

AI in Financial Applications: Utopia and Reality

Yi Zhou

Carey Business School, Johns Hopkins University, Washington, D.C., USA

Email: yzhou286@jh.edu

Manuscript received April 30, 2026; accepted May 30, 2026; published September 2, 2026.

Abstract—Generative AI, in particular, is increasingly reshaping financial markets by transforming information processing, altering data-generating processes, and redefining decision-making structures. However, its implications remain inherently dual, combining efficiency gains with emerging risks. This paper examines the application of AI in finance through a “utopia versus reality” framework, focusing on the gap between its theoretical potential and practical outcomes. Using a literature review approach, it synthesizes existing research on AI’s role in information infrastructure, financial data transformation, and market behavior. The findings suggest that generative AI enhances information efficiency and expands the scope of usable data, enabling new forms of analysis and decision-making. At the same time, it introduces structural challenges, including diminished signal reliability, instability in data-generating processes, shifts in labor demand, and rising concerns over trust stemming from bias, opacity, and limited verifiability. Besides, the emergence of self-evolving and agentic AI transforms these risks from isolated model errors into dynamic and systemic concerns. Overall, while AI enhances efficiency and expands the scope of financial analysis, its inherent limitations and risks highlight a persistent gap between its theoretical potential and practical outcomes.

Keywords—artificial intelligence, generative ai, financial markets

I. INTRODUCTION

The rapid advancement of artificial intelligence is reshaping industries across the board, and finance is no exception. To understand the impact of AI applications in finance, it is necessary to first outline the broader trajectory of AI development.

The conceptual origins of AI can be traced back to Turing (1950), with the term formally coined at the Dartmouth Conference in 1956 (McCarthy, Minsky, Rochester, & Shannon, 2006). Early AI systems (1950s–1980s) centered on rule-based approaches, relying on logical reasoning, expert systems, and manually encoded knowledge to solve problems, but lacked the ability to learn from data. A turning point came in 1986, when Rumelhart *et al.* (1986) extended backpropagation to multi-layer neural networks, enabling machines to learn from data iteratively. From the 1990s onward, advances in computational power and the rise of internet-scale data shifted AI toward data-driven paradigms in which statistical learning models automatically extract patterns from large datasets, marking a transition from explicit rule encoding to empirical learning.

Subsequently, as foundational architectures including Long Short-Term Memory networks (LSTMs), Deep Learning frameworks, Generative Adversarial Networks (GANs), and the Transformer were successively proposed and refined, AI entered the era of Large Language Models (LLMs) and generative AI systems. The 2022 release of ChatGPT marked generative AI’s mainstream arrival. The

current frontier is moving toward multimodal agents, shifting from passive question-answering toward autonomous task planning.

This paper examines AI applications in finance from a dual perspective of “utopia” and “reality,” focusing on the gap between its theoretical potential and practical implementation. Using a literature review approach, it provides a balanced evaluation of the broader benefits, impacts, and risks of AI in financial systems, while briefly noting emerging concerns related to current developments. This insight becomes increasingly vital as generative and autonomous AI continue to reshape financial markets. As AI evolves toward generative and autonomous systems, understanding this gap is essential for financial practitioners, researchers, and regulators.

II. AI APPLICATION IN FINANCE: UTOPIA

A. Generative AI as an Information Infrastructure

Generative AI is increasingly functioning as an information infrastructure within financial markets, fundamentally reshaping how information is produced, processed, and distributed.

By lowering the cost of information acquisition and processing, generative AI enhances the capacity of financial analysts and investors to collect, generate, analyze, and disseminate information, thereby improving both the quantity and quality of information output. Evidence from the temporary ban of ChatGPT in Italy shows that restricting access to generative AI led to a decline in analyst forecast volume and accuracy, accompanied by a widening of bid–ask spreads; these effects reversed once access was restored (Bertomeu, Lin, Liu, & Ni, 2025). This finding suggests that generative AI improves information processing efficiency and reduces information asymmetry, facilitating more efficient price discovery.

Beyond processing efficiency, generative AI also increases research and decision-making productivity. Generative AI enables the rapid integration of large-scale information, while large language models support tasks such as hypothesis generation and research design (Eisfeldt, & Schubert, 2025), thereby rendering previously infeasible analyses tractable.

B. AI and the Transformation of Financial Data

Following generative AI’s emergence as an information infrastructure in financial markets, it’s also transforming the nature of financial data itself.

First, AI expands the range of available data. Traditional financial research relied predominantly on structured numerical data, such as prices, yields, trading volumes, and accounting metrics, valued for their standardization and tractability. With the application of LLMs, unstructured information such as earnings call transcripts, analyst reports, and regulatory disclosures has become increasingly quantifiable

and analyzable. As a result, financial research is shifting from low-dimensional “small model” frameworks to high-dimensional, information-driven paradigms (Hong & Wang, 2024).

More importantly, generative AI alters the data-generating process in financial markets. As AI becomes embedded in information extraction, forecasting, and trade execution, market participants are no longer passive observers of data but active participants in its generation. In this sense, financial data evolves from an exogenous input into an endogenous outcome shaped by the interaction between AI systems and market behavior. Empirical evidence supports this mechanism: firms with high AI exposure earned significantly higher market valuations following the release of ChatGPT (Eisfeldt, Schubert, & Zhang, (2023), and hedge funds adopting generative AI earned higher abnormal returns attributable primarily to enhanced information processing (Sheng, Sun, Yang, & Zhang, 2025). Since asset prices themselves represent a critical component of financial data, AI has begun to reshape the data generation mechanism from the inside out.

III. AI APPLICATION IN FINANCE: REALITY

From a reality perspective, while AI’s application in finance has enhanced efficiency and expanded data sources, it has simultaneously introduced new structural concerns.

A. *Structural Risk of AI in Finance*

First, generative AI can generate plausible but inaccurate or unverifiable content, creating risks of hallucination. In financial markets, such erroneous outputs may be misinterpreted as credible signals, thereby distorting expectations and asset prices. Existing research suggests that generative AI reduces the cost of producing high-quality signals, paradoxically weakening their informativeness and increasing screening errors (Cowgill, Hernandez-Lagos, & Wright, 2026). Experimental evidence further indicates that AI may produce overconfident yet flawed recommendations, leading users to decisions that worsen performance (Otis, Clarke, Delecourt, Holtz, & Koning, 2024). Therefore, AI acts not merely as an information-processing tool but also as a source of noise that amplifies mispricing risks.

Second, generative AI destabilizes the data environment underlying financial research. As AI models are continuously updated, identical inputs may produce different outputs across time, challenging reproducibility. More critically, AI introduces an endogenous data-generating process: models are trained on existing data while simultaneously generating new data that feeds back into the system (Hong & Wang, (2024). This self-reinforcing loop implies that financial data is no longer a passive record of economic activity but rather a product of technology–market interaction, thereby invalidating the assumption of a stable data environment.

Finally, generative AI reshapes the role of human capability in research and decision-making. While AI lowers the cost of accessing sophisticated tools, it also increases the risk of misuse when users lack sufficient understanding. Evidence suggests that AI adoption generates heterogeneous effects: high-ability users benefit significantly, whereas low-ability users may experience performance deterioration (Otis *et al.*, 2024). Rather than eliminating skill differences,

generative AI tends to amplify existing advantages and disadvantages, magnifying performance gaps across individuals.

B. *AI and Structural Change in the Financial Labor Market*

The impact of generative AI on the financial labor market is uneven and task-dependent, reflecting differences in skill requirements and in the ability to effectively utilize AI tools (Acemoglu, Autor, Hazell, & Restrepo, 2022). Rather than causing widespread job displacement, AI reshapes both task allocation and the nature of human capabilities.

Generative AI reduces the importance of routine, information-processing tasks while increasing the value of higher-order cognitive skills. Traditionally, financial roles emphasized data collection, organization, and preliminary analysis, those tasks at which LLMs perform particularly well and are therefore more easily replaced. In contrast, skills requiring contextual understanding and complex decision-making remain less easily replaceable by current AI systems.

At the same time, AI alters career development pathways. By lowering entry thresholds for foundational tasks, AI acts as a substitute for training-intensive work traditionally performed by junior employees, weakening the learning-by-doing career development model. Evidence suggests that lower-wage roles are more exposed to AI-driven productivity shocks, while higher-skilled roles remain relatively insulated (Eisfeldt & Schubert, 2025).

This transformation aligns with broader labor economics evidence that technological progress boosts demand for higher-order cognitive skills and social competencies such as decision-making and communication (Deming, 2017). Consequently, generative AI does not eliminate financial labor demand but reshapes the skill premium by increasing the value of human-AI complementarity. In financial markets, the core competency is shifting from processing information to effectively interpreting and leveraging AI-generated outputs.

C. *The Trust Issues*

Beyond structural risks and labor market disruptions, the most consequential challenges of AI in finance arise from its effects on trust.

First, the challenge posed by AI in financial markets is shifting from information generation to information credibility. Generative AI outputs derive from learned statistical patterns rather than verificatory reasoning grounded in real-world causal relationships (Hong & Wang, 2024), which limits their inherent verifiability. This weakens the epistemic foundation upon which financial decisions rely. Moreover, recent research shows that LLMs can simulate human behavior for large-scale data generation (Cowgill, Hernandez-Lagos, & Wright, (2026).

This practice risks undermining the scientific integrity of empirical research and eroding trust in research findings, as synthetically generated behavioral data cannot be independently verified against real human responses.

Besides, AI systems are inherently susceptible to bias embedded in training data, and Machine learning systems can replicate and reinforce existing patterns of discrimination, leading to biased outcomes in applications such as credit

allocation and investment decisions (Ferrara, Sellitto, Ferrucci, Palomba, & De Lucia, 2024). As AI becomes increasingly embedded in financial systems, these biases may scale, affecting not only individual decisions but also broader patterns of resource allocation.

Additionally, financial data serves both as a privacy asset and an information asset. Generative AI relies on large-scale data for training and inference; therefore, it raises the risk of data leakage. Large language models may reconstruct or indirectly expose sensitive information embedded in training data, potentially revealing proprietary data or investment strategies. This creates concerns not only about data privacy but also about informational asymmetry and market fairness. As a result, the use of AI in finance is no longer purely a technical issue, but a central concern of Responsible AI, involving data governance, information security, and institutional trust.

D. Emerging Issues: Self-Evolving AI and Agentic AI

As AI advances, current research is shifting from static large language models toward more dynamic systems, including self-evolving AI and Agentic AI, and meanwhile, related concerns arise. Unlike traditional models with fixed parameters, self-evolving AI continuously updates through feedback during deployment, enabling adaptive decision-making in complex environments. However, this adaptability introduces the risk of objective drift, where models gradually deviate from their intended goals through iterative optimization, potentially leading to unpredictable or harmful outcomes. Existing research shows that self-evolving agents may abandon alignment constraints and converge toward unaligned strategies during interaction (Han, Xiong, Liu *et al.*, 2026). These risks are further compounded by the black-box nature of AI systems, which makes their internal logic difficult to observe and evaluate. In financial markets, where even minor changes in model parameters can be rapidly amplified by trading behavior, such dynamics substantially increase the potential for systemic risk.

At the same time, the emergence of AI agents marks a shift of AI from an analytical tool into an autonomous decision-maker. While this enhances efficiency, it also introduces significant risks related to control and alignment, as autonomous systems may operate beyond effective human oversight. Evidence suggests that under outcome-driven optimization, AI agents may weaken ethical and safety constraints (Li, Fung, Weiss *et al.*, 2026). Real-world incidents of agents overriding user instructions and executing unintended actions further highlight the fragility of control in autonomous AI systems. Taken together, self-evolving and agentic AI transform AI risk from isolated model errors into dynamic and systemic failures. This shift underscores the urgent need for governance frameworks capable of ensuring long-term alignment and stability.

IV. CONCLUSION

This paper examines the dual trajectory of AI's integration into finance, spanning an aspirational vision of generative AI as a transformative information infrastructure and a more sobering reality defined by structural, institutional, and emerging risks. On the utopian side, generative AI has

demonstrably enhanced information processing capacity, expanded the scope of usable financial data, and shifted research paradigms toward high-dimensional, information-driven approaches. On the reality side, however, the same technologies introduce hallucination risks, endogenous data generation mechanisms, labor market disruption, trust deficits rooted in bias and opacity, and the emerging governance challenges posed by self-evolving AI and autonomous agents.

As a literature review, this paper is necessarily limited by the scope and recency of existing empirical studies, many of which remain preliminary or context-specific. The rapid pace of AI development also means that findings may quickly become outdated. Future research will benefit from longitudinal studies tracking AI's market-level effects over time, as well as comparative institutional analyses of how different regulatory environments shape the responsible deployment of AI in finance.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

REFERENCES

- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. 2022. Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics*, 40(S1): 293–340.
- Bertomeu, J., Lin, Y., Liu, Y., & Ni, Z. 2025. The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting & Economics*, 80(1): Article 101782.
- Cowgill, B., Hernandez-Lagos, P., & Wright, N. 2026. *Does AI Cheapen Talk? Theory and Evidence from Global Entrepreneurship and Hiring*. CESifo Working Paper No. 12508.
- Deming, D. J. 2017. The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4): 1593–1640.
- Eisfeldt, A. L., & Schubert, G. 2025. Generative AI and Finance. *Annual Review of Financial Economics*, 17(1): 363–393.
- Eisfeldt, A. L., Schubert, G., & Zhang, M. B. 2023. *Generative AI and Firm Values*. NBER Working Paper No. w31222.
- Ferrara, C., Sellitto, G., Ferrucci, F., Palomba, F., & De Lucia, A. 2024. Fairness-aware machine learning engineering: how far are we? *Empirical Software Engineering*, 29(1): 9.
- Han, S., Xiong, K., Liu, J., *et al.* 2026. *Alignment tipping process: How self-evolution pushes LLM agents off the rails*. arXiv Preprint, arXiv: 2510.04860.
- Hong, Y., & Wang, S. 2024. How Will ChatGPT and Large Language Models Transform the Research Paradigm in Economics? *China Journal of Econometrics*, 4(1): 1–25.
- Li, M. Q., Fung, B. C. M., Weiss, M., *et al.* 2026. *A benchmark for evaluating outcome-driven constraint violations in autonomous AI agents*. arXiv Preprint, arXiv: 2512.20798.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. 2006. A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Magazine*, 27(4): 12–14.
- Otis, N., Clarke, R., Delecourt, S., Holtz, D., & Koning, R. 2024. The Uneven Impact of Generative AI on Entrepreneurial Performance. Available at SSRN: <https://ssrn.com/abstract=4671369>
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986. Learning representations by back-propagating errors. *Nature*, 323(6088): 533–536.
- Sheng, J., Sun, Z., Yang, B., & Zhang, A. L. 2025. Generative AI and Asset Management. *Review of Financial Studies*, forthcoming.
- Turing, A. M. 1950. Computing machinery and intelligence. *Mind*, 59(236): 433–460.

Copyright © 2026 by the authors. This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted

use, distribution, and reproduction in any medium, provided the original work is properly cited ([CC BY 4.0](#)).