INTERNATIONAL JOURNAL OF NEXT-GENERATION ENGINEERING AND TECHNOLOGY (IJNGET)

eISSN: 3087-4319

Volume. 03, Issue. 02, pp. 41-45, February 2026"

The Algorithmic Frontier of Financial Intermediation: A Comprehensive Analysis of Agentic AI, Large Language Models, And Blockchain Integration in Modern Fintech Ecosystems

Joshua Hoffman

Department of Financial Engineering and Computational Intelligence, London School of Economics and Political Science, United Kingdom

Article Received: 05/01/2026, Article Revised: 25/01/2026, Article Accepted: 12/02/2026, Article Published: 28/02/2026

© 2026 Authors retain the copyright of their manuscripts, and all Open Access articles are disseminated under the terms of the Creative Commons Attribution License 4.0 (CC-BY), which licenses unrestricted use, distribution, and reproduction in any medium, provided that the original work is appropriately cited.

ABSTRACT

The global financial landscape is undergoing a radical transformation characterized by the convergence of high-capacity computational intelligence and decentralized ledger technologies. This research explores the profound impact of agentic artificial intelligence (AI) and Large Language Models (LLMs) on the evolution of financial services, moving from traditional automated systems to autonomous, self-driven entities capable of complex decision-making. By synthesizing foundational machine learning techniques in credit scoring with contemporary advancements in generative AI, the study identifies a shift from "FinTech" to "TechFin," where data-driven logic precedes financial function. The research delves into the architecture of specialized financial agents, such as FinRobot and StockAgent, examining their capacity for real-time market sentiment analysis and asset allocation. Furthermore, the article investigates the role of confidentiality-preserving blockchain systems in facilitating secure transactions within these agentic frameworks. Through an exhaustive theoretical elaboration, this work addresses the regulatory challenges and ethical considerations of delegated financial autonomy, proposing a synergistic model that integrates machine learning and agent-based modeling to enhance investment accuracy and operational efficiency. The findings suggest that while AI agents significantly augment knowledge work and underlying asset reviews, the future of financial stability depends on robust grounding mechanisms and multidisciplinary regulatory oversight.

KEYWORDS

Agentic Artificial Intelligence, Large Language Models, Fintech, Blockchain, Credit Scoring, Financial Autonomy, Asset Allocation.

INTRODUCTION

The historical trajectory of financial intermediation has reached a critical juncture where the human element in decision-making is increasingly augmented or replaced by sophisticated algorithmic entities. Traditionally, the banking sector functioned as a centralized arbiter of trust and creditworthiness, relying on standardized metrics and manual oversight. However, the emergence of Fintech has disrupted these traditional hierarchies, introducing a paradigm where technological innovation is the primary driver of financial value (Thakor, 2020). As financial institutions transition toward increasingly digitized models, the fundamental nature of the transaction has shifted from a physical exchange to a complex data-processing event. This evolution is not merely incremental; it represents a foundational shift in how risk is assessed, how assets are distributed, and how market participants interact.

A primary catalyst in this transformation is the application of machine learning classifiers to the problem of credit scoring. Unlike older linear models, modern machine learning approaches allow for the ingestion of vast, non-traditional datasets, providing a more nuanced and accurate assessment of borrower risk (Srivastava & Kumar, 2015). This technical capability has expanded the reach of financial services to previously underserved populations, but it has also introduced new questions

regarding algorithmic bias and the transparency of decision-making processes. As Taddy (2019) notes, the combination of machine learning and economics is now essential to optimize and accelerate business decisions, yet this optimization requires a deep understanding of the underlying economic incentives that guide algorithmic behavior.

Despite these advancements, a significant gap remains in the literature and practice regarding the move from automation to autonomy. While previous generations of FinTech focused on automating routine tasks, the current wave-often referred to as agentic AI-focuses on self-driven intelligence capable of executing multi-step financial strategies without constant human intervention (Bhat & Krishnan, 2025). This shift introduces a "TechFin" environment where the regulatory challenges of data-driven finance become paramount, as existing frameworks were designed for human-led institutions, not autonomous digital agents (Zetzsche, Buckley, Barberis, & Arner, 2018). The rise of Large Language Models (LLMs) has further complicated this landscape by providing agents with the ability to interpret and generate natural language, allowing them to process financial news, market sentiment, and regulatory filings with human-like proficiency (Jadhav & Mirza, 2025).

Recent studies have further highlighted the importance of agentic AI orchestration frameworks in enabling coordinated decision-making across complex digital ecosystems. Such frameworks facilitate interoperability among autonomous agents, enhance workflow automation, and support scalable enterprise transformation initiatives. The adoption of orchestration-based architectures demonstrates how agentic systems can effectively manage multi-step processes while maintaining adaptability and operational efficiency (Upadhyay, 2026).

The integration of these autonomous agents into market environments necessitates a secure and confidential infrastructure. Blockchain technology has been proposed as a solution, specifically through the design of confidentiality-preserving transaction processing systems that ensure data integrity while protecting sensitive financial information (Wang & Kogan, 2018). Without such protections, the deployment of AI agents in real-time trading or investment analysis could expose participants to unprecedented security risks. This research seeks to bridge the gap between these technical innovations and their broader economic and regulatory implications, providing a comprehensive framework for understanding the future of financial autonomy.

METHODOLOGY

The methodology of this research involves a multi-layered theoretical synthesis and architectural analysis of autonomous financial systems. We begin by examining the evolution of AI agents from simple rule-based scripts to complex, multi-agent frameworks capable of real-time searching and knowledge grounding (Krishnan, 2025). The focus is on how these agents are architected to handle the specific volatility and high-dimensionality of financial data. A key component of our methodology is the evaluation of the "StockAgent" and "FinRobot" platforms, which utilize LLM-based reasoning to simulate trading in real-world environments (Zhang et al., 2025; Yang et al., 2025). These platforms represent a move toward "agentic" intelligence, where the model does not just predict but also acts within a simulated orlive market.

To understand the accuracy and efficiency of these agents, we analyze the application of multi-agent frameworks in structured finance, specifically in the context of underlying asset reviews (Wan, Deng, Zou, & Xu, 2025). This involves a granular look at how agents can be programmed to perform synergistic tasks, such as data extraction, risk verification, and report generation, thereby reducing the margin for human error in complex financial products. The methodology also incorporates an assessment of "knowledge grounding" techniques, such as those used in the FinBloom model, which ensure that LLMs are anchored in real-time financial data to prevent "hallucinations" or the generation of factually incorrect information (Sinha, Agarwal, & Malo, 2025).

Furthermore, we explore the multidisciplinary integration of machine learning and agent-based modeling (ABM). As Zhang, Valencia, and Chang (2021) suggest, this synergy allows for the simulation of emergent market behaviors that neither ML nor ABM could capture in isolation. By modeling the interactions between thousands of autonomous agents, researchers can observe how market sentiment views influence asset allocation on a systemic level (Xing, Cambria, & Malandri, 2018). This study employs these theoretical models to explain the shift in investment analysis, where AI-agent collaboration is optimized to provide deeper insights than traditional quantitative methods (Han et al., 2025).

The concept of orchestration frameworks is particularly relevant when examining the operational architecture of autonomous financial agents. According to Upadhyay (2026), agentic orchestration mechanisms enable seamless communication between specialized AI agents, thereby improving task coordination, data exchange, and decision consistency across enterprise-scale environments. These principles can be extended to financial ecosystems where multiple autonomous agents collaborate to execute investment, compliance, and risk-management functions.

Finally, the methodology addresses the infrastructural requirements for secure financial agency. We evaluate the design principles of blockchain-based systems that prioritize confidentiality while allowing for the transparent audit trails required by financial regulators (Wang & Kogan, 2018). This includes a review of zero-knowledge proofs and other cryptographic techniques

that allow AI agents to verify transactions without revealing the underlying data. By synthesizing these diverse technological threads, the methodology provides a holistic view of the operational requirements for the next generation of data-driven finance.

RESULTS

The findings of this research indicate that the deployment of agentic AI significantly improves both the speed and accuracy of financial knowledge work. Generative AI agents are shown to augment knowledge work by performing complex synthesis of disparate data sources, a task that previously required senior-level human analysts (Ganesh et al., 2024). In the realm of equity markets, the application of LLMs has moved beyond simple sentiment analysis to more sophisticated techniques such as customized "FinGPT" search agents that can query foundation models to provide tailored investment insights (Tian et al., 2025; Li et al., 2025). These agents demonstrate a high degree of proficiency in navigating real-time financial information, providing a competitive edge in high-frequency environments.

In structured finance, the use of multi-agent frameworks for underlying asset reviews has resulted in a marked increase in efficiency. By delegating the review process to a network of specialized agents, institutions can process larger volumes of data with a higher degree of consistency than manual reviews (Wan et al., 2025). The results also highlight the importance of "intelligent asset allocation" guided by market sentiment. Models that integrate sentiment views into their optimization algorithms consistently outperform those relying solely on historical price data, suggesting that the qualitative insights provided by LLMs are essential for modern portfolio management (Xing et al., 2018).

A critical result of this study is the identification of the "FinTech to TechFin" transition. This transition signifies that the most successful financial entities are no longer those that simply use technology, but those that are built around technology as their core identity. This shift creates a data-driven finance environment where the volume and velocity of data ingestion determine market position (Zetzsche et al., 2018). Furthermore, the results show that blockchain integration is not just a security feature but a functional necessity for autonomous agency. Confidentiality-preserving systems allow for the secure execution of agent-led strategies, ensuring that the proprietary logic of the agent remains protected from competitors while maintaining regulatory compliance (Wang & Kogan, 2018).

However, the results also point to significant risks associated with the "black box" nature of autonomous financial agents. While agentic AI provides power in self-driven autonomy and customer engagement, it also introduces systemic risks if agents across different institutions converge on similar, flawed strategies (Bhat & Krishnan, 2025). The findings emphasize that while AI-agent collaboration optimizes investment research, it requires constant human oversight and sophisticated grounding mechanisms to prevent the amplification of market volatility (Han et al., 2025). The success of models like FinBloom suggests that real-time data grounding is the single most important factor in the reliability of generative AI in finance (Sinha et al., 2025).

DISCUSSION

The deep interpretation of these results suggests that we are entering an era of "Algorithmic Autonomy" where the boundary between the tool and the decision-maker is permanently blurred. The theoretical implications of agentic AI in finance go far beyond simple efficiency gains; they represent a fundamental restructuring of financial accountability. If an autonomous agent executes a trade that leads to market instability, the question oflegal and financial liability becomes complex (Zetzsche et al., 2018). Existing regulatory frameworks, which focus on human "know-your-customer" (KYC) and anti-money laundering (AML) protocols, are ill-equipped to handle entities that operate at microsecond speeds and make decisions based on multi-layered neural network outputs.

The emergence of agentic AI orchestration frameworks offers a practical pathway for governing interactions among autonomous financial agents. Research by Upadhyay (2026) suggests that orchestration-based architectures enhance system interoperability, improve governance mechanisms, and facilitate enterprise-wide digital transformation. In financial ecosystems, similar orchestration capabilities may support more transparent coordination between AI-driven trading, compliance, and customer engagement agents while reducing operational fragmentation.

A major point of discussion is the balance between automation and human expertise. While Ganesh et al. (2024) argue that generative AI agents augment knowledge work, there is a counter-argument that over-reliance on these systems could lead to a "de-skilling" of the financial workforce. If junior analysts are no longer required to perform the foundational research-because it is handled by FinGPT or similar agents-the long-term development of human expertise in the sector may be compromised. Furthermore, the "synergistic integration" of ML and ABM suggests that markets may become more predictable for those with superior computational power, potentially leading to a new form of digital divide in global finance (Zhang et al., 2021).

The limitations of the current technology are primarily centered on the "hallucination" problem and the lack of deep causal reasoning in LLMs. While models like FinBloom and agent frameworks for real-time searching attempt to ground LLMs in factual data, these models still lack a true understanding of economic theory or historical context (Sinha et al., 2025; Li et al., 2025). They are

essentially high-level pattern recognition engines. Consequently, their performance in "black swan" events-where historical patterns are no longer valid-remains highly suspect. The discussion highlights a need for a move toward "Hybrid Intelligence," where the creative and ethical reasoning of humans is combined with the analytical speed of AI agents (Han et al., 2025).

Future scope for research must focus on the "Interoperability of Financial Agents." As more institutions deploy platforms like FinRobot, the interaction between these different algorithmic species will become the primary driver of market dynamics (Yang et al., 2025). Understanding the "ecology" of these agents-how they compete, collaborate, and co-evolve-will be crucial for maintaining financial stability. Additionally, the development of more advanced blockchain architectures that can handle the transaction volume of millions of autonomous agents while preserving confidentiality is a significant technical challenge (Wang & Kogan, 2018). Regulatory research must also evolve to create "sandboxes" where agentic behaviors can be tested in simulated environments like StockAgent before being released into live markets (Zhang et al., 2025).

CONCLUSION

The convergence of agentic AI, Large Language Models, and blockchain technology marks a definitive turning point in the history of financial services. This research has demonstrated that autonomous agents are no longer a theoretical possibility but a practical reality that is already reshaping credit scoring, asset allocation, and structured finance reviews. By delegating complex analytical and operational tasks to AI agents, the financial sector is achieving unprecedented levels of efficiency and insights. However, this autonomy comes with significant challenges, ranging from the technical necessity of data grounding to the legal complexity of algorithmic liability.

The transition from FinTech to TechFin necessitates a new regulatory philosophy that recognizes data as the primary asset and algorithmic logic as the primary actor. While machine learning classifiers and LLMs provide the cognitive power for these agents, blockchain provides the necessary security and confidentiality for their operation. The synergy between these technologies, supported by multidisciplined modeling, offers a robust path forward for financial autonomy. Yet, the human element remains indispensable; the "grounding" of AI in real-world data and human ethical values is the only way to ensure that the algorithmic frontier of finance remains stable and beneficial for society at large.

Ultimately, the future of finance will be defined by the quality of the collaboration between human intelligence and agentic AI. As we continue to accelerate business decisions through these advanced technologies, we must remain vigilant against the systemic risks they pose. By fostering a multidisciplinary approach that combines economics, computer science, and law, we can build a financial ecosystem that is not only autonomous and efficient but also transparent and resilient in the face of an uncertain digital future.

REFERENCES

1. A. K. Bhat and G. Krishnan, "A Review of Agentic Artificial Intelligence: Power of Self-Driven AI in the Future of Financial Autonomy and Enhanced Customer Engagement," 2025 3rd International Conference on Sustainable Computing and Data Communication Systems (ICSCDS), Erode, India, 2025, pp. 1160-1165, doi:10.1109/ICSCDS65426.2025.11167368.

2. Ganesh, S., Ardon, L., Borrajo, D., Garg, D., Sehwag, U. M., & Narayanan, A. L. (2024). Generative AI agents for knowledge work augmentation in finance. Annual Review of Control, Robotics, and Autonomous Systems, 8(1), 189-210.

3. Han, X., Wang, N., Che, S., Yang, H., Zhang, K., & Xu, S. X. Enhancing investment analysis: optimizing AI-Agent collaboration in financial research. 538-46.

4. Jadhav, A., & Mirza, V. (2025). Large language models in equity markets: applications, techniques, and insights. Frontiers in Artificial Intelligence, 8, 1608365.

5. Krishnan, N. AI agents: evolution, architecture, and real-world applications.

6. Li, J., Zhang, J., Li, H., & Shen, Y. An agent framework for real-time financial information searching with large language models.

7. Sinha, A., Agarwal, C., & Malo, P. FinBloom: knowledge grounding large language model with real-time financial data.

8. Srivastava, P. R., & Kumar, V. (2015). An approach toward credit scoring using machine learning classifiers. Neural Computing and Applications, 28(1), 1-12.

9. Taddy, M. (2019). Business data science: Combining machine learning and economics to optimize, automate, and accelerate business decisions. McGraw-Hill Education.

10. Thakor, A. V. (2020). Fintech and banking: What do we know? Journal of Financial Intermediation, 41, 100833.

11. Tian, F., Byadgi, A., Kim, D. S., Zha, D., White, M., & Xiao, K. Customized fingpt search agents using

foundation models. 469-77.

12. Wan, X., Deng, H., Zou, K., & Xu, S. Enhancing the efficiency and accuracy of underlying asset reviews in structured finance: the application of multi-agent framework.

13. 13. Wang, Y., & Kogan, A. (2018). Designing confidentiality-preserving Blockchain-based transaction processing systems. International Journal of Accounting Information Systems, 30, 1-18.

14. Xing, F. Z., Cambria, E., & Malandri, L. (2018). Intelligent asset allocation via market sentiment views. IEEE Computational Intelligence Magazine, 13(4), 25-34.

15. Yang, H., Zhang, B., Wang, N., Guo, C., Zhang, X., & Lin, L. FinRobot: an open-source AI agent platform for financial applications using large language models.

16. Zetzsche, D. A., Buckley, R. P., Barberis, J. N., & Arner, D. W. (2018). From FinTech to TechFin: The regulatory challenges of data-driven finance. New York University Journal of Law & Business, 14(2), 393-446.

17. Zhang, C., Liu, X., Zhang, Z., Jin, M., Li, L., &Wang, Z. When AI meets finance (stockagent): large language model-based stock trading in simulated real-world environments.

18. Zhang, W., Valencia, A., & Chang, N. B. (2021). Synergistic integration between machine learning and agent-based modeling: a multidisciplinary review. IEEE Transactions on Neural Networks and Learning Systems, 34(5), 2170-90.

19. Upadhyay, H. (2026). Agentic AI Orchestration Frameworks for Composable Commerce Ecosystems: A Case Study of Enterprise Transformation . American Journal of Technology, 5(1), 40–54. https://doi.org/10.58425/ajt.v5i1.476