Graph Neural Networks for Cross-Market Financial Prediction and Systemic Risk Modeling

Shawn Nolan
Department of Computer Science and Information Systems, Bradley University
cmaestre@bradley.edu

Arthur Prescott
Department of Economics and Finance, University of South Alabama
tarisaka@southalabama.edu

Abstract

The global financial ecosystem is characterized by an intricate web of interdependencies where shocks in one asset class or geographic region propagate rapidly across traditional market boundaries. Conventional predictive models, often rooted in time-series analysis of isolated variables, frequently fail to capture the relational dynamics and topological shifts inherent in these interconnected systems. This paper investigates the utility of Graph Neural Networks (GNNs) as a foundational architecture for cross-market financial prediction and systemic risk modeling. Unlike standard deep learning approaches that treat financial data as Euclidean sequences, GNNs operate directly on the non-Euclidean graph structures defined by supply chains, ownership networks, and correlation matrices. We conduct a system-level analysis of GNN deployment, emphasizing the structural trade-offs between graph density and computational scalability. The discussion extends to the socio-technical infrastructures required to sustain high-fidelity relational modeling, addressing critical issues of data governance, algorithmic fairness, and the environmental sustainability of large-scale graph processing. Furthermore, we examine the policy implications of GNN-driven risk assessment, arguing that while these models provide superior detection of contagion pathways, they also introduce novel systemic vulnerabilities related to model convergence and data poisoning. By synthesizing perspectives from graph theory, systems engineering, and financial policy, this research proposes a robust framework for integrating relational intelligence into the global financial stability architecture.

Keywords:
Graph Neural Networks, Systemic Risk, Cross-Market Prediction, Financial Infrastructure, Algorithmic Governance, Relational Learning, Socio-Technical Systems.

1. Introduction

The conceptualization of financial markets has undergone a radical shift from a collection of discrete, efficient entities toward a vision of a singular, highly coupled global network. In this contemporary landscape, the "intelligence" of a financial system is defined not only by its ability to process local price signals but also by its capacity to decode the hidden relationships between disparate asset classes, geopolitical actors, and corporate entities. Standard predictive paradigms, dominated by recurrent and convolutional neural networks, excel at extracting features from temporal sequences but remain fundamentally limited in their ability to model the explicit and implicit connections that bind the market together. Graph Neural Networks (GNNs) represent a significant architectural breakthrough in this regard, offering a mathematical and computational framework designed to operate on the very topology of the financial world.

This paper explores the integration of GNNs into the core of financial forecasting and risk management infrastructures. We move beyond a purely technical evaluation of predictive accuracy to examine the systemic requirements for deploying relational models in high-stakes, regulated environments. The shift to GNNs necessitates a rethinking of data pipelines, moving from flat tabular structures to complex, heterogeneous graphs that incorporate corporate hierarchies, global trade flows, and sentiment-based correlations. This transition introduces profound questions regarding the robustness of graph-based representations, the governance of the data used to construct these networks, and the physical infrastructure needed to manage the high computational costs of message-passing algorithms.

As GNNs become more prevalent in institutional risk assessment, they begin to influence the behavior of the markets they monitor. This reflexivity creates a socio-technical feedback loop where the model’s identification of a "contagion pathway" may lead to preemptive actions that either mitigate the risk or, paradoxically, accelerate its propagation. Our analysis addresses these complexities by situating GNNs within the broader context of global financial policy and sustainability. This introduction serves as a foundation for a deep dive into how relational intelligence can be harnessed to build a more resilient and transparent financial system, while simultaneously acknowledging the new classes of risk that such advanced technologies introduce.

2. Theoretical Evolution: From Correlation to Relational Topology

The evolution of systemic risk modeling has historically mirrored the increasing complexity of global trade. Early models often relied on simple correlation matrices, which, while useful for identifying synchronous price movements, failed to provide insights into the causal mechanisms or directional flows of market stress. The limitations of these linear, symmetric measures became glaringly apparent during the 2008 financial crisis, where the collapse of the subprime mortgage market cascaded through an opaque network of derivatives and interbank lending that traditional models were not designed to visualize. The subsequent shift toward network science in finance allowed for a more structural view of "too-interconnected-to-fail" institutions, yet these early network models were often static and lacked the capacity for complex feature learning.

Deep learning initially addressed the feature-learning gap but did so primarily through Euclidean architectures. These models assume that data points exist in a structured grid or sequence, which is an poor approximation for the messy, irregular relationships found in financial systems. Graph Neural Networks solve this by employing a message-passing paradigm where each node—representing a stock, a sector, or a national economy—aggregates information from its neighborhood. This allows the model to learn representations that are sensitive to the local and global topology of the market. The theoretical shift from correlation-based modeling to relational topology represents a move from observing what is happening to understanding where and how it is likely to spread.

However, the theoretical promise of GNNs is complicated by the "heterogeneity" and "dynamic nature" of financial graphs. Unlike a static social network, a financial graph changes its edges based on shifting market sentiment, regulatory changes, and economic reports. A robust GNN architecture must therefore incorporate dynamic edge weights and temporal components to account for the velocity of market shifts. This section argues that the theoretical foundation of GNNs in finance must be built on "relational inductive biases"—the assumption that the relationship between entities is as informative as the entities themselves. This perspective is essential for modeling the non-linear "phase transitions" that occur during periods of extreme market stress.

3. Architectural Trade-offs and the Scalability of Message-Passing

In the design of a GNN-based system for cross-market prediction, the engineer is confronted with a fundamental trade-off between the depth of information aggregation and the stability of the model. Deep GNNs, which aggregate information from many "hops" away in the graph, theoretically capture long-range dependencies across the global economy. Yet, in practice, these models often suffer from "over-smoothing," a phenomenon where the features of all nodes converge to a similar value, rendering the model's output useless for distinguishing between high-risk and low-risk entities. Preventing over-smoothing while maintaining the benefits of a global perspective requires sophisticated architectural interventions, such as skip connections, normalization layers, and attention mechanisms that weight the importance of different neighbors.

The computational scalability of GNNs represents a second critical trade-off. Standard message-passing algorithms require the storage and processing of the entire graph, which, for a global financial network incorporating millions of nodes and billions of edges, exceeds the memory capacity of even the most advanced hardware. Systems researchers have addressed this through "graph sampling" and "subgraph-based training," where the model learns from smaller, representative portions of the network. While these techniques reduce the compute burden, they introduce the risk of missing critical "weak ties"—distant or infrequent connections that may nonetheless serve as conduits for systemic contagion. Choosing the right sampling strategy is thus as much an economic and risk-management decision as it is a technical one.

Furthermore, the "spectral" versus "spatial" design of GNNs offers a distinct set of trade-offs. Spectral GNNs utilize the graph Laplacian to operate in the frequency domain, offering a powerful way to identify structural patterns but struggling with graphs of varying sizes or structures. Spatial GNNs, which define the convolution directly on the graph's physical neighbors, are more flexible and scalable but can be less mathematically rigorous in their feature extraction. For a cross-market system that must integrate everything from micro-cap stocks to sovereign debt, a hybrid spatial-temporal architecture is often necessary. This section emphasizes that the architectural choice is never neutral; it reflects a specific set of assumptions about how risk propagates and which market participants are most relevant to the prediction.

4. Data Governance and the Construction of Financial Knowledge Graphs

The effectiveness of a Graph Neural Network is entirely dependent on the quality and comprehensiveness of the underlying graph. In a financial context, constructing this graph—often referred to as a "Financial Knowledge Graph"—is a monumental data governance challenge. Unlike price data, which is readily available and standardized, relational data is often fragmented, proprietary, and highly sensitive. Building a graph that includes corporate ownership requires navigating complex legal disclosures; mapping supply chains requires access to private trade data; and modeling sentiment-based links involves the high-velocity processing of news and social media.

This fragmentation creates a risk of "informational siloing," where only the largest financial institutions have the resources to build a truly global knowledge graph. From a systems perspective, this raises concerns about market fairness and transparency. If a GNN-driven risk model at a major bank identifies a systemic vulnerability using private data, the rest of the market remains "blind" to that risk until it manifests as a crisis. Policy interventions may be required to mandate the creation of "public-interest graphs"—anonymized, high-fidelity relational datasets that can be used by regulators and smaller institutions to monitor systemic health. Such an infrastructure would serve as a public good, similar to a weather monitoring system or a public health database.

Moreover, the governance of the graph construction process must account for "edge reliability." Not all connections in a financial graph are equal; a correlation based on historical price movements is fundamentally different from a structural link like a credit default swap or a shared board member. A robust GNN system must incorporate "uncertainty quantification" at the edge level, acknowledging that the graph itself is a probabilistic estimation of reality. Governance frameworks must mandate that firms disclose not just their model's predictions, but the "topology of their assumptions"—the specific network structure that led to a given risk assessment. This transparency is essential for building social trust in automated risk-management systems.

5. Infrastructure, Deployment, and the Environmental Cost of Graph Computing

The deployment of Graph Neural Networks for real-time financial monitoring requires a massive investment in specialized physical infrastructure. GNN operations are notorious for their irregular memory access patterns, which are poorly suited for traditional central processing units (CPUs) and even many standard graphics processing units (GPUs). To run GNNs at the scale and speed required for cross-market prediction, firms are increasingly turning to "Graph Processing Units" (IPUs) and custom application-specific integrated circuits (ASICs) designed to optimize the sparse matrix operations inherent in message-passing. This specialized hardware must be housed in high-density data centers with ultra-low-latency links to market exchanges, creating a significant barrier to entry for smaller market participants.

The environmental sustainability of this infrastructure is a growing concern. Training a large-scale, dynamic GNN requires an order of magnitude more energy than traditional time-series models due to the complexity of the data aggregation steps. In the context of the global transition to Net Zero, the financial industry must confront the carbon footprint of its predictive power. A "Green GNN" framework would involve the development of energy-efficient aggregation methods, the use of "knowledge distillation" to compress large graph models into smaller, more efficient versions for deployment, and the strategic location of data centers in regions with high renewable energy capacity.

Furthermore, the deployment phase involves a continuous "GraphOps" (Graph Operations) cycle. Because the financial world is constantly evolving, the graph must be updated in near real-time as new corporate filings are released, new trade deals are signed, and new market correlations emerge. This requires a streaming data infrastructure capable of performing "incremental graph updates" without requiring a full retraining of the model. The reliability of this infrastructure is a matter of systemic importance; a failure in the graph-update pipeline could lead to a model using outdated topology, potentially missing the onset of a contagion event. This section highlights that the resilience of a GNN system is as much a function of its physical and logistical support as it is of its neural architecture.

6. Algorithmic Fairness and the Bias of Topology

The concept of "fairness" in a Graph Neural Network is uniquely complex because bias can emerge not just from node attributes (e.g., the sector of a stock) but from the very structure of the graph. If a GNN is trained on a financial graph where certain regions or industries are more "central" or highly connected than others due to historical luck or colonial legacy, the model will naturally learn that these entities are "systemically important." This can lead to a self-fulfilling prophecy where the AI directs more liquidity and lower risk premiums toward dominant players, while systematically under-valuing or over-estimating the risk of emerging markets or innovative startups that exist on the "periphery" of the graph.

Correcting for "topological bias" requires a proactive approach to graph engineering. This might involve "graph de-biasing" techniques, such as adding synthetic edges to under-represented regions or utilizing "fairness-aware" message-passing that prevents the model from over-emphasizing the features of highly connected nodes. However, there is a fundamental tension between "accuracy" and "fairness" in systemic risk modeling. If a certain institution is truly a central node in the interbank lending network, the model must identify it as a systemic risk. The challenge is to distinguish between "structural importance" that is grounded in economic reality and "topological noise" that is a reflection of historical bias.

This social dimension of risk modeling also touches upon the "deskilling" of the regulatory workforce. As GNNs provide increasingly sophisticated visualizations of systemic contagion, there is a risk that human regulators will defer entirely to the machine's topological intuition. It is essential to maintain a "human-in-the-loop" framework where the GNN serves as a decision-support tool rather than an autonomous judge. Regulators must be trained to interrogate the graph's structure—asking why a certain edge exists and what would happen if a different connectivity assumption were made. By treating fairness as a first-order system property, we can ensure that GNNs contribute to a more equitable and stable global economy.

7. Model Convergence, Systemic Fragility, and Policy Responses

A profound systemic risk associated with the institutional adoption of Graph Neural Networks is the phenomenon of "model convergence." If the majority of the world's systemically important financial institutions (SIFIs) and regulatory bodies utilize similar GNN architectures trained on the same foundational knowledge graphs, they are likely to reach identical conclusions about market risk. This synchronization can lead to "crowded trades" and simultaneous risk-shedding, where every autonomous system attempts to exit the same "pathway of contagion" at once. This collective behavior can turn a predicted risk into a realized crisis, exhausting liquidity and accelerating the very collapse the models were designed to prevent.

Addressing this fragility requires a new set of policy tools. Central banks and regulators might consider "algorithmic diversity mandates," where firms are required to disclose their general model DNA and are incentivized to use diverse graph-construction techniques. Furthermore, the implementation of "relational circuit breakers" may be necessary. These would be mechanisms that detect when a synchronized movement is occurring across a specific graph topology and temporarily halt trading or require manual human intervention to break the feedback loop. The goal is to prevent the market from becoming a "monoculture of intelligence" that is highly accurate in normal times but catastrophically fragile during periods of extreme uncertainty.

Policy responses must also address the threat of "data poisoning" in financial graphs. Because GNNs rely on the relationships between nodes, a malicious actor could theoretically influence the model's predictions by creating "fake" relationships—such as through wash-trading to create artificial correlations or by setting up shell companies to manipulate the supply-chain graph. Protecting the integrity of the global financial knowledge graph is thus a matter of national and economic security. We argue for the creation of an "International Graph Oversight Board" that would be responsible for verifying the quality of foundational relational data and developing standards for graph-based risk reporting.

8. Cross-Domain Comparisons: Learning from Biology and Cybersecurity

To refine the use of GNNs in finance, it is instructive to look at other domains that manage complex relational systems, such as computational biology and cybersecurity. In biology, GNNs are used to model protein-protein interaction networks to predict how a virus might propagate through a cellular system. The "resilience" of a biological system is often found in its redundancy and modularity—properties that financial engineers could strive to replicate by incentivizing a more modular global banking structure that prevents a single node failure from cascading.

In cybersecurity, GNNs are employed to detect lateral movement within a network—identifying how an attacker moves from one server to another to reach a target. This is remarkably similar to how market stress moves from a commodity market to an equities market and finally to the banking sector. By adopting "zero-trust" architectures from cybersecurity, financial systems could be designed to isolate specific graph partitions during a crisis, preventing the "lateral movement" of financial contagion. These cross-domain insights suggest that the future of systemic risk modeling lies in treating the financial market as a dynamic "infrastructure" that requires the same level of protection and resilience as a power grid or a telecommunications network.

The comparison also highlights the importance of "anomaly detection" over simple "trend prediction." In both biology and cybersecurity, the most critical signal is often the one that doesn't fit the established relational pattern. For finance, this means using GNNs to identify "topological anomalies"—sudden changes in the connectivity of a market that have no obvious economic cause. These anomalies are often the first signs of a looming crisis or a sophisticated case of market manipulation. By leaning into these interdisciplinary perspectives, we can move toward a more "immune-system" like approach to financial stability, where GNNs act as the sensors that detect and neutralize systemic pathogens.

9. Forward-Looking Perspectives: Toward Autonomous Resilience

As we look toward the next decade, the role of GNNs in finance will likely evolve from "monitoring" systemic risk to "actively mitigating" it through autonomous resilience. We anticipate the development of "Self-Healing Market Graphs," where the system identifies potential contagion pathways and automatically suggests—or even executes—portfolio rebalancing or liquidity injections to dampen the propagation of stress. This level of autonomy would represent a massive leap in market efficiency, but it also raises profound questions about the "moral agency" of AI systems. If an autonomous GNN redirects capital away from a struggling economy to protect the global network, who is responsible for the social consequences?

The future will also see the rise of "Federated Graph Learning," where multiple institutions contribute to a shared global risk model without ever sharing their underlying private data. Using techniques like differential privacy and secure multi-party computation, a "global graph" could be trained to identify systemic vulnerabilities while maintaining the commercial confidentiality of the participants. This would solve the data governance dilemma and provide a more comprehensive view of global risk than any single institution could achieve alone.

Finally, we anticipate a shift from "graph prediction" to "graph intervention." Instead of asking "what will happen to this network," future models will ask "how can we change this network to make it more stable." This involves using GNNs in combination with reinforcement learning to design better financial regulations and more resilient institutional structures. The goal is to move from a reactive posture to a generative one, where we use our understanding of relational intelligence to build a financial system that is fundamentally designed for stability, fairness, and long-term sustainability.

10. Conclusion

Graph Neural Networks represent a transformative tool for cross-market financial prediction and systemic risk modeling. By moving beyond the limitations of Euclidean architectures and isolated price signals, GNNs allow us to visualize and manage the global economy as the interconnected network it truly is. However, as this paper has argued, the successful deployment of relational intelligence requires more than just better algorithms; it demands a robust socio-technical framework that addresses architectural trade-offs, data governance, physical infrastructure, and algorithmic fairness.

We have explored the potential of GNNs to detect contagion pathways while highlighting the systemic dangers of model convergence and topological bias. We have also emphasized the need for sustainable computing and international policy coordination to manage the risks associated with this powerful technology. As the financial world becomes increasingly coupled and automated, the ability to decode and design the topology of risk will be the defining skill of the twenty-first-century financial engineer. By treating the market as a "system of systems," we can leverage the power of Graph Neural Networks to build a more resilient, transparent, and equitable future for the global financial ecosystem.

References

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

2. Arratia, A. (2014). Computational Finance: An Introductory Course with R. Atlantis Press.

3. Zhang, T. (2025). A Knowledge Graph-Enhanced Multimodal AI Framework for Intelligent Tax Data Integration and Compliance Enhancement. *Frontiers in Business and Finance*, *2*(02), 247-261.

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

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

6. 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).

7. Bronstein, M. M., et al. (2017). Geometric deep learning: Going beyond Euclidean data. IEEE Signal Processing Magazine, 34(4), 18-42.

8. 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).

9. Bruna, J., et al. (2014). Spectral networks and locally connected networks on graphs. International Conference on Learning Representations.

10. 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.

11. Defferrard, M., Bresson, X., & Vandergheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. Advances in Neural Information Processing Systems.

12. Tang, Y., Kojima, K., Gotoda, M., Nishikawa, S., Hayashi, S., Koike-Akino, T., ... & Klamkin, J. (2020, February). InP grating coupler design for vertical coupling of InP and silicon chips. In *Integrated Optics: Devices, Materials, and Technologies XXIV* (Vol. 11283, pp. 33-38). SPIE.

13. 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.

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

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

16. 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.

17. Tang, Y., Kojima, K., Gotoda, M., Nishikawa, S., Hayashi, S., Koike-Akino, T., ... & Klamkin, J. (2020). Design and Optimization of Shallow-Angle Grating Coupler for Vertical Emission from Indium Phosphide Devices.

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

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

20. Liu, T. (2026). A Comparative Study of Transformer-Based and Classical Models for Financial Time-Series Forecasting. *Journal of Risk and Financial Management*, *19*(3), 203.

21. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. Advances in Neural Information Processing Systems.

22. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. International Conference on Learning Representations.

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

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

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

26. 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).

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

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

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

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

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

32. Veličković, P., et al. (2018). Graph attention networks. International Conference on Learning Representations.

33. 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).

34. Wu, Z., et al. (2020). A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems, 32(1), 4-24.

35. Xu, K., et al. (2019). How powerful are graph neural networks? International Conference on Learning Representations.

36. Yang, Y., et al. (2020). Financial knowledge graph for risk management. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

37. Zhang, S., et al. (2019). Graph convolutional networks: A comprehensive review. Computational Intelligence Magazine, 14(4), 16-44.

38. Zhou, J., et al. (2020). Graph neural networks: A review of methods and applications. AI Open, 1, 57-81.