.

2.3 Traditional ecological knowledge and community learning

The knowledge at issue is mostly not propositional. Judging whether to open a water gate before a cold night; knowing which corner of a field drains too quickly; reading a compost pile by hand. These resist codification. Polanyi's (1962) account of tacit knowledge begins from the premise that competent judgment exceeds what its holder can state, and treats knowing as skilled integration within action, not information detachable from practice. Lave and Wenger (1991) conceptualise learning as participation in social practice: through legitimate peripheral participation, newcomers enter real activity at the margins and move toward fuller membership. Traditional ecological knowledge, on this account, is not a closed archive. Berkes et al. (2000) show it is generated, accumulated, and transmitted through social mechanisms and used to interpret environmental feedback, which makes such systems comparable to adaptive management under uncertainty.

This constrains what preservation can mean. If traditional knowledge is defined by its capacity to revise under feedback, preservation cannot consist in arresting it. A practice held fixed has become a display of a practice rather than the practice. Thornbury's (1994) study of Japan's Cultural Properties Protection Law makes the danger explicit: designation can turn folk performance into a fixed cultural icon, while its continued vitality depends on leaving room for change.

One question must be settled before this article uses the word "tradition." Practices presented as immemorial are often recent constructions, and a claim to continuity can be a claim to authority made in the present.

Hobsbawm and Ranger (1983) show that traditions presented as ancient may be recent constructions that establish continuity with a selectively imagined past. Robertson's (1988) study of postwar *furusato-zukuri* gives this critique a specifically Japanese form: governments, planners, and advertisers transformed the nostalgic image of the rural "old village" into a politically meaningful symbol of authenticity and collective belonging.

Aizu sits inside the imagery Robertson describes, and so does any account written about it.This article treats the critique as a constraint, and a useful one. It does not claim an unbroken four-hundred-year practice; Section 4.1 documents the opposite. It claims that a documented 1684 text describes a method of testing, that farmers today independently describe the same orientation, and that whether an unbroken line connects them is neither establishable nor necessary. Holding to that narrower claim is what keeps the argument clear of the nostalgia Robertson analyses, since what is at issue is not the antiquity of a practice but the survival of a relationship, and relationships can be newly built without being fraudulent.

Together these literatures supply the mechanism succession scholarship lacks. If agricultural knowledge is largely tacit, transmission needs co-presence, repetition, and permission to join real work, not simply a relationship of inheritance. That generates a testable expectation: transmission should survive the collapse of family inheritance wherever co-presence and permission can be reconstructed by other means, and fail wherever they cannot.

3. Methodology

3.1 Why transmission requires ethnography

Two anthropological problems meet in this case. Enculturation asks how a person acquires the competences and dispositions of the community they are born into, a process Herskovits (1948) distinguished from the cross-cultural borrowing of acculturation. Cultural reproduction asks how a way of life persists as the people carrying it are replaced, and Bourdieu and Passeron (1977) treat the transmission of inherited dispositions as the mechanism through which a social order reproduces itself across generations. Rice cultivation in Aizu reproduced itself through the household for generations, and that mechanism is failing. Asking what replaces it is a question about cultural reproduction once its ordinary vehicle has been removed.

This shapes the method. Much of what a farmer here knows is not propositional but embodied, distributed across judgment, timing, and trained perception of a particular landscape. Such knowledge is not reliably elicited by questioning, because informants perform it more readily than they state it, and because asking someone to articulate a competence often produces a rationalisation. It becomes visible in the doing. Participant observation is therefore what makes the object visible at all, and staying through a full cycle of the agricultural year follows from the same reasoning.

The study aims at an emic account: transmission in the terms used by those conducting it, including their own criteria for who is worth teaching. Those criteria proved central (Section 4.2). Where the analysis moves to etic terms, such as modelling the fixed apprenticeship term as a costly signal, these are the author's constructions and are marked as such.

3.2 Research design

This is a qualitative case study centred on Aizu-Wakamatsu, with the natural-farming community 無の会 (Mu no Kai) as its principal case, not its entire subject. Fieldwork ran for three months in 2026. I stayed with a household connected to the community and combined semi-structured interviews, participant observation, and analysis of government, municipal, and historical documents. Access was obtained by chain referral.

All personal names in this paper are pseudonyms, and identifying details of individuals have been altered or withheld. Where a person's published writing is quoted, it is cited to the organisation rather than to the individual author, so that the pseudonyms cannot be resolved through the reference list.

Across three months I spoke with thirty-two people. Three of these conversations were recorded semi-structured interviews, conducted with prior informed consent obtained verbally before recording began, after I explained the purpose of the research, my institutional context, and how the recording would be used. All three were adults speaking about their own professional work: Atsuhiko Uno, a member of 無の会 (8 July 2026); Hibiki Tokumoto, an agricultural operator in Tottori (2 August 2026); and Shigeru Kobayashi, a rice farmer in the municipality of Aizu-kitaaizu-machi (14 August 2026). Interviews were conversational rather than uniform, which let each speaker define the problem in his own terms and correspondingly limits direct comparison.

The remaining twenty-nine further conversations were informal and unrecorded, arising within participant observation: shared work at 田植え (rice transplanting), a 町おこし (community revitalization) project, walks with older residents, shared meals, and direct produce sales. Beyond the three recorded interviews, it comprised the founde of 無の会 Kenji Kojima, the apprentice Masataka Saito, further members, municipal staff, a primary school teacher, and local farmers outside 無の会; three further members of the community; the household I stayed with; at least two municipal staff; one primary school teacher; one volunteer guide to local history; at latest three farmers working outside 無の会; adult staff at a daycare centre; and the remainder met through direct produce sales and activities.

Three further distributions describe the sample. Five of the thirty-two were women, 16 per cent, and all three recorded interviews were with men. Ages were recorded for twenty-nine of the thirty-two: six in their thirties, eight in their fifties, ten in their sixties, and five aged seventy-five or over, so fifteen of the twenty-nine were sixty or older. Six of the thirty-two were members or remote members of 無の会 and twenty-six were not, though most of the twenty-six were reached through the community's own network rather than independently of it. None of these proportions describe the farming population in Aizu-Wakamatsu. A sample of thirty-two obtained by chain referral through one community describes the reach of that chain, and the distributions are reported here so that the claims in section 4 can be read against the shape of the material that produced them. Where a distribution bears directly on a specific claim, it is raised again at that point.

No informal conversation is quoted verbatim: such material appears as summarized field observation attributed by role rather than name, which is the primary safeguard against reputational risk in a community where roles are identified. For most of the twenty-nine, I know only approximate occupation and role, not names. No participant was a minor or misled about my purpose. Field notes and transcripts were coded thematically, and the categories organizing section 4 came out of the material and were not imposed on it.

3.3 Positionality and reflexivity

I am Japanese, raised in Tokyo through middle school and subsequently educated abroad, and a pre-university student, not a credentialed researcher. Five features of that position bear on the findings.

My linguistic competence was partial in a way that matters here. I conducted every conversation in Japanese without an interpreter, but in Tokyo-Japanese. The register in which informal knowledge actually travels, the side between two farmers who have known each other for forty years, remained substantially closed to me across three months.

Participation was also the point at which I could measure my own distance from the material, and that is worth reporting rather than assuming. Early in the fieldwork I watched a farmer adjust the water level in a paddy by eye, judging a quantity rather than measuring it, and I could not see what he was seeing. The judgment was visible as an action and unavailable to me as knowledge, which is the article's object, stated as a personal failure. Entry into shared work was slow for reasons that were social before they were technical: I participated in 4 mornings of 田植え and six days of 町おこし, plus some help in picking up and selling vegetables in person. What that delay taught me is the same finding section 4.2 reports from the receiving end, arrived at from a much shorter distance. I was given explanations after I had shown up repeatedly, not because I had become competent.

The access chain ran almost entirely through one community, so the sample is closer to a map of 無の会's social network than of Aizu-Wakamatsu. People indifferent or hostile to natural farming were unlikely to reach me, and the finding that community apprenticeship succeeds where institutional training fails was heard from inside the party that benefits from that comparison.

My informants also agree with each other to a degree that should be explained rather than taken as confirmation. That convergence may reflect the access chain, or my questions, or the retrospective character of accounts given by people explaining their own life choices. I guard against the first by including Kobayashi and Tokumoto, who sit outside 無の会 and, in Tokumoto's case, outside the region entirely. The convergence survives that test, which is why I retain it. A reader should nonetheless register that this study found no dissenting voice on its central point, which is not the same as there being none.

Finally, I arrived with a prior commitment to an endangerment framing and a romantic disposition toward agricultural labour. A narrative in which knowledge survives through devoted apprentices was more appealing to me than one of simple disappearance, and I was predisposed to find it.

4. Discussion

The discussion develops three themes across the interviews instead of reporting each in turn. Succession and transmission have separated (4.1); trust governs whether the substitute channels function (4.2); and what passes through them is a method rather than a fixed technique (4.3). Section 4.4 states the answer these produce.

4.1 Succession and transmission have separated

For generations, receiving the land and learning to work it were one event. They are now two, and each can fail without the other.

Masataka Saitō inherited it without being taught. Born in 1984 in Niigata, he spent roughly twenty years as an automotive mechanic in Saitama before he and his siblings unexpectedly took over a relative's paddy. He did not then farm it. He entrusted it to a neighbouring farmer while he decided whether he was ready, then joined 無の会 in April 2025 on a fixed two-year term, intending to farm independently from 2027. His account of what he did and did not inherit is precise. He remembered the visible sequence of rice farming, planting, harvest, the shape of the year, from growing up around it. He lacked the judgment underneath: when to adjust water level, how to respond to a given week's weather, how to read a change in soil. His relatives had farmed with chemical inputs while he wanted to farm without them, so even the sequence he had absorbed belonged to a different method. Observation produced familiarity, not competence, and he treated the decision to farm as separate from and later than the decision to accept ownership.

Kobayashi runs the opposite way: knowledge and no heir. He is seventy, has farmed for fifty years since beginning in his father's paddies at twenty, and now works about six hectares, his family's land combined with paddies entrusted by neighbours who could no longer work them. Neither of his children will take over, and he is honest that his own generation told them farming was hard and leaving was better.

> 「でも、家の子どもでなければならないとも思っていません。やりたい若い人がいるなら、土地のことも水のことも、失敗したことも、全部教えたい」
>
> "But I don't think it has to be my own child. If there's a young person who genuinely wants to do this, I want to teach them everything: the land, the water, my own failures, all of it." (Kobayashi, personal interview, 2026; working translation)

Kobayashi's case inverts the sequence Uchiyama et al. (2008) describe, in which delegation of managerial responsibility follows the identification of a family successor. Here ownership arrived first and unbidden, and he sought instruction outside the family, from a teacher unrelated to him. The intangible transfer those authors treat as the substance of succession is still occurring. It has detached from the succession event.

What this does not yet explain is why. Kobayashi wants to teach. Saitō wanted to learn. The willingness is present on both sides, so the failure is not one of willingness. It is that the situations in which teaching used to happen have been removed.

Two removals account for most of it. The first is educational. The clearest account of decline I obtained came from a primary school teacher who does not farm. One primary school serves the neighbourhood, with roughly five children per grade and in some years two or three.

There is no middle or upper secondary school in the area, so when children finish primary education, the family moves, or the child leaves for a larger city. Departure is not total: many return during school breaks to help grandparents farm, which leaves labour without residence and presence without succession. Most writing on rural Japanese agriculture, including this project's initial framing, treats youth departure as an evaluation of farming: the work is hard, the income unreliable. That explanation assumes the young person is present to evaluate. Departure here is imposed at eleven or twelve by the absence of a classroom. By the time succession becomes a live question, the potential successor has spent six or ten years elsewhere. The farm did not lose a candidate to unattractive conditions. It lost one to an education system. Agricultural subsidy and consolidation support operate on the farm; neither operates in the classroom, where the timing of departure is set.

The second removal is the collapse of shared labour. Kobayashi contrasted 1951 with the present, and although his figures conflate two units, the corrected comparison holds: MAFF records 200.7 direct labour hours per 10 *a* in 1951 falling to 21.16 in 2021, which converts to roughly 251 eight-hour labour-days per hectare reduced to about 26 (MAFF, n.d.-a; e-Stat/MAFF, n.d.-b). Two hundred fifty-one days of shared physical labour is also two hundred fifty-one days of proximity, in which an older farmer's judgment about wind, water, and soil is visible to whoever works alongside, asks about, corrects, and absorbs without either party calling it teaching. Twenty-six days compresses that by nearly a factor of ten. Kobayashi refuses to read this as a case against machinery: "Machines exist to help people. But if the machines come in, and the people leave, and finally even the paddies disappear, then what was it all for?"
Not every disappearance is a loss of this kind. A reciprocal planting group may vanish because a transplanter is faster, and little knowledge goes with it. A canal-clearing group cannot be replaced that way, because the task concerns a shared system no machine manages. Only the second is a transmission loss.

This matters for what the standard indicators can detect. Once these occasions are removed, transmission no longer follows descent and instead follows whoever remains inside an economic structure that still puts two people in the same field, which means the count of successors, farms, or consolidated hectares can hold steady, or improve, while the thing those counts are used to indicate has already gone.

4.2 Trust is the gate

If the occasions have gone, the knowledge should follow. In one part of this field site it has not, and the reason identifies the condition on which non-familial transmission depends. 無の会 farms roughly 15–17 hectares of paddy and 4–5 hectares of upland field across about eighty parcels. Its founder, Norio Kojima, born in Aizu in 1950, is the only member born in the region. He grew up in a rice-farming family but did not take over that operation, teaching high school English before returning to farming by another route. Every other member arrived from elsewhere, from a business school, financial-sector IT, twenty years of automotive repair, office work, documentary filmmaking. Their stated motives are unusual for an agricultural workforce: opposition to war, regenerative ecology, distrust of employed labour.

This inflow is not merely observed. It is a stated strategy. The organisation maintains a tiered structure for drawing outsiders into agricultural work short of full-time commitment: remote members visiting several days monthly while keeping outside employment, a formal *han-nō han-shō* ("half-farming, half-business") programme, and short-term seasonal helpers called *muraibito* (無來人). More than two hundred *muraibito* visited in 2024, and several subsequently relocated permanently (無の会, n.d.-a). Kojima describes the resulting dependence directly: "無の会 is not something that belongs only to its members. We have experienced change and growth that simply would not exist without the visitors who came" (Kojima, 2026).

Transmission also runs outward. Selling produce directly is only partly transactional: customers ask how something was grown, what to do with an unfamiliar vegetable, whether this year's rice differs from last year's, and answering is a small repeated act of teaching to people who will never farm.

The same region contains a formal, publicly funded route into farming that works far less well, and the difference is the most important thing this fieldwork found. Uno described institutional attempts, through JA, to train new and unconnected farmers, and described them as frequently failing for two compounding reasons. This passage reconstructs his account. Incoming farmers typically lack any network among established local farmers, so formal instruction does not connect them to the relationships through which knowledge travels. Older farmers frequently decline to teach anyone they judge insufficiently serious. Trust must be earned inside the village before knowledge is offered, and that earning is difficult.

Saitō corroborates the identical mechanism from the receiving end, independently. At farms where he worked before finding 無の会, he was used mainly as unskilled labour and rarely told why decisions were taken. Kojima's initial warning that farming would be difficult conveyed that substantive instruction would follow only after he had shown consistency across a full agricultural cycle. One must prove one will still be there in month eleven before anyone explains why month three's decision mattered. The two-year term, in his framing, concerns proving trustworthiness rather than the time required to absorb information.

Kobayashi states the same proposition in a different register:「最後は人なんだよ。人が人を呼ぶんです」—"In the end it's people. People bring people." His criteria for who is suited to receive teaching are not technical. Whether someone can absorb a bad season without blaming anyone, whether they can converse with the people around them, whether they can persist without an answer arriving promptly. Education does not appear.

His account also corrects the institutional picture. He has consulted JA for decades on fertiliser, varieties, machinery financing, and post-harvest inspection and shipping, noting that "there are things one household cannot do alone." The finding is therefore not that institutional agriculture fails generally. Institutions do well at the parts that can be written down and organised: storage, grading, distribution, credit. They fail at teaching, which cannot be. These are different tasks.

The composition of the sample bears directly on this comparison, and the claim should be read against it. Six of the thirty-two people in this study were members or remote members of 無

の会, and most of the remaining twenty-six were reached through its network, so the finding that community apprenticeship succeeds where institutional training fails is drawn largely from one side of the comparison it makes. No one who runs a JA training programme, and no one who completed one, is represented here. What holds the claim up is that Sakai and Terao sit outside 無の会 entirely and describe the same mechanism, and that Endō describes it from the receiving end at farms unconnected to the community. It stands on independent corroboration rather than on a balanced sample, and a study designed around the institutional side could overturn it.

This refines rather than confirms Lave and Wenger (1991). In their formulation, the newcomer's movement inward tracks accumulating competence: participation is legitimate from the outset. Saitō describes something the model does not anticipate. At his earlier placements, participation was peripheral without being legitimate in the relevant sense; he performed real work but was told nothing of why. What altered his position was not competence but persistence, demonstrated before competence was at issue. A fixed multi-year term therefore functions as a costly signal of intended permanence rather than as a training duration, which explains why shorter and cheaper formal programmes fail at the point where they should succeed. Formal channels can supply instruction; they cannot supply trust, and instruction without trust does not move knowledge of the kind Polanyi (1962) describes.

4.3 What is transmitted is a method, not a recipe

Once the gate opens, what passes through it? The answer determines whether outsiders can receive this knowledge at all, and it is the second half of the research question.

The *Aizu Nōsho* is a record of experiments. In 1684, Sase Yojiemon, a farmer and village leader in what is now Aizu-Wakamatsu, gathered his own trials alongside practices learned from other farmers, and because many farmers struggled with the written text, he recast the knowledge as *waka* poetry so it could be memorised (MLIT, n.d.). Remembering was tied to revising. This is the historical case Berkes et al. (2000) predict but rarely observe in documented form: Sase recorded his testing procedure and then engineered a mnemonic so the procedure, not only its conclusions, would survive transmission.

Kojima has published nearly the same argument, disclaiming both the rejection of modern agricultural science and the uncritical celebration of organic methods, and turning to Sase directly:「私たち会津の農民は『古きを訪ねて新しきを知る』ことです」— "We Aizu farmers live by 'inquiring into the old to know the new'" (Kojima, 2026). Farmers in an age of mechanisation, he continues, should extract what the old text discovered and realise it with modern tools instead of imitating the past.

Kobayashi arrives at the same position without reference to any text:

> 「農業は答えを覚える仕事じゃない。その年の田んぼを見て、自分で考えて、失敗したら直す仕事です」
>
> "Farming is not a job of memorising answers. It is a job of looking at that year's paddy, thinking for yourself, and correcting courses when you fail." (Kobayashi, personal interview, 2026; working translation)

Fifty years of experience, on this account, produces a habit of testing, not a stock of answers. Saitō applies the same logic prospectively: he does not intend to reproduce his relative's methods but to begin from the soil and water he actually has, test what he has learned on a limited area, and reduce chemical inputs gradually. He described 無の会 as a model to adapt rather than a manual to copy onto different land, which approximates the *Aizu Nōsho*'s own argument applied by someone who has not read it. What Kojima teaches him is procedural. Reading compost by temperature, smell, moisture, and texture. Managing water and weeds without defaulting to chemical intervention. Adjusting decisions paddy by paddy.

Tokumoto reached the same conclusion independently, from a different region and background. Operating roughly 100 hectares in Tottori, he came to farming late and describes his first years as sustained failure for a specific reason: he began at large scale without first learning from anyone who already knew how to farm in that place. What altered his trajectory was visiting organic farms across Japan and abroad and asking each what they did and why. His claim that "farming is a science" reads in isolation as a rejection of tradition; read against how he arrived at it, it is the *Aizu Nōsho* procedure. Both he and Uno insist that whoever intends to farm land must learn what is already known there before claiming standing to alter it. Descent confers that standing on neither man.

This is why the substitute channel can work at all. Were the content fixed technique, the loss of family succession would be terminal and the arrival of outsiders useless, since no outsider can receive a technique calibrated to land they did not grow up on. A method of observing, testing, and revising is portable to a newcomer in a way a recipe is not.

That portability is what this region has to rely on, because the alternatives are close to it. Roughly 61.8 percent of Japan's cultivated land is flat enough for large-scale consolidation, and much of Aizu's *chūsankan* terrain falls outside it (Ueno, 2023), so scale cannot substitute for local knowledge past a certain point. The sensing technology Tokumoto uses in Tottori requires capital that a small or mid-sized Aizu operation does not have. Where a flatter or wealthier region might buy its way partly around the transmission problem, this one cannot. Person-to-person teaching is not the sentimental option here; it is close to the only one.

4.4 Answering the question

Knowledge of rice cultivation in Aizu-Wakamatsu is transmitted, where it is transmitted at all, through three channels that have partially replaced family inheritance: structured apprenticeship between unrelated adults, in which a fixed term functions as a demonstration of reliability; voluntary in-migration produced by deliberately built infrastructure operating at a scale of hundreds of people annually; and community participation running through seasonal returnees, half-farming families, and consumers. None of this is inheritance. All of it is transmission.

As for what this reveals about tradition, it has never required descent here. It has required someone taught seriously enough to test, revise, and eventually teach again. This blocks one conclusion in particular: that preserving tradition means protecting elderly people and old methods from change. What these accounts describe is a teaching relationship that must exist

between at least two people, irrespective of their relation. An untaught heir loses as much as an empty field.

5. Conclusion

This study asked how rice cultivation knowledge is transmitted once family inheritance ceases to function as its primary channel. The finding is that the two now operate separately, and that where transmission continues it runs through constructed channels resting on a condition the succession literature does not model: trust sufficient for an experienced farmer to think a newcomer worth teaching.

The disciplinary significance is that agricultural decline should not be modelled as a single curve. Demographic loss and knowledge loss are separable and governed by different mechanisms. The mechanism behind knowledge loss identified here is the disappearance of *occasions*: the collapse from roughly 251 shared labour-days a year to about 26, and the closure of a schooling pathway that removes children at eleven or twelve. Policy instruments operate on ownership and scale. None operate on occasions. It also suggests an answer to the question put in section 1. A family is not what this knowledge requires. It requires two people in the same field often enough, one of whom has decided the other is worth teaching.

Three limitations bound these claims. The three speakers quoted at length are all men in formal or titled roles, and the work historically performed by women in Japanese farming households carries its own history of instruction that this study did not reach. Transmission was more often reported than observed: Kojima teaching Saitō is the central case here, and I witnessed none of it. And the account of institutional training failure rests on a partially damaged recording and was obtained entirely from people outside the institution concerned, though Saitō's independent testimony from the receiving end does some of the work a clean recording would have done.

Further research follows from each. The gendered channel requires dedicated study. The institutional side requires the counter-account this study lacks. Most consequentially, whether knowledge calibrated to a specific place survives transmission to people who did not grow up in it is answerable only across a longer timescale: Saitō begins farming independently in Shibata in 2027, on inherited land, using a method learned from an unrelated teacher two prefectures away. The contribution offered here is a distinction rather than a prognosis. Succession and transmission were treated as one variable because, for most of the period the relevant scholarship describes, they were empirically fused. In Aizu-Wakamatsu they are no longer fused, and tools built for the fused case misdescribe what is occurring. Tradition in this region has required a teaching relationship and the occasions on which one can form. Whether enough such occasions remain is an empirical question, and a different one from whether enough farmers' children remain.

Code and Data Availability Statement

This study generated no code and no quantitative dataset. Primary materials comprise audio recordings of three semi-structured interviews, the author's transcriptions and translations,

and handwritten fieldnotes. These are not publicly deposited. Informal conversations in the fieldnotes were held on the understanding described in section 3.2, and depositing them would break it. Transcripts of the three named and consenting participants may be made available to the editors on request.

Works Cited

Aizu-Wakamatsu City Agricultural Revitalization Council. (2024). *Aizu-Wakamatsu City paddy-field profitability enhancement vision for fiscal year 2024* [in Japanese]. https://www.city.aizuwakamatsu.fukushima.jp/docs/2015091700041/file_contents/vision2024.pdf

Berkes, F., Colding, J., & Folke, C. (2000). Rediscovery of traditional ecological knowledge as adaptive management. *Ecological Applications, 10*(5), 1251–1262. https://doi.org/10.1890/1051-0761(2000)010[1251:ROTEKA]2.0.CO;2.

BliegeBird, Rebecca, and EricAlden Smith. "Signaling Theory, Strategic Interaction, and Symbolic Capital." *Current Anthropology*, vol. 46, no. 2, 2005, pp. 221–48. *JSTOR*, https://doi.org/10.1086/427115. Accessed 6 Sept. 2026.

Bourdieu, P., & Passeron, J.-C. (1977). *Reproduction in education, society and culture* (R. Nice, Trans.). Sage.

e-Stat / Ministry of Agriculture, Forestry and Fisheries. (n.d.-b). *Long-run statistical series on farm labour hours* [in Japanese]. https://www.e-stat.go.jp/dbview?sid=000329029.

Egawa, A. (2005). *Support system for new entrants to agriculture from non-farming households from the viewpoint of farm management and rural life*. *PRIMAFF Annual Report 2004*, 32. https://www.maff.go.jp/primaff/e/publications/annual/attach/pdf/200510_p_ar04_11.pdf

Feldhoff, T. (2013). Shrinking communities in Japan: Community ownership of assets as a development potential for rural Japan? *Urban Design International, 18*(1), 99–109. https://doi.org/10.1057/udi.2012.26

Hobsbawm, E., & Ranger, T. (Eds.). (1983). *The invention of tradition*. Cambridge University Press.

Inoue, T., Koike, S., Yamauchi, M., & Ishikawa, Y. (2022). Exploring the impact of depopulation on a country's population geography: Lessons learned from Japan. *Population, Space and Place, 28*(5), e2543. https://doi.org/10.1002/psp.2543

Klien, S. (2020). *Urban migrants in rural Japan: Between agency and anomie in a post-growth society*. State University of New York Press. https://doi.org/10.2307/jj.18253018

Kojima, N. (2026、February 7). 振り返ってみると [Looking back]. Originally published in 会津嶺 *[Aizune]*、November 2025 issue [in Japanese]. Republished by 自然農法無の会 on note.com. https://note.com/munokai/n/n2053c5ac5a38

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press. https://doi.org/10.1017/CBO9780511815355

Matanle, P., & Sato, Y. (2010). Coming soon to a city near you! Learning to live "beyond growth" in Japan's shrinking regions. *Social Science Japan Journal, 13*(2), 187–210. https://doi.org/10.1093/ssjj/jyq013

McGreevy, S. R., Kobayashi, M., & Tanaka, K. (2019). Agrarian pathways for the next generation of Japanese farmers. *Canadian Journal of Development Studies, 40*(2), 272–290. https://doi.org/10.1080/02255189.2018.1517642

Ministry of Agriculture, Forestry and Fisheries. (n.d.-a). *Direct labour hours required for rice cultivation, by year* [Chart 5-2] [in Japanese]. https://www.maff.go.jp/j/council/seisaku/kensho/attach/pdf/5siryo-2.pdf

Ministry of Land, Infrastructure, Transport and Tourism. (n.d.). *Aizu Nōsho* [in Japanese]. https://www.mlit.go.jp/tagengo-db/common/001653708.pdf

Polanyi, M. (1962). Tacit knowing: Its bearing on some problems of philosophy. *Reviews of Modern Physics, 34*(4), 601–616. https://doi.org/10.1103/RevModPhys.34.601

Robertson, J. (1988). Furusato Japan: The culture and politics of nostalgia. *International Journal of Politics, Culture, and Society, 1*(4), 494–518. https://doi.org/10.1007/BF01390685

Thornbury, B. E. (1994). The Cultural Properties Protection Law and Japan's folk performing arts. *Asian Folklore Studies, 53*(2), 211–225. https://doi.org/10.2307/1178644

Tsutsumi, M. (2001). Succession of stem families in rural Japan: Cases in Yamanashi Prefecture. *International Journal of Japanese Sociology, 10*(1), 69–79. https://doi.org/10.1111/1475-6781.00008

Uchiyama, T., Lobley, M., Errington, A., & Yanagimura, S. (2008). Dimensions of intergenerational farm business transfers in Canada, England, the USA and Japan. *The Japanese Journal of Rural Economics, 10*, 33–48. https://doi.org/10.18480/jjre.10.33

Ueno, T. (2023, October 12).「無の会は何者か、無の会はどこへ行くのか」["Who is 無の会, and where is it going?"] [in Japanese]. 自然農法無の会, note.com. https://note.com/munokai/n/nd6d817f82053

無の会 [Mu no Kai]. (n.d.-a). これからの"農"を模索する [Exploring the future of "agriculture"] [in Japanese]. 自然農法無の会. https://munokai.com/pages/explore-farming

Retail Investors vs Institutional Investors - The Gamestop Saga By Alexander Yan

Introduction

In recent years there have been a multitude of cases where institutional investors around the world have performed the action known as 'short squeeze'. Short squeeze is the process in which the price of an asset rises sharply, causing traders who had sold short to close their positions. It occurs when a security has a significant amount of short sellers, meaning lots of investors are betting on its price falling (Mitchell 2024).

Out of these events that have happened in recent memory, such as the 2008 Porsche vs Volkswagen or the big short of Herbalife (Gaw 2025). The recurring trend seems to point towards the involvement of two or more institutional investors, with retail investors playing no significant role in the operation. However, despite this tendency, a special episode known as the short squeeze of the company 'Gamestop' occurred in 2021. From January 13, 2021, Gamestop shares saw a sudden and drastic increase in price and in return volatility. The run-up was reported to have been led by a large increase in trading by retail investors using the Robinhood Financial platform, organized via social media, in particular the *WallStreetBets* chat forum on Reddit (Malz 2021). This completely shut out the institutional investors who were holding short positions, marking a shift from the growing domination of those markets by large institutional investors (Fisch 2022).

Before delving into the event that occurred in 2021 regarding GameStop, it is important to understand the two sides that were involved in the trades. Institutional investors are entities or organisations that pool together large sums of money to invest in a variety of asset classes, including but not limited to stocks and ETF's. These include pension funds, mutual funds, insurance companies and hedge funds (Lawler 2023). Whereas a retail investor is an individual who invests their own money into stocks, bonds, mutual funds, and other similar vehicles. Unlike institutional investors, these individuals are not professional traders or investment entities but are everyday people investing on their own behalf, looking to grow their savings or prepare for future financial goals (Lawler 2023).

In August of 2019, GameStop was not a high-value company and the stock was slightly south of $4.00 per share (Salvucci 2023). Covid-19 seemed to have been a 'blessing in disguise' for the company. Despite the company being on the brink of bankruptcy due to the pandemic, at the start of 2021, something truly unprecedented yet miraculous happened. Many retail investors put their resources together and invested in the company, attempting to take advantage of the low prices. By the end of January 2021, the prices of GameStop skyrocketed and seemed to have reached an all-time high, hovering around $500 per share (Salvucci 2023). Although the share is back to around $13-15 per share in 2024 (Hood et al. 2024), the 2021 events signalled a triumph for retail investors over the institutional ones who were intent on shorting the stock to even lower prices.

Institutional vs Retail Investors

The common trend across the world points towards the fact that institutional investors make up the majority of all shares regardless of the company. This can first be seen within the US, where a study from Bloomberg Intelligence in 2020 found that only 19.5% of all investors were retail (Arora 2022). Undoubtedly, this shows the dominance that institutional

investors have in the American market. Similarly, for the UK, only 13.5% of UK shares are owned by retail investors (Sharesoc 2023). The rest of Europe does not shy away from this trend either, and according to a recent report in 2021, 72% are institutional investors, leaving 28% to be retail investors (Stastica Research Department 2024). However, according to Tim Focas, the majority of the investors are passive investors and only 3% are actively investing and trading in day-to-day stocks. When moving across and into Asia, this trend seemingly does not disappear and according to The Economics Times, only around 7.5% of investors listed on the National Stock Exchange are retail investors in India (ET Bureau 2023). However, on the surface, this dominance by institutional investors seems to have been stopped in South Korea because roughly 60% of the total stocks that are owned in South Korea are owned by retail investors (Park 2023). However, it is important to note that short selling is banned in South Korea, which is an important way in which institutional investors make money. Japan, on the other hand, does follow the trend that appears to hold true for the rest of the world. According to Statista, in 2024, only around 18% of all investors in Japan were retail investors (Statistica Research Department 2025).

On one hand, there are people who believe that retail investors should accept the fact that institutional investors make up the majority of financial markets around the world. They simply believe that retail investors do not have enough expertise, enough information and enough time when compared to the people whose entire lives revolve around the concept of investment and trading. On the other hand, there is a majority of people who believe that retail investors can do just as much if not more than global institutional investors. If the market is controlled entirely by institutional investors, then whenever things go haywire, retail investors that may have not contributed to the problem are the ones who endure the most severe consequences as they generally lack the means to regain the wealth they have just lost. In other words, when the greed of institutional players outmatches their ability to contain risks, small retail investors also stand to incur losses for no apparent reason of their own. This was accentuated during the 2008 financial crisis, when the only people who managed to ride the episode out were institutional, and not retail investors. While most financial institutions, bankers, and executives have been able to regain their wealth, the impact of these events on retail investors was devastating as forty-seven percent of all American Households owned some stocks or bonds (Black 2009, 2).

Deconstructing the Incident

The motives behind the events at Gamestop still remain unclear. Throughout the history of Wall Street the stock market has always been largely controlled by institutional investors such as investment banks, commercial banks, hedge funds, mutual funds, asset management firms, insurance companies, broker-dealers, currency and commodity traders, financial institutions, and so on (Picardo and Anderson 2023).The events of Gamestop however was this process of retail investors coming together to buy into the stock, completely reversing the plans of the hedge funds and money managers that had been trying to short sell (Dziak 2024). In a competitive society fueled by money, it is easy for one to assume that these investors simply wanted to make profits on their own behalf. Due to the lack of expertise that the retail investors possessed, short-term thinking heavily occupied their mind and therefore as soon as the price of Gamestock reached an all-time high, the

majority of the investors sold the stocks that they had, which accounted for large amounts of profits. This could have been the goal the entire time: to make money.

Yet at the same time, all these motives are retrospective. Despite the fact that making profits seems like the ultimate goal for these investors, the primary goal can also be read as the idea of revolutionising society and the stock market as a whole. The driving motive behind this event was not so much about the retail investors making profits. Rather it was to show how through coordinated efforts they were able to challenge large institutional players (BOOKMAP 2024). It was the first major event in the recent history of Wall Street when there was a collective unity that co-existed between all the retail investors. These investors were able to push the desire for their own needs to one side and collectively work towards overpowering the institutional investors who were always seen as the ones determining the fate of financial markets.

The ideology of capitalism has always been ingrained into the economy in the US. American capitalism has flourished by providing incentives for innovation and bold risk-taking by entrepreneurs and investors (Manyika, Pinkus, and Tuin 2020). In fact, many analysts view the United States as the peak capitalist economy—the economy whose institutions other countries should emulate (Freeman 2003, 19). Usually what happens in the market is that power is typically concentrated in the hands of institutional investors, hedge funds, private equity firms and so on. As these entities have massive experience in the field, they are easily able to exert a heavy influence on the market. Retail investors, on the other hand, often have limited power and face barriers to entry. However, the GameStop saga directly contradicted the power dynamic, even if fleetingly. This event demonstrated that under certain circumstances, retail investors can have a larger impact on the market and challenge the abilities of institutional investors. This event proved to be pivotal in the scheme of the market as it directly challenged the norms of capitalism that had existed for an extended period. Furthermore, the Gamestop saga also demonstrated critically the fragility of the financial market. The value of Gamestop deviated from its original value, which was driven by the trading of retail investors. This highlights how the market prices can detach from underlying economic realities, illustrating the irrationality and unpredictability the stock market possesses. This heavily contradicts the normal capitalistic market in which the market is expected to not only be efficient but also very rational. The stock prices are usually a great reflection of the fundamental value of that company. In the long term, a Wall Street adage says that "the fundamentals" always win. That means a stock's price eventually settles where it should, based in part on how much profit a company is producing (Torstar Syndication Services 2021). Through data analysis, the value is usually brought to bear on financial performance, future prospects, as well as entire market conditions. However, the sudden disruption of this proved to be a crucial moment.

Gamestop also teaches us a vital lesson about capitalism and how the only way to change capitalism is through capitalism. There is no external way of manipulating the financial market, whether that is demonstrations, protests or legal changes. The institutional investors live and breathe the stock market, so only through playing the ‘game’ better than they do could one actually succeed over them. To put it another way, we can see ourselves living in a late, mature phase of hyper-digitised capitalism–Gamestop capitalism, if you will–where the limits to social action are circumscribed within the existing order of things.

Regarding the matter, I believe that the problem that exists is not the fact that people can fight within the system, which indeed is quite important within the market. Rather, the main problem that exists is the simple fact that this is the only way in which the retail investors are able to succeed. There is simply no other way retail investors can fight within the current system, as it is completely foreclosed by the institutional investors that have the power and, more importantly, want to maintain the status quo. The fact of the matter is that there is no other method for retail investors to succeed. Otherwise, we would have seen more individual involvement in financial markets around the world. Moreover, all the feasible methods have proved that they can be just as easily shut down as they are thought of. The ways that might work involve major changes to the stock market as a whole and potentially the government, which would, in a sense, halt the entire economy.

However, there are various other reasons that could be exposed for investment in Gamestop. Firstly, the stock market is widely regarded as one of the most profitable and best forms of investment, and it has gotten to the point where it has almost become a necessity for moneyed people in society to invest in the market. If one does not invest in the stock market, they have a chance of losing the value of their money. Additionally, investing in the stock market can allow a person to retire early; it is an alternative place where one can park their money, thus allowing them to receive warranted and passive income. Moreover, this investment has proved to be the best form as it has outweighed other forms of investments such as real estate. Mutual funds, ETFs, and retirement accounts offer professional management, but their returns are always second to what the stock market offers (Hartford Funds Management Group 2025). Though there are other investment strategies such as whiskey, real estate and artwork, amongst other alternative investments, the barrier of entry into these methods is comparatively higher. The amount of money required to begin the investment is far out of the range of many people in society. Whiskey and artwork on the other hand have too little of a supply to be a stand-out form of investment. If there are not enough people to do it, the market simply is not large enough and therefore, not enough money exists in it. Lastly, investment banks and private equity firms may seem like a great form of investment due to the return being quite high. However, trusting and allowing other people with your assets and then losing it all might be a recurring fear for many with only a bare minimum income to invest. The problem with investment banks and these firms is the fact that despite profits being relatively lucrative, the risk is too high for the reward that is given back to the person. There simply is not enough in return when compared to the sheer amount of risk that one puts in. When comparing the stock market investment to the rest of these, it proves to be the most widespread and easily the most reliable. This could be another reason why the events of Gamestop happened in the first place: a tentative materialisation of the desire to participate in the stock market.

Since we have already established that the place in the stock market is not a leisure activity in this day in age but a necessity to beat inflation, it begs the question of whether or not GameStop promises and enables this inclusion in the stock market. Before we ask this question, we must wonder, what true meaning and purpose did GameStop really have? It showcased the ability of the social network and how people can come together via social media, such as Reddit, to achieve massive feats. However, it also showed with great clarity that social media itself is nowhere near enough for a social movement to be successful.

Consider the entrenched capital of successful large investment banking and private equity firms. They all have a concrete and efficient organisation that keeps their company intact. There is a hierarchy which allows these companies to understand their management and have deserving people make decisive decisions for the entire company. Additionally, these companies all have long-term goals, which are accepted by the people working in these firms. When compared to GameStop, it is evident that this was not the case. Many investors started to quit and sell their stock halfway through the process as they were thinking about the profits that they could generate, which left many people disappointed whose goal was to demonstrate a lesson towards the instructional investors and wider society. These social movements, if we were to see GameStop as one, do not think about the long term, and this showcases a plaguing obstacle to their success. On one hand, large firms and companies have long-term plans, with contingency back-ups, to keep themselves in the market with their diverse and intricate planning. On the other hand, most social movements do not have any clarity at all about what their plan is for the future, and this leaves them in an extremely vulnerable position, which leads to many individuals starting to doubt the process, ultimately making the entire process fizzle out of existence. What this has taught us is that social media platforms are a tremendous place to bring people together and communicate. However, for a social movement to be successful, something needs to happen beyond social media, which can be seen as the first step in a series of stepping stones. Organisation, leadership, and goal-setting need to be well thought-out by everyone participating and shared amongst the movement. Social media simply cannot provide this and therefore an alternative must be found.

However, the failure of the GameStop saga can also be seen as an act of active manipulation. It started off as just a couple of people coming together with Reddit, with the initial goal being to rebel without a cause to demonstrate to institutional investors that even the retail investors can exert a heavy influence on the stock market. However, this threat seemed to grow exponentially and it was not long before there were too many threads on Reddit overwhelming the original idea until it fizzled away, and people did not know what was happening. People joined the movement, and some may argue that some people joined with mischievous intent. Andrew Ross Sorkin said on live television that the entire saga was a Ponzi Scheme (CNBC Televsion, n.d.). The stock market is a place of numbers that entail fine margins. If people can successfully manipulate the numbers, then it can drive a massive movement straight into the ground, and this is what ended up happening. The movement became too widespread and people with the mala fide intentions joined the movement and ultimately led to its failure. However, this risk of ‘active manipulation’ must have been accepted by the people knowing all too well the risk that comes with collaborating with strangers on social media. Despite knowing this risk, people still acted on it, which shows how this failure cannot be blamed on this active manipulation, as the risk of this happening was completely warranted. Furthermore, given the number of people that joined, this overwhelming movement can only be read as a result of the success it originally had, which led to many people jumping on the bandwagon to make profits. Overall, we can point to a variety of different reasons why GameStop failed. It was not simply one factor that led to its downfall, but rather an accumulation of many different facets.

Moreover, GameStop demonstrated the importance of social capital and economic

capital. A social movement always requires economic capital, as money is required for any successful movement (Rogers 2023). In a majority of cases, this form of capital comes from donations which people give their money to a movement. But a popular fear is that this money is rarely used to further the movement and is never to be seen again. This can deter a lot of people from wanting to join movements, as it means losing their money. What made GameStop such a large and successful movement was the sheer fact that the people who put the money into the movement received more money in return. Social capital is also an absolute necessity in any social movement. It always requires a large number of people collectively working together towards a common aim. In the beginning, this was apparent when retail investors who had never really been considered a threat altered the entire market. Through successful collaboration and a common goal, the idea of capturing social capital was emphasised, and people were able to make a noticeable change. However, the vital nature of social capital was even more evident when there seemed to be fractures within the unity of these retail investors and the goals that had once been united started to shift in various directions. It proved that if there is no common goal, organisation or leadership that can allow people to work together efficiently, the efforts put into the work of organising will prove to be futile. Without collectivisation, social movements will never be able to experience sustained success.

Overall, Gamestop can be seen by many as a great success as it was able to prove to the world the strength of retail investors and allow them to make some profits. However, it did not produce any sustained leadership or organisation. Therefore, the very reasons it was possible in the first place were likely the reasons owing to which it eventually failed. The movement consolidated a degree of social and economic capital, and it was very easy for people to join and participate. However, the very nature of the movement being easy to join also meant that quitting could be just as easy. Combining this with the lack of organisation and leadership could have been the Achilles' heel of a great campaign.

What to make of Gamestop?

There have been similar movements to that of the Gamestop Saga in recent years, mostly in the political sphere: for instance, the Arab Spring and the Occupy Wall Street movement. The Arab Spring movement happened in 2010, mostly in the Middle Eas,t where people called for political reform due to unrest and a want for greater civil liberties. This revolution started in Tunisia and Egypt, then spread to many other Arab countries like Libya and Yemen (da Vinci and Ali's 2024). There were many similarities that existed between the Arab Spring and the Gamestop Saga, however, the most glaring similarity lies in the use of social media. Both events used social media to encourage people to join their movement. In the case of the Arab Spring, different types of social media were used to facilitate communication and organise revolts. It was also used by the government to encourage civilians to join the protest as well as raise global awareness. Additionally, neither movement had a true leader and a sense of hierarchy, which led to decreased organisation and collaboration.

The Occupy Wall Street movement occurred in 2011 in New York’s famous financial district – Wall Street. Here, a group of young activists decided to protest the US economy. They used the popular phrase, ‘We are the 99 percent ’, referring to the difference in wealth

between the richest percentile of U.S. citizens and the rest of the population (Volle 2025). The reason for the failure of the Occupy Wall Street Movement was similar to the reasons for the failure of GameStop. There was no leadership and sense of direction in the movement. This made it difficult for the movement to have a unified message and a common goal, which failed to bring about concrete change.

Both these movements showed how the major flaws that the Gamestop movement possessed had already been actualised in recent times, making people question whether or not Gamestop could be avoided entirely as their approach was very similar to past events that failed, implying in a manner that they had not learnt anything from it. However, it is naive to overlook that although no real change was brought, there was a sense of reform or hope that was being sparked. It indicated that change was indeed possible, even though unlikely. In criticising this episode excessively, it is also easy for us to forget the fact that the reason for early success within these movements stemmed from a sense of unity, not hierarchy, along with great communication through the extensive use of social media. Despite this, the end was inevitable. It was a failure. Although the sense of no hierarchy was appealing in the beginning. It led to the movements lacking a barrier to entry. Anybody could join easily to help the movement, but it was also quite easy for people to leave. There was no leader to truly unite people in the long term and to set goals for everyone. Could Nike or Apple survive without a CEO? Or does a successful movement require a leader?

After comparing these two events to Gamestop, it is personally a massive shame to see such a big movement plummet to the floor due to the poor execution by the people involved. After many events like this have happened in the past, across different fields, it is apparent that the mistakes that occurred in Gamestop could have been entirely avoided. If Gamestop could have developed into something extremely successful, the profits the retail investors would have gained would have been historic. This could have also paved the path for future investors looking to effect similar change. Instead, the legacy of the movement is but a faint memory, except for the reassurance that if something like this could have happened once, it can happen again. That being said, it is plausible that this situation may be politically galvanised at a different point in the future. It was the first time that everybody, not just the financial world, had witnessed something where the retail investors trumped institutional investors, signifying the power of collectivisation and the possibility of reenacting a similar resurgence in the future, hopefully with clearer goals and a better organisational structure. Yet, it is valid to wonder if retail investors may want to achieve this again, and the answer may be, not quite. Firstly, given the unexpectedness of the event, it is hard to believe that an event like this can ever occur again in the face of greater vigilance by institutional players. The main reason lies in the fact that Gamestop heavily relied on the COVID-19 pandemic. A big reason for the massive initial success was correlated to the pandemic, with more people turning to online gaming. More involvement in social media also generated unprecedented levels of traction. The unique set of circumstances presented by the pandemic is hard to replicate. This indicates how this event is extremely unlikely to happen again, considering the bleak chances of recreating the conditions that had enabled it in the first place. In fact, even if the circumstances were to be opportune, it is still futile to believe that an event like this can repeat. This brings me to my second point, which is that despite this being a revolutionary event in some eyes, the truth remains that the institutional

investors escaped this fiasco with barely a scratch. Despite having lost about $US20 billion ($26 billion) on their Gamestop positions, there were able to recover their wealth in record time (Nine Entertainment Company 2021) . Even if institutional investors such as banks, hedge funds and other financial companies incur losses more regularly, their wealth pool is large enough to cause a speedy recovery, not to mention the hedging strategies which they anyway employ by default.

Gamestop was an unparalleled opportunity, one for retail investors to change the course of the financial world's contemporary history. Whereas, it was another tough day at the job for institutional investors who recouped their losses in a matter of days, with the idea of this event long at the back of their minds. It is a shame to see such a big movement plummet due to the poor execution by the people involved–indeed, the people involved had no name or face to them. When comparing the Gamestop Saga to similar events such as the Arab Spring and the Occupy Wall Street movement, it is apparent that Gamestop made mistakes that they could have pre-empted. If this movement could be prolonged, it could have easily been turned into an extremely successful movement—not the least because of the extremely high profits at stake.. When looking at it from a wider lens, we can deduce that Gamestop was a good sign for the financial world as it provided everyone with something that hadn't quite been seen before. It was the first time that everybody had seen the retail investors beat institutional investors, and therefore shone a new light on something that was still much of a grey area prior to this episode. To put it another way, it signalled a safety valve to dissipate pent-up institutional capital whenever its concentration becomes too unbearable for the overall economic health of the market. However, it is hard to believe that an event like this can occur again. The main reason, at the cost of repetition, lies in the fact that Gamestop heavily relied on the world pandemic. Due to the pandemic, not only did more people turn to online gaming, but the social media influence surrounding the event also increased. This unique set of circumstances is hard to replicate. Additionally, if by some miracle something like this were to happen again, it is futile to believe that it will have much of a knock on institutional investors. Institutional investors, by nature, can tolerate more losses than retail investors. Institutional investors understand that they can recover their losses through their extremely diverse portfolios. Of course, the stock of Gamestop was a massive opportunity for both parties. However, for the retail investors, it might have been a once-in-a-lifetime chance to make extreme profits, whereas for the institutions, it was just another day at the office.

Conclusion

There has always been this secret tension that has existed in the background of the financial sector that is not talked about enough. This paper specifically unmasked the tension that exists between institutional investors and retail investors, looking at it through the lens of the Gamestop saga that occurred in 2021. This paper aimed to first give a little background about the event and how it truly unfolded, exploring the surprising nature of the capricious price rise.

In conclusion, I would like to emphasise that this paper is not a definitive statement of everything that occurred during the Gamestop saga. I do recognise that there are certain aspects in which this paper lacks. While many research papers may choose to focus on the depth of the situations, this paper prefers more breadth, focusing more on the problems

within the financial sector. When writing this paper, I had a choice of focusing solely on Gamestop to the exclusion of its larger social implications. However, I chose to use this event as an example to explore the totality of the financial market and the tension between retail and institutional investors. Additionally, this paper focuses not on the quantitative side of finance, such as the day-to-day market standings, or even the financial models that caused the Gamestop stock to fluctuate. Instead, this paper engages with the qualitative aspects, taking data not as mere numbers, but as a narrative that helps us interpret what investors are trying to do. Instead of specificity, this paper seeks to grasp the entire situation. It focuses on the meanings and motivations behind actions and not specific numbers. After all, there is no way the financial market can be encapsulated into only one event or one stock. It is a world of its own.

Works Cited

Arora, Krishan. 2022. “The Rise Of The Retail Investor.” https://www.forbes.com/councils/forbesagencycouncil/2022/11/04/the-rise-of-the-retail-investor/.

Black, Barbara. 2009. “PROTECTING THE RETAIL INVESTOR IN AN AGE OF FINANCIAL UNCERTAINTY.” (July).

BOOKMAP. 2024. “What Really Happened with the GameStop Short Squeeze? A Deep Dive into the Market Phenomenon.” https://bookmap.com/blog/what-really-happened-with-the-gamestop-short-squeeze-a-deep-dive-into-the-market-phenomenon.

CNBC Televsion, dir. n.d. *Some people see GameStop as a stick-it-to-the-man opportunity: Sorkin*. https://www.youtube.com/watch?v=kSqnwa3EQWA. da Vinci, Leonardo, and Ben Ali's. 2024. “Arab Spring | History, Revolution, Causes, Effects, & Facts.” Britannica. https://www.britannica.com/event/Arab-Spring.

Dziak, Mark. 2024. “GameStop short squeeze.” EBSCO. https://www.ebsco.com/research-starters/business-and-management/gamestop-short-squeeze.

ET Bureau. 2023. https://economictimes.indiatimes.com/opinion/et-editorial/retail-investors-live-long-an d-prosper/articleshow/102842314.cms.

Fisch, Jill. 2022. “GameStop and the Reemergence of the Retail Investor.” Freeman, Richard. 2003. “Changing the Guard: The Rise of the United States to Peak Capitalist Economy.” (January), 19. https://www.nber.org/system/files/chapters/c10301/c10301.pdf.

Gaw, Kathryn. 2025. “Biggest short squeezes: what were the 5 worst squeezes for short sellers?” IG. https://www.ig.com/en/trading-strategies/what-were-the-biggest-short-squeezes-in-hi story--210507. Hartford Funds Management Group. 2025.

Hood, Cooper, Nathan Graham, Brennan Klein, Adam Bentz, Ryan Northrup, Colin McCormick, and Matthew Rudoy. 2024. “How High GameStop's Stock Price Got & What It Costs Now.” Screen Rant. https://screenrant.com/dumb-money-gamestop-stock-price-high-cost-now/.

Lawler, Jasper. 2023. “Retail and Institutional Investors: Definition, Comparison, Difference.” Trading 212. https://www.trading212.com/learn/investing-101/retail-and-institutional-investors

Malz, Allan M. 2021. “The GameStop Episode: What Happened and What Does It Mean?” Cato Institute. https://www.cato.org/cato-journal/fall-2021/gamestop-episode-what-happened-what-d oes-it-mean.

Manyika, James, Gary Pinkus, and Monique Tuin. 2020. “Rethinking the future of capitalism in America.” McKinsey. https://www.mckinsey.com/featured-insights/long-term-capitalism/rethinking-the-fut

ur e-of-american-capitalism.
Mitchell, Cory. 2024. “Short Squeeze: Definition, Causes, and Examples.” Investopedia. https://www.investopedia.com/terms/s/shortsqueeze.asp.
Nine Entertainment Company. 2021. “Defending the shorts: GameStop saga goes beyond good and evil."” (Februrary). https://link.gale.com/apps/doc/A650436968/STND?u=hkharrow&sid=bookmark-STN D&xid=5b24a567.
Park, Anna J. 2023. “Retail investors take up 64% of Korean stock market.” The Korea Times. https://www.koreatimes.co.kr/www/biz/2025/02/602_363158.html.
Picardo, Elvis, and Somer Anderson. 2023. “Why Wall Street Is a Key Player in the World's Economy.” Investopedia. https://www.investopedia.com/articles/investing/100814/wall-streets-enduring-impact economy.asp.
Rogers, Cathy. 2023. “The case for funding social movements.” https://www.alliancemagazine.org/blog/the-case-for-funding-social-movements/.
Salvucci, Jeremy. 2023. “An in-depth timeline of the GameStop short squeeze saga.” TheStreet. https://www.thestreet.com/investing/stocks/a-timeline-of-the-gamestop-short-squeeze
Sharesoc. 2023. “UK STOCK MARKET STATISTICS SOME BACKGROUND INFORMATION.” https://www.sharesoc.org/investor-academy/advanced-topics/uk-stock -market-statisti cs/. Stastica Research Department. 2024. https://www.statista.com/statistics/368484/europe-market-share-institu tional-retail-clie nts/#:~:text=As%20of%202021%2C%20the%20share,management% 20stood%20at %2072%20percent.
Statistica Research Department. 2025. “Distribution of individual investors in Japan as of July 2024, by value of securities owned.” https://www.statista.com/statistics/1370156/japan-private-investors-breakdown-by-val ue-of-securities-held/.
Torstar Syndication Services. 2021. “"GameStop saga is over, but effect on Wall Street could be longer-lasting.” February 21st, 2021. https://go.gale.com/ps/retrieve.do?tabID=T004&resultListType=RESULT_LIST&seahResultsType=SingleTab&retrievalId=118b180c-7ec3-46cc-b47f-4939dee48a96&hitount=556&searchType=BasicSearchForm&currentPosition=9&docId=GALE%7CA65 2522033&docType=Article&sort=Rel.
Volle, Adam. 2025. “Occupy Wall Street | 2011, Definition, Movement, & Purpose.” Britannicahttps://www.britannica.com/topic/Occupy-Wall-Street.