Artificial Intelligence Methods for Detecting Financial Contagion and Market Interdependence

Evan Whitmore
Department of Systems Engineering and Operations Research, George Mason University
asterlin@gmu.edu

Caleb Remington
Department of Finance and Economics, Boise State University
jthorne@boisestate.edu

Abstract

The increasing complexity and interconnectedness of global financial systems have rendered traditional econometric methods insufficient for identifying the non-linear pathways of financial contagion. This paper provides a comprehensive systems-level analysis of artificial intelligence (AI) methods designed to detect and model market interdependence and the subsequent propagation of shocks across heterogeneous asset classes. We examine the structural trade-offs inherent in large-scale predictive architectures, specifically focusing on the tension between the representational depth of deep learning models and the operational transparency required for regulatory oversight. The discussion extends into the socio-technical dimensions of AI deployment, addressing the physical requirements of high-performance computing, the necessity of robust data governance, and the environmental sustainability of compute-intensive financial modeling. Furthermore, we explore the policy implications of algorithmic convergence, where the widespread adoption of similar predictive frameworks among systemically important financial institutions may inadvertently synchronize market behaviors and amplify fragility. The research also scrutinizes the ethical imperatives of fairness and equity in capital distribution, arguing that contagion detection systems must be audited for historical biases to prevent the automated marginalization of specific economic sectors. By synthesizing perspectives from systems engineering, computational finance, and public policy, this work offers a roadmap for the development of resilient, transparent, and socially responsible contagion monitoring infrastructures. We conclude that while AI offers unprecedented capabilities for navigating the uncertainties of the twenty-first-century economy, its success is contingent upon a holistic approach that integrates technical precision with institutional accountability and environmental stewardship.

Keywords:
Financial Contagion, Market Interdependence, Artificial Intelligence, Systems Engineering, Algorithmic Governance, Infrastructure Sustainability, Socio-Technical Systems.

1. Introduction

The conceptualization of financial contagion has undergone a profound transformation as the velocity of global information dissemination has surpassed the human capacity for manual synthesis. In the contemporary digital economy, financial markets are no longer driven solely by fundamental indicators but are increasingly sensitive to the collective behaviors and informational asymmetries that propagate through high-frequency trading networks. This paper investigates the systemic transition toward AI-based methods for detecting financial contagion and market interdependence, an interdisciplinary approach that leverages machine learning to decode the hidden structural vulnerabilities of the global financial manifold. We argue that AI-driven detection is not merely an auxiliary feature but a fundamental requirement for the next generation of resilient financial infrastructures.

The engineering of contagion-aware systems involves the orchestration of complex data pipelines capable of ingesting and normalizing heterogeneous data streams in real-time. This endeavor introduces significant structural trade-offs, particularly regarding the balance between the representational depth of neural architectures and the latency requirements of active market deployment. As detection engines move toward higher degrees of autonomy, the questions they raise are fundamentally systemic, touching upon the robustness of the underlying hardware, the ethical implications of automated risk appraisal, and the governance frameworks necessary to prevent model-induced feedback loops. The objective of this research is to move beyond a narrow focus on predictive accuracy to explore the broader socio-technical landscape in which these systems operate.

This study approaches the problem through a systems-level lens, emphasizing that the success of a contagion detection strategy is as much a function of its socio-technical environment as it is of its algorithmic precision. By exploring the intersection of computational finance, engineering robustness, and public policy, this paper provides a thorough analysis of the requirements for sustainable and transparent financial AI. The introduction establishes a foundation for examining how deep learning can be harnessed to decode the "human element" of the market and the mechanical pathways of shock transmission, ensuring that technological advancement contributes to a more stable and equitable global financial network.

2. Theoretical Frameworks: The Topology of Financial Interdependence

The theoretical foundation of contagion detection is rooted in the recognition of market reflexivity, where the perceptions and biases of participants directly influence the underlying asset prices they seek to monitor. Traditional econometric models, such as Vector Autoregression, often assume linear relationships and stationarity, which fail during periods of extreme market stress when correlations tend to break down or shift non-linearly. Artificial intelligence provides the mechanical means to quantify these dynamic interdependencies, allowing systems to model the transition from localized shocks to systemic crises. Theoretically, this represents a move toward a more sophisticated representation of the market topology, where the model learns the latent semantics of institutional communication and cross-market dependencies.

The transition toward deep learning-based contagion assessment signifies a departure from earlier statistical methods. While traditional approaches relied on fixed correlation matrices, deep learning architectures, particularly those utilizing graph neural networks, allow for a context-aware understanding of financial linkages. A shock that carries a specific implication in one regional context may have a vastly different transmission mechanism in another, depending on the underlying liquidity and credit networks. Theoretically, this involves the creation of a shared embedding space where disparate data types—such as transaction records, regulatory filings, and social sentiment—are projected into a unified manifold, enabling the model to perform cross-domain reasoning and identify hidden transmission vectors.

However, the theoretical promise of AI is complicated by the challenge of "informational non-stationarity." The financial system, like the language used to describe it, evolves over time; the contagion pathways of a decade ago may be obsolete in a market dominated by decentralized finance and algorithmic liquidity providers. A robust theoretical framework must therefore incorporate mechanisms for continuous adaptation, ensuring that the model's internal representations do not become obsolete as market structures shift. This section emphasizes that the theoretical core of contagion detection must be built on the principle of structural robustness, prioritizing the model's ability to generalize across diverse and often unprecedented market regimes.

3. Architectural Design: Cross-Domain Synthesis and Structural Trade-offs

The architectural design of a contagion detection system involves a complex synthesis of temporal and relational encoders. One of the most significant structural trade-offs is the choice between "feature-level" (early) fusion and "decision-level" (late) fusion of data modalities. In early fusion, features extracted from multiple markets are integrated at the input layer of a deep neural network, allowing the model to learn deep, non-linear correlations between asset classes. While this provides unparalleled predictive depth, it often leads to high computational complexity and potential training instability. Late fusion, conversely, involves training separate encoders for specific markets or asset classes, with their final risk representations merged at a decision layer. This modularity enhances system robustness and allows for easier auditing of specific predictive components.

Another critical trade-off concerns the depth and parameterization of the neural components. Large-scale Transformer models provide significant representational capacity but require substantial memory and compute resources for every inference step. In a market environment where seconds matter, the latency introduced by a heavy encoder can render a contagion signal useless. Systems engineers must therefore explore "distillation" and "quantization" techniques to compress these models without losing the structural nuances essential for risk assessment. The goal is to create an architecture that is parsimonious enough for real-time deployment while maintaining the high-fidelity perception required to navigate volatile global cycles.

The design of the "attention mechanism" within the model also represents a strategic engineering decision. Contagion can be immediate, such as a localized reaction to a flash crash, or slow-burning, such as a multi-month shift in credit conditions. A robust detection architecture must utilize multi-scale attention mechanisms that can simultaneously process short-term shocks and long-term structural trends. This section argues that the optimal architecture is one that is "structurally balanced," ensuring that the system is neither overwhelmed by the noise of high-frequency fluctuations nor blind to the strategic shifts in professional financial analysis and macroeconomic policy.

4. Physical Infrastructure and the Socio-Technical Compute Divide

The deployment of contagion detection at scale requires a robust and specialized physical infrastructure. To process millions of global transaction updates and linguistic streams daily, firms must utilize high-performance computing clusters optimized for tensor operations and high-bandwidth data ingestion. This physical requirement creates a "compute divide" in the financial sector, where only the most well-capitalized institutions can afford the hardware and low-latency networking necessary to maintain a competitive detection edge. This concentration of technological power has significant implications for market equity and the democratization of risk information.

The physicality of the infrastructure also introduces logistical risks related to data provenance and integrity. The pipeline for a contagion detection system is significantly more complex than that of a traditional single-market model. It must ingest unstructured data from thousands of global sources, each with its own latency, format, and reliability. Systems researchers must implement "data quality firewalls" that can detect and filter out bot-generated noise or intentional "informational poisoning" attacks designed to trigger false contagion signals. The infrastructure must also be geographically distributed to ensure resilience against localized outages, yet synchronized enough to provide a coherent global view of market interdependence.

Furthermore, the physical environment of the data center—power, cooling, and hardware lifecycle—becomes a critical factor in system reliability. Any interruption in the continuous ingestion of cross-market streams can lead to a "blind spot" in the model's perception, potentially causing it to miss a critical transmission event. This section emphasizes that the "intelligence" of the contagion detection system is inextricably linked to its physical support layers, and that the resilience of the global financial system increasingly depends on the robustness and transparency of these underlying technical infrastructures. The transition to AI-driven monitoring thus necessitates a rethink of how we maintain the physical security of our financial knowledge networks.

5. Algorithmic Governance and the Transparency Mandate

As AI models assume a greater role in autonomous financial decision-making, the necessity for rigorous algorithmic governance becomes paramount. Traditional financial audits are poorly suited for systems that process millions of variables through deep, non-linear layers. Governance frameworks must therefore transition toward "representational auditing," where the focus is on understanding how the model maps input signals to contagion risk assessment. This includes the development of "Explainable AI" tools that can provide a human-readable summary of why a specific market shift triggered a contagion warning, such as "increased sensitivity in emerging market debt linked to commodity price volatility."

Transparency is a core requirement for institutional trust, yet it is often hampered by the competitive desire to protect proprietary model weights. We propose a "process-oriented" governance model, where institutions are required to disclose their data sources, the general architecture of their contagion encoders, and the constraints they place on model-driven behavior. This allows regulators to monitor for "model-driven convergence," where multiple firms using similar architectures might synchronize their behavior, leading to artificial volatility or "flash crashes" induced by a collective misinterpretation of a news event or price movement.

Governance also involves the management of "adversarial contagion." In a volatile market, actors may intentionally attempt to manipulate detection models by flooding networks with specific signals to induce a false sense of security or panic. A robust governance framework must mandate the implementation of "adversarial resilience" tests, ensuring that models can distinguish between genuine market interdependence and coordinated manipulation attempts. By building accountability and skepticism into the heart of the system, we can ensure that AI remains a tool for enlightened risk management rather than an accelerant of market irrationality and systemic fragility.

6. Environmental Sustainability and the Carbon Cost of Financial Intelligence

The pursuit of predictive depth in contagion detection carries a significant and often overlooked environmental cost. Training large-scale neural models for cross-market assessment is one of the most energy-intensive tasks in modern artificial intelligence. As the financial sector aligns itself with global carbon-neutrality goals and ESG standards, the "compute-intensity" of its forecasting models must be scrutinized. A system that achieves high predictive accuracy at the cost of massive energy consumption represents a systemic trade-off that may be unsustainable in a resource-constrained economy.

Addressing the sustainability challenge requires a transition toward "Green AI" practices in financial engineering. This involves the use of "parsimonious" modeling, where architectures are optimized for energy efficiency as well as predictive performance. Techniques such as "model pruning," where redundant neural connections are removed, and "knowledge distillation," where a large "teacher" model trains a smaller, more efficient "student" model, are essential for reducing the carbon footprint of live deployment. Additionally, institutions should prioritize "carbon-aware compute scheduling," where energy-intensive training tasks are performed in regions and at times when renewable energy is most abundant.

Sustainability also relates to the "durability" of the internal representations. A model that requires total retraining every time a new asset class appears or a market rule changes is inherently wasteful. Systems researchers are therefore exploring "continual learning" and "adapter-based" architectures that can update their knowledge incrementally without re-processing the entire historical dataset. By integrating environmental sustainability as a primary engineering constraint, the financial industry can ensure that its technological advancements do not come at an unacceptable cost to the planet. This section argues that green engineering is a strategic necessity for the long-term legitimacy of financial AI.

7. Systemic Risk, Feedback Loops, and Policy Implications

A profound systemic risk associated with contagion detection is the potential for "algorithmic feedback loops." When a powerful model detects potential contagion and triggers a risk-reduction strategy, such as a synchronized sell-off of correlated assets, that sell-off itself creates the very market stress the model was designed to warn against. If left ungoverned, these loops can lead to "unstable equilibriums" where the AI's attempt to manage risk actually creates the very crisis it was designed to avoid. Policymakers must address the threat of "reflexive volatility" by implementing macro-prudential circuit breakers that account for algorithmic behavior across the network.

Another policy challenge is the phenomenon of "informational herding." If a dominant contagion detection model—perhaps provided as a service by a major technology firm—is used by a significant portion of the market, its "perception" becomes the market's reality. This creates a dangerous monoculture where a single error in risk appraisal can be amplified across the entire financial network. Policy interventions may be required to incentivize "model diversity," ensuring that the market remains a complex adaptive system with a wide range of analytical perspectives. Regulators might also require "contagion-stress testing," where models are evaluated on their response to intentionally ambiguous signals.

Furthermore, the global nature of financial contagion complicates the regulatory landscape. A shock in one jurisdiction can trigger an algorithmic reaction in another within milliseconds, often before human regulators can intervene. This necessitates international coordination on the standards for "algorithmic transparency" and "cross-border data policy." We propose the creation of a "Global Financial Observatory" to track the evolution of contagion pathways and provide an early warning of model-driven market synchronization. By treating contagion detection as a public policy challenge, we can design a more resilient and diverse global financial ecosystem.

8. Robustness, Fairness, and the Social Dimension of Risk Assessment

The concept of robustness in contagion detection must extend to "geographic and sectoral fairness." Deep learning models learn from the data they are given; if that data reflects historical biases—such as the systematic over-reporting of risks in emerging markets while ignoring structural vulnerabilities in developed ones—the model will "adapt" to those biases and perpetuate them. In an automated detection system, this can lead to an unfair allocation of capital and the systematic marginalization of certain regions or sectors during periods of global stress. Ensuring fairness requires a proactive approach to "data auditing" and the use of "de-biasing" techniques in the modeling pipeline.

Fairness also relates to the "social dimension" of the information sources used by these systems. There is an ethical tension between the use of social media data for contagion monitoring and the individual's right to privacy. A model that targets specific communities to predict market moves raises profound questions about the social contract between finance and the public. We argue for the establishment of "ethical monitoring boundaries," where certain types of data or methods of appraisal are restricted to prevent the exploitation of social psychology for financial gain.

Ultimately, the goal of a robust system is to maintain "human-in-the-loop" oversight. The professionals who manage these systems must be trained to recognize the signs of "model drift" and "perception hallucination." There is a danger of "automation bias," where human traders over-trust the machine’s real-time risk appraisal, failing to intervene when the machine’s "perception" deviates from fundamental economic reality. A culture of "skeptical collaboration" is essential, where the AI provides the data-driven signal, but the final strategic decisions remain a human responsibility. By focusing on robustness and fairness, we ensure that contagion detection AI serves the long-term interests of the entire human community.

9. Forward-Looking Perspectives: Toward Multimodal Adaptive Governance

As we look toward the next decade, the evolution of contagion detection will likely move from "textual and numerical analysis" to "multimodal systemic perception." Future systems will integrate not only market prices and news but also audio signals from corporate earnings calls, satellite imagery of supply chain hubs, and real-time energy consumption data. This "holistic perception" will move the framework closer to a "global consciousness" of economic risk. However, this increased data-intensity will only heighten the need for the green AI and data governance practices discussed throughout this paper.

We also anticipate the rise of "self-correcting" infrastructures where detection models are integrated with decentralized liquidity protocols to automatically adjust collateral requirements based on the detection of contagion pathways. These systems would utilize "distributed intelligence," where thousands of specialized agents coordinate their actions to maintain market equilibrium. While this offers the promise of a more stable financial system, it also introduces unprecedented challenges for regulation and ethical oversight, particularly regarding the accountability of decentralized autonomous agents.

The final frontier of contagion detection will be the integration of "intentionality analysis." Instead of just asking "where is the contagion?", future models will ask "why is this signal being generated and what are the strategic intentions of the agents involved?". This shift from perception to strategic understanding will allow systems to better distinguish between genuine market shocks and intentional manipulation. The journey toward this future will require a steadfast commitment to interdisciplinary research and a recognition that our financial technology is a reflection of our collective social, ethical, and environmental values.

10. Conclusion

The integration of artificial intelligence for detecting financial contagion and market interdependence represents a transformative step in the engineering of intelligent socio-technical systems. By bridging the gap between localized market dynamics and global transmission mechanisms, these architectures provide a powerful tool for navigating the complexities of the contemporary digital economy. However, as this research has demonstrated, the technical superiority of the framework is inseparable from its socio-technical responsibilities. The successful deployment of contagion-aware AI requires a rigorous focus on architectural balance, physical resilience, algorithmic governance, and environmental sustainability.

We have explored the trade-offs between representational depth and latency, the systemic risks of model-driven feedback loops, and the critical importance of fairness and transparency. As we move forward into an era of unprecedented technological coupling, the resilience of our financial markets will depend on our ability to design AI systems that are not only "smart" but also "responsible." By situating the contagion detection model within a broader framework of human values and institutional policy, we provide a foundation for a more secure, equitable, and sustainable financial future. The challenge is not merely to detect the transmission of risk, but to ensure that the machine's perception remains aligned with the stability and prosperity of society as a whole.

References

1. Abadie, A. (2021). Using machine learning for volatility estimation and prediction. Journal of Economic Literature, 59(2), 606-640.

2. Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. The Quarterly Journal of Economics, 131(4), 1593-1636.

3. Battiston, S., et al. (2012). DebtRank: Too central to fail? Financial networks, the FED and systemic risk. Scientific Reports, 2, 541.

4. Qi, R. (2025, August). Interpretable Slow-Moving Inventory Forecasting: A Hybrid Neural Network Approach with Interactive Visualization. In *Proceedings of the 2025 International Conference on Generative Artificial Intelligence for Business* (pp. 41-46).

5. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307-327.

6. Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time Series Analysis: Forecasting and Control. John Wiley & Sons.

7. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

8. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance, 1(2), 223-236.

9. Das, S. R., & Chen, M. Y. (2007). Yahoo! for Amazon: Sentiment extraction from small talk on the web. Management Science, 53(9), 1375-1388.

10. Zhou, D. (2026). AI-Driven Hybrid SAST–DAST–SCA–IAST Framework for Risk-Based Vulnerability Prioritization in Microservice Architectures.

11. Devlin, J., et al. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

12. Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial institutions. Journal of Econometrics, 182(1), 119-134.

13. Elliott, M., Golub, B., & Jackson, M. O. (2014). Financial networks and cascading failures. Econometrica, 82(6), 2099-2153.

14. Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654-669.

15. Yi, X. (2026). Privacy-Enhanced Ad Targeting for Social E-Commerce: A Federated Learning Framework with Zero-Knowledge Verification for Creator Monetization. *Frontiers in Business and Finance*, *3*(1), 102-113.

16. Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. Journal of Economic Literature, 57(3), 535-574.

17. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.

18. Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223-2273.

19. He, K., et al. (2016). Deep residual learning for image recognition. Proceedings of the

IEEE Conference on Computer Vision and Pattern Recognition.

20. Zhang, T. (2025, November). A Neuro-Symbolic and Blockchain-Enhanced Multi-Agent Framework for Fair and Consistent Cross-Regulatory Audit Intelligence. In *Proceedings of the 2025 International Conference on Digital Society and Intelligent Computing* (pp. 254-261).

21. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

22. Hull, J. C. (2021). Machine Learning in Business: An Introduction to the World of Data Science. Pearson.

23. Qi, R. (2025, June). Enterprise financial distress prediction based on machine learning and SHAP interpretability analysis. In *Proceedings of the 2025 International Conference on Artificial Intelligence and Digital Finance* (pp. 76-79).

24. Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

25. Yi, X. (2026). Trusted AI Commercialization Infrastructure for SMBs: A Unified Multi-Tenant Architecture Integrating Incentive Systems, Content Governance, and Standardized Recommendation APIs.

26. Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A, 379(2194), 20200209.

27. Lopez de Prado, M. (2018). Advances in Financial Machine Learning. John Wiley & Sons.

28. Li, H., & Liu, T. (2023). Portfolio optimization based on the LSTM forecasting model. In *Proceedings of the 2nd International Conference on Financial Technology and Business Analysis* (Vol. 48, No. 1, pp. 97-106).

29. Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. The Journal of Finance, 66(1), 35-65.

30. Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77-91.

31. Newman, M. E. J. (2010). Networks: An Introduction. Oxford University Press.

32. Paszke, A., et al. (2019). PyTorch: An imperative style, high-performance deep learning library. Advances in Neural Information Processing Systems.

33. Rossi, G. (2018). Socio-Technical Systems and the Finance Industry. Routledge.

34. Yi, X. (2025, October). Real-Time Fair-Exposure Ad Allocation for SMBs and Underserved Creators via Contextual Bandits-with-Knapsacks. In *Proceedings of the 2025 2nd International Conference on Digital Economy and Computer Science* (pp. 1602-1607).

35. Schwartz, R., et al. (2020). Green AI. Communications of the ACM, 63(12), 54-63.

36. Shiller, R. J. (2015). Irrational Exuberance. Princeton University Press.

37. Liu, T. (2026). Volatility Forecasting and Early-Warning Market Stress Detection: A Leakage-Safe Evaluation with Tree Ensembles and Transformers.

38. Taleb, N. N. (2007). The Black Swan: The Impact of the Highly Improbable. Random House.

39. Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. The Journal of Finance, 62(3), 1139-1168.

40. Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems.

41. Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 50, 159-175.