Graph Neural Networks for Modeling Systemic Risk in Financial Networks

Julianne S. Fairfield
Department of Systems Engineering, New Mexico Institute of Mining and Technology
j.fairfield@nmt.edu

Randall T. Vance
School of Computing and Informatics, University of Louisiana at Lafayette
rvance@louisiana.edu

Abstract

The global financial system is a highly coupled, non-linear network of institutions, markets, and sovereign entities where the connectivity between agents is as critical as their individual solvency. Traditional risk assessment frameworks, which largely rely on localized balance sheet analysis or linear correlation matrices, have proven insufficient for capturing the topological shifts and cascading contagion pathways inherent in modern systemic crises. This paper explores the integration of Graph Neural Networks (GNNs) as a foundational computational architecture for modeling systemic risk within these complex networks. By leveraging the message-passing paradigm, GNNs allow for the extraction of relational features that describe the structural importance and vulnerability of nodes within a non-Euclidean financial landscape. We conduct an extensive system-level investigation of GNN deployment, emphasizing the structural trade-offs between graph density and predictive robustness. The discussion encompasses the socio-technical infrastructures required for real-time risk monitoring, the data governance challenges associated with proprietary network disclosures, and the environmental sustainability of high-compute graph processing. Furthermore, we analyze the policy implications of GNN-driven risk assessment, addressing concerns regarding algorithmic fairness, model-driven market convergence, and the necessity for transparency in automated regulatory oversight. By synthesizing insights from systems engineering, graph theory, and financial policy, this research proposes a resilient framework for utilizing relational intelligence to safeguard global financial stability.

Keywords:
Graph Neural Networks, Systemic Risk, Financial Stability, Contagion Modeling, Algorithmic Governance, Socio-Technical Infrastructure, Relational Learning.

1. Introduction

The conceptualization of financial stability has undergone a radical transformation following the systemic failures of the early twenty-first century. It is now widely recognized that the primary driver of market collapse is often not the failure of a single entity, but the structural fragility of the network that binds institutions together. In this context, systemic risk is defined by the probability that a localized shock will propagate through interbank lending, shared asset holdings, and derivative contracts to cause a widespread cessation of financial functioning. Predicting these events requires an analytical framework capable of processing the intricate, multi-layered relationships that define the global economy. Graph Neural Networks (GNNs) represent the most significant architectural breakthrough in this regard, offering a means to model the very topology of financial contagion.

This paper situates the use of GNNs within the broader context of large-scale systems research. Unlike standard machine learning approaches that treat financial data as discrete temporal sequences, GNNs operate directly on the graph structures that constitute the financial world. We move beyond a narrow focus on algorithmic accuracy to examine the systemic requirements for deploying relational models in high-stakes, regulated environments. The transition to graph-based modeling necessitates a fundamental rethinking of data pipelines, moving from flat tabular structures to complex knowledge graphs that incorporate corporate ownership, supply chain dependencies, and cross-border capital flows.

The deployment of GNNs in the financial stability architecture introduces profound questions regarding the robustness of graph representations and the governance of the underlying data. As these models begin to influence the behavior of the markets they monitor, we must also consider the reflexive feedback loops they create. This introduction serves as the foundation for a deep inquiry into how relational intelligence can be harnessed to build a more resilient financial system, while simultaneously acknowledging the novel classes of systemic vulnerability that such advanced technologies introduce. Through a detailed analysis of infrastructure, sustainability, and policy, we aim to provide a roadmap for the future of automated systemic risk oversight.

2. Theoretical Evolution: From Institutional Solvency to Topological Resilience

Historically, the assessment of financial risk was an institution-centric endeavor. Regulators focused on capital adequacy ratios and liquidity buffers, operating under the assumption that if each individual node in the system were solvent, the system as a whole would remain stable. This "micro-prudential" approach, however, ignores the externalities generated by institutional interconnectedness. The realization that the structure of the network itself is a source of risk led to the development of "macro-prudential" policy, which seeks to identify systemically important financial institutions based on their centrality and connectivity. Early macro-prudential models utilized static network measures such as PageRank or eigenvector centrality to identify "too-interconnected-to-fail" actors.

While these early network models provided valuable insights, they were fundamentally limited by their inability to model complex, non-linear feature interactions and the dynamic nature of financial edges. Graph Neural Networks solve this by employing a message-passing mechanism where each node aggregates information from its neighbors, allowing for the learning of high-level representations that capture both local vulnerabilities and global structural patterns. This theoretical shift from institutional solvency to topological resilience represents a move toward understanding the financial system as a dynamic organism where the health of a node is inseparable from the health of its neighborhood.

However, the theoretical promise of GNNs is challenged by the problem of "edge uncertainty." Unlike a physical power grid, the edges in a financial graph—representing loans, contracts, or correlations—are often opaque and subject to rapid change. A robust theoretical framework for financial GNNs must therefore incorporate probabilistic edge weights and temporal dynamics. This section argues that the theoretical foundation of systemic risk modeling must move toward "relational inductive biases," where the model's assumptions are grounded in the known structural properties of financial contagion, such as the tendency for shocks to follow paths of high leverage and shared asset exposures.

3. Architectural Design and Systemic Trade-offs

The design of a GNN-based system for systemic risk monitoring involves a series of rigorous architectural trade-offs. The most significant of these is the trade-off between graph depth and representational stability. In theory, deeper GNNs, which aggregate information from many "hops" away in the graph, can capture long-range contagion pathways that span the entire global economy. In practice, however, deep GNNs frequently 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 a global perspective requires sophisticated interventions, such as the use of attention mechanisms to weight the importance of different neighbors and the integration of residual connections.

Computational scalability represents a second critical trade-off. Standard message-passing algorithms require the processing of the entire adjacency matrix, which, for a global financial network incorporating millions of corporate and institutional nodes, exceeds the capacity of conventional hardware. Systems researchers address this through graph sampling and subgraph-based training, where the model learns from representative portions of the network. While these techniques reduce the compute burden, they introduce the risk of missing "weak ties"—distant or infrequent connections that may nonetheless serve as critical conduits for systemic shock during a crisis. Choosing the right sampling strategy is thus as much a 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 that change in size or topology. Spatial GNNs, which define convolutions directly on the graph's physical neighbors, are more flexible and better suited for the dynamic environment of financial markets. For a system intended to monitor cross-market risk in real-time, a hybrid spatial-temporal architecture is often necessary. This section emphasizes that architectural choices in GNN design reflect specific assumptions about how risk propagates and which market participants are most relevant to systemic stability.

4. Data Governance and the Financial Knowledge Graph

The efficacy of a Graph Neural Network is entirely dependent on the quality and comprehensiveness of the underlying graph. In the financial sector, constructing this "knowledge graph" is a monumental data governance challenge. Unlike price data, which is readily available and standardized, relational data regarding interbank lending, derivative counterparties, and corporate ownership is often fragmented, proprietary, and highly sensitive. Building a graph that captures the true state of global interconnectedness requires navigating a complex web of legal disclosures and privacy regulations across multiple jurisdictions.

This fragmentation creates a risk of informational siloing, where only the largest financial institutions and regulators have the resources to build high-fidelity knowledge graphs. From a systems perspective, this raises concerns about market transparency and fairness. If a GNN-driven risk model at a major central bank identifies a systemic vulnerability using private data, the rest of the market remains unaware of 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 without compromising proprietary secrets.

Moreover, the governance of graph construction must account for "edge reliability." A connection based on a historical price correlation 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 the automated systems that govern financial stability.

5. Infrastructure, Deployment, and the Physicality of Graph Computing

The deployment of Graph Neural Networks for real-time financial monitoring requires a massive investment in specialized physical infrastructure. GNN operations involve irregular memory access patterns that 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 systemic risk assessment, firms are increasingly turning to specialized hardware such as Graph Processing Units or custom-designed ASICs optimized for sparse matrix operations. 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 an urgent concern. Training large-scale, dynamic GNNs 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 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" 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. 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 Topological Centrality

The concept of fairness in a GNN is uniquely complex because bias can emerge from the structure of the graph itself rather than just the node attributes. 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 legacy economic structures, 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 over-estimating the risk of emerging markets or innovative startups that exist on the periphery of the graph.

Correcting for this "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 GNNs is the phenomenon of "model convergence." If the majority of the world's systemically important financial institutions 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 an algorithmic monoculture 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 artificial relationships—such as through wash-trading to create fake 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 Biological and Cyber Resilience

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 replicate by incentivizing a more modular banking structure that prevents a single node failure from cascading through the entire network.

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 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 does not 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 an "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 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 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 we can 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 modeling systemic risk in financial networks. 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, 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. 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). AUBIQ: A generative AI-powered framework for automating business intelligence requirements in resource-constrained enterprises. *Frontiers in Business and Finance*, *2*(01), 66-86.

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

6. 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).

7. Bronstein, M. M., et al. (2017). Geometric deep learning: Going beyond Euclidean data.

IEEE Signal Processing Magazine, 34(4), 18-42.

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

9. 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.

10. 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.

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. 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. 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.

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

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

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

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

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

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

23. Qi, R. (2025, July). DecisionFlow for SMEs: A lightweight visual framework for multi-task joint prediction and anomaly detection. In *Proceedings of the 2025 International Conference on Economic Management and Big Data Application* (pp. 899-903).

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

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

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

27. 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).

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

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

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

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

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

33. 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).

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.