Graph Neural Networks for Systemic Risk Propagation Modeling in Multi-Layer Financial Markets

Brandon L. Fleming
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
blfleming@missouri.edu

Manav Katra
School of Computing, Clemson University, Clemson, SC, USA.
mbatra@clemson.edu

Damien D. Bichards
Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
damienwork@colostate.edu

Sameer M. Gokhale
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
sameermgokhale@uc.edu

Abstract

The increasing inter connectivity of modern financial markets has amplified the potential for systemic risk, where distress in one institution or market layer can cascade into widespread instability. Traditional models of financial contagion often rely on single-layer network abstractions or linearized propagation assumptions that fail to capture the complex, non-linear dependencies that characterize today's multi-asset, multi-jurisdictional financial architecture. This paper proposes and analyzes a framework based on Graph Neural Networks (GNNs) designed to model systemic risk propagation across multi-layer financial markets. We examine the structural trade-offs involved in constructing such models, including the representation of heterogeneous node and edge types, the aggregation of information across layers, and the integration of temporal dynamics. From a systems perspective, we discuss the data infrastructure required to support real-time risk monitoring, the computational sustainability of large-scale GNN deployments, and the governance challenges associated with model transparency, fairness, and regulatory oversight. The paper further explores how GNN-based risk models can enhance stress testing, early-warning systems, and macro-prudential policy design. By situating GNNs within the broader socio-technical context of financial regulation and infrastructure, we highlight open problems in robustness, interpretability, and cross-institutional coordination. Our analysis suggests that while GNNs offer significant advantages over classical network models, their deployment in high-stakes financial environments demands careful attention to data leakage, model calibration, and systemic fairness. We conclude by outlining a research agenda that bridges graph representation learning, financial network theory, and responsible AI governance.

Keywords

Graph Neural Networks, systemic risk, financial contagion, multi-layer networks, machine learning in finance, financial infrastructure, risk governance.

1. Introduction

The financial system is increasingly understood as a complex adaptive network in which institutions, markets, and instruments are linked through direct exposures, collateral chains, and common asset holdings. These linkages, while facilitating liquidity and risk sharing under normal conditions, also create pathways for distress to propagate in ways that are difficult to anticipate using linear or univariate models. The 2008 global financial crisis, the 2020 COVID-19 market turmoil, and the 2023 banking sector stress events have all underscored the need for more sophisticated computational tools capable of capturing non-linear contagion dynamics across multiple layers of financial activity. Graph Neural Networks (GNNs), a class of deep learning models designed to operate on graph-structured data, have emerged as a promising approach for modeling such systemic risk propagation precisely because they can learn relational dependencies and hierarchical patterns directly from network topology and node features [1, 2]. However, the application of GNNs to multi-layer financial markets introduces a host of architectural, infrastructural, and governance challenges that extend well beyond algorithmic optimization.

This paper provides a comprehensive systems-level analysis of how GNNs can be designed, deployed, and governed for systemic risk modeling in multi-layer financial markets. Rather than focusing solely on technical performance metrics, we emphasize the structural trade-offs inherent in representing financial networks as graphs, the data and computational infrastructure required to sustain such models in production, and the policy implications of using learned representations for regulatory decision-making. We draw on insights from network science, machine learning, financial economics, and socio-technical systems theory to argue that the successful integration of GNNs into financial risk management hinges not only on model accuracy but also on transparency, fairness, robustness to distributional shift, and alignment with existing regulatory frameworks. The paper is organized as follows. Section 2 reviews the foundational literature on financial contagion and graph-based modeling. Section 3 describes the architectural components of a GNN-based systemic risk model tailored to multi-layer markets. Section 4 examines the data infrastructure and governance considerations. Section 5 discusses deployment, sustainability, and model robustness. Section 6 explores policy implications and future directions. Section 7 concludes.

2. Background and Related Work

Systemic risk in financial networks has been studied extensively using classical network models, beginning with the seminal work of Allen and Gale on financial contagion through interbank lending networks [3]. Their model demonstrated that the structure of interbank exposures can either amplify or dampen the propagation of shocks, depending on the degree of connectivity and the presence of incomplete markets. Later contributions by Gai and Kapadia extended this analysis to show that even a single bank failure can trigger a cascade if the network is both densely connected and opaque [4]. These models, while influential, typically assume that shocks propagate linearly through balance sheet exposures and that the network is static and single-layered. In reality, financial markets consist of multiple overlapping layers, including equity, fixed income, derivatives, foreign exchange, and repo markets, each with distinct risk transmission mechanisms. Early warning systems that rely on aggregated single-layer measures often miss cross-layer contagion, particularly when stress in one market is transmitted through collateral rehypothecation or margin calls in another.

The development of multi-layer network models has provided a more realistic framework for understanding contagion. Researchers have shown that the multiplex nature of financial linkages can lead to unexpected amplification effects, where a shock in one layer triggers distress in another through common counterparties or overlapping portfolios [6]. However, these multi-layer models remain computationally expensive and often require hand-specified propagation rules that may not generalize to novel market conditions. Machine learning approaches, particularly those based on deep learning, offer the ability to learn propagation dynamics directly from historical data. Recent work has demonstrated the effectiveness of transformer-based models for financial time series forecasting [7] and attention-enhanced reinforcement learning for portfolio optimization [8]. Yet these methods typically treat each asset or institution independently, ignoring the relational structure that is central to contagion.

Graph Neural Networks bridge this gap by explicitly incorporating relational inductive biases into the learning architecture. Kipf and Welling introduced the Graph Convolutional Network (GCN), which aggregates feature information from a node's neighbors through a learned filter [9]. Since then, numerous variants have emerged, including Graph Attention Networks (GATs) that weigh neighbor contributions dynamically [10] and GraphSAGE that samples neighborhoods for scalability [11]. In financial applications, GNNs have been applied to fraud detection [12], credit risk assessment [13], and interbank network analysis [14]. However, these applications typically operate on a single graph layer and do not directly address the multi-layer nature of systemic risk. The challenge of modeling cross-layer dependencies remains an active area of research, with recent proposals incorporating hypergraphs or multiplex GNN architectures [15, 16]. Our work builds on these foundations by examining the system-level implications of deploying such models in real-world financial infrastructures.

3. Graph Neural Network Architecture for Multi-Layer Systemic Risk Modeling

The architecture of a GNN for systemic risk propagation in multi-layer markets must address several fundamental design decisions. First, the representation of nodes and edges across layers. In a multi-layer financial network, nodes may represent banks, hedge funds, central counterparties, or sovereign entities, each with distinct roles and risk profiles. Edges might represent direct interbank loans, derivative contracts, collateral links, or ownership structures. A homogeneous GNN that treats all nodes and edges identically would ignore these critical differences. To capture heterogeneity, a multiplex GNN can be constructed where each layer corresponds to a distinct financial market or instrument type, and nodes that appear in multiple layers are linked across layers through inter-layer edges [17]. The message-passing mechanism must then aggregate information both within a layer (intra-layer) and between layers (inter-layer). For example, a bank's exposure in the derivatives layer may affect its viability in the interbank lending layer, which in turn influences its ability to meet margin calls in the repo layer. A properly designed GNN can learn these cross-layer dependencies by propagating information through both types of connections.

A second critical architectural choice is the treatment of temporal dynamics. Systemic risk events unfold over time, with initial shocks cascading through the network in a sequence of failures, liquidity freezes, and asset fire sales. Static GNNs that operate on a single snapshot of the network cannot capture these dynamics. Temporal GNNs, such as those based on recurrent units or attention over time, allow the model to learn how network states evolve [18]. In our context, a temporal multiplex GNN could be trained on historical time series of balance sheet data, market prices, and transaction networks to predict the probability of default cascades or the propagation of stress indices. Recent work on volatility forecasting and early-warning market stress detection has shown that leakage-safe evaluation protocols are crucial in such settings, as walk-forward constraints prevent the model from accessing future information that would not be available in real time [19]. Similarly, attention-enhanced reinforcement learning approaches have been used to dynamically adjust portfolio weights in response to changing market conditions [8], illustrating the potential for adaptive risk management.

A third architectural dimension involves the aggregation function. Standard GNNs use mean, sum, or max pooling to combine neighbor features, but in financial networks, the impact of a neighbor may depend not only on its own characteristics but also on the size of the exposure relative to the node's capital. A learnable attention mechanism, similar to that used in GATs, can assign higher weights to neighbors that pose greater contagion risk [10]. However, attention weights must be calibrated to avoid overfitting to rare events. Regularization strategies, such as dropout on edges or graph spectral smoothing, can help improve generalizability. Furthermore, the model must be robust to missing or noisy data, which is common in financial networks where private bilateral exposures are often unobserved. Techniques such as variational graph autoencoders [20] or link prediction layers can be used to infer missing edges and estimate the latent network structure.

From a systems perspective, the architecture also determines the computational footprint and scalability. Multi-layer temporal GNNs can be extremely large, especially when applied to national or global financial systems with thousands of institutions and millions of edges. Scalable training requires mini-batch sampling strategies, as in GraphSAGE [11], and distributed computation across multiple GPUs. Moreover, the model must be updated frequently to reflect changing market conditions, which raises questions about the sustainability of continuous retraining. Energy consumption and latency constraints may limit the feasible update frequency, especially for real-time risk monitoring applications. These trade-offs must be explicitly considered during the design phase, as they affect not only model performance but also the operational cost and carbon footprint of the system.

4. Data Infrastructure and Governance

The effectiveness of a GNN-based systemic risk model is fundamentally constrained by the quality, granularity, and accessibility of financial data. Building a reliable multi-layer network requires data on balance sheets, interbank exposures, derivative positions, collateral pools, and market prices across multiple jurisdictions and asset classes. Much of this data is proprietary, held by individual institutions, central banks, or trade repositories, and is often subject to confidentiality agreements. Creating a unified graph that spans the entire financial system necessitates significant data sharing infrastructure, which in turn raises privacy and competitive concerns. Regulators, such as the European Systemic Risk Board and the U.S. Financial Stability Oversight Council, have taken steps to collect granular data through initiatives like the European Market Infrastructure Regulation (EMIR) and the Dodd-Frank Act, but these datasets are often fragmented and not harmonized across borders.

Data governance frameworks must address issues of consent, anonymization, and data lineage. For instance, if a central bank collects interbank loan data from all domestic banks, it must ensure that the aggregated graph does not reveal proprietary information about individual institutions. Differential privacy techniques can be applied to node or edge features before releasing them to researchers or model developers [21]. However, strong privacy guarantees may degrade the utility of the graph for contagion modeling, particularly if rare but systemically important edges are obscured. A second governance challenge concerns temporal consistency. Financial data are reported at different frequencies: balance sheets are quarterly, market prices are continuous, and transaction data may be reported daily. Aligning these disparate timestamps into a coherent graph sequence requires interpolation or imputation, which introduces uncertainty. The model must be designed to propagate this uncertainty through its predictions, and governance policies should mandate the disclosure of imputation methods and their impact on risk estimates.

Another layer of governance relates to model fairness and bias. Systemic risk models that rely on historical data may inadvertently encode discriminatory patterns, for example, by underestimating the risk of smaller or less-connected institutions, or by penalizing institutions from emerging markets due to data sparsity. These biases can have real-world consequences if the model is used to set capital requirements or trigger intervention measures. Fairness-aware GNN architectures have been proposed, incorporating adversarial debiasing or reweighting techniques to ensure that predictions are not systematically skewed along dimensions such as institution size, geography, or asset class [22]. However, fairness constraints often conflict with predictive accuracy, and regulators must decide on acceptable trade-offs. In the context of systemic risk, the primary goal is system-wide stability, which may justify some degree of uneven impact across institutions as long as the overall risk is mitigated. Nonetheless, transparency about these trade-offs is essential for maintaining trust in the regulatory process.

Finally, governance must address the interpretability of GNN models. Regulators and risk managers require explanations for why a model predicts a particular cascade scenario or flags an institution as high-risk. GNNs are inherently less interpretable than linear network models, but techniques such as graph saliency maps, counterfactual explanations, and prototype-based reasoning can provide insights into the model's decision process [23]. The challenge is that explanations themselves must be robust and not misleading. For example, a saliency map that highlights a particular edge as important may change drastically with a small perturbation in the data. Standards for model validation and explanation fidelity are still evolving, and the financial sector will need to collaborate with the machine learning research community to develop auditing frameworks that are both technically sound and practically feasible.

5. Deployment, Sustainability, and Robustness

Deploying a GNN-based systemic risk model in a live production environment introduces several practical challenges that extend beyond the research prototype. The model must operate within strict latency constraints, particularly if it is used for real-time monitoring or intraday risk assessment. While full cascade simulations may take minutes or hours, GNN inference can be sped up using model quantization, pruning, and specialized hardware. However, these optimizations can reduce model accuracy, necessitating careful benchmarking against a baseline. A second concern is model robustness to distributional shift. Financial markets evolve over time, with new instruments, regulatory changes, and structural breaks rendering historical data less representative. A model trained on data from a period of low volatility may fail to generalize to a crisis [24]. Continuous online learning, periodic retraining, and ensemble methods can help, but they also increase computational cost and introduce the risk of overfitting to recent events.

Sustainability, in the sense of long-term operational viability, depends on the cost of data acquisition, model training, and infrastructure maintenance. Large-scale GNNs require substantial computational resources, which may be a burden for smaller regulatory agencies or central banks in developing economies. Cloud-based solutions and federated learning approaches could democratize access, but they raise additional concerns about data sovereignty and cybersecurity. Moreover, the environmental impact of training deep learning models has become a growing concern, and financial institutions with net-zero commitments may face pressure to limit the carbon footprint of their risk analytics. Lightweight GNN architectures, such as those using simplified message-passing schemes or low-rank approximations, offer a path toward more sustainable deployment [25].

Robustness also encompasses adversarial resilience. An adversary with knowledge of the GNN model could potentially craft transactions or reporting data to manipulate the model's predictions, for example, to artificially lower a bank's systemic risk score. Adversarial training on perturbed graphs has been studied in the context of node classification and link prediction [26], but its application to systemic risk models is nascent. The financial system's high-stakes nature demands that models be tested against a range of adversarial scenarios, including data poisoning and evasion attacks. Red-teaming exercises, where independent researchers attempt to fool the model, can help identify vulnerabilities before deployment.

6. Policy Implications and Future Directions

The adoption of GNNs for systemic risk modeling has the potential to reshape macro-prudential policy by providing regulators with a dynamic, data-driven view of financial stability. One immediate application is in stress testing. Current regulatory stress tests, such as the U.S. Comprehensive Capital Analysis and Review (CCAR) and the European Banking Authority (EBA) stress tests, rely on large-scale linear or reduced-form models that do not account for network feedback effects. Incorporating GNN-based contagion modules could enhance the realism of these tests, capturing second-round effects such as fire sales and liquidity hoarding. However, regulators must be cautious about model complexity; a black-box GNN could be difficult to audit and challenge. A hybrid approach, where GNN outputs are used as inputs to more transparent structural models, may strike a useful balance.

Another policy area is the design of early-warning systems. Traditional early-warning indicators, such as credit spreads and volatility indices, often miss emerging systemic vulnerabilities that are latent in the network structure. GNNs can learn to extract early signals of contagion from the evolving graph, potentially providing earlier and more accurate warnings. Research on stress index construction using principal component analysis and arbitrage pricing theory has shown the value of multi-factor approaches [27]. Combining such indices with GNN-based network features could improve detection. However, as with any machine learning system, false positives and false negatives carry significant costs. Over-alerting may lead to unnecessary regulatory interventions, while false negatives may leave the system exposed. Calibrating the decision threshold requires a thorough understanding of the cost function, which is itself a policy decision.

International coordination is another critical dimension. Financial markets are globally interconnected, yet regulatory authorities operate within national boundaries. A GNN trained on domestic data alone may miss spillover effects from foreign markets. Cross-border data sharing agreements, such as the ones under the Financial Stability Board, can facilitate the construction of global financial networks, but political and legal barriers remain. Federated learning, where models are trained across jurisdictions without sharing raw data, offers a technological solution, but it introduces challenges of model convergence and accountability. Future research should explore how distributed GNNs can be designed to respect local privacy laws while still producing globally coherent risk estimates.

Finally, the ethical implications of using AI in financial oversight deserve sustained attention. Algorithmic decision-making in systemic risk can concentrate power in the hands of those who control the models, potentially leading to new forms of systemic inequality. Ensuring that model development is inclusive and transparent, with input from diverse stakeholders, is essential for public legitimacy. Academic researchers, central banks, and industry practitioners must collaborate to establish best practices for model validation, explainability, and fairness. The journey from theoretical GNN architectures to operational risk systems is long, but the potential rewards in terms of financial stability are substantial.

7. Conclusion

Graph Neural Networks present a powerful framework for modeling systemic risk propagation in multi-layer financial markets, capturing the relational and temporal dependencies that classical models miss. This paper has examined the architectural trade-offs involved in designing such models, the data and governance infrastructure required for their deployment, and the policy implications for financial stability. While GNNs offer clear advantages in representing complex contagion dynamics, their practical adoption is contingent on addressing challenges related to data privacy, model interpretability, computational sustainability, fairness, and adversarial robustness. The financial system's high-stakes environment demands that these models be developed with rigorous validation, transparent governance, and a commitment to serving the public interest. As research progresses, interdisciplinary collaboration between graph machine learning, financial economics, and regulatory policy will be essential to translate computational advances into resilient financial infrastructure.

References

1. Allen, F., & Gale, D. (2000). Financial contagion. Journal of Political Economy, 108(1), 1–33.
2. Gai, P., & Kapadia, S. (2010). Contagion in financial networks. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 466(2120), 2401–2423.
3. Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., ... & Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261.
4. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In International Conference on Learning Representations (ICLR).
5. Liu, T. (2026). PCA-APT Stress Index for Market Drawdowns.
6. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (NeurIPS), 30.
7. Xue, P., & Ye, Y. (2026). Attention-enhanced reinforcement learning for dynamic portfolio optimization. Intelligent Systems with Applications, 200622.
8. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. In International Conference on Learning Representations (ICLR).
9. Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. In Advances in Neural Information Processing Systems (NeurIPS), 30.

10. Liu, T. (2022, December). Financial Constraint'Impact on Firms' ESG Rating Based on Chinese Stock Market. In 2022 4th International Conference on Economic Management and Cultural Industry (ICEMCI 2022) (pp. 1085-1095). Atlantis Press.

11. Hu, L., & Shen, Y. (2026). A predictive analytics approach for forecasting global stock index returns using deep learning techniques. Decision Analytics Journal, 100685.

12. 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.

13. Liu, T. (2026). Interpretable Machine Learning for Volatility Forecasting Under Realistic Walk-Forward Constraints.

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

15. Eisenberg, L., & Noe, T. H. (2001). Systemic risk in financial systems. Management Science, 47(2), 236–249.

16. Brin, S., & Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. Computer Networks and ISDN Systems, 30(1–7), 107–117.

17. Caccioli, F., Shrestha, M., Moore, C., & Farmer, J. D. (2014). Stability analysis of financial contagion due to overlapping portfolios. Journal of Banking & Finance, 46, 233–245.

18. Poledna, S., Molina-Borboa, J. L., Martínez-Jaramillo, S., van der Leij, M., & Thurner, S. (2015). The multi-layer network nature of systemic risk and its implications for the costs of financial crises. Journal of Financial Stability, 20, 70–81.

19. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems, 32(1), 4–24.

20. Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2019). How powerful are graph neural networks? In International Conference on Learning Representations (ICLR).

21. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (pp. 308–318).

22. Chouldechova, A., & Roth, A. (2020). The frontiers of fairness in machine learning. In Machine Learning for Social and Behavioral Change (pp. 1–24). Cambridge University Press.

23. Ying, R., Bourgeois, D., You, J., Zitnik, M., & Leskovec, J. (2019). GNNExplainer: Generating explanations for graph neural networks. In Advances in Neural Information Processing Systems (NeurIPS), 32.

24. Gulrajani, I., & Lopez-Paz, D. (2020). In search of lost domain generalization. In International Conference on Learning Representations (ICLR).

25. Chen, J., Zhu, J., & Song, L. (2018). Stochastic training of graph convolutional networks with variance reduction. In International Conference on Machine Learning (ICML), 80, 941–949.

26. Zügner, D., Akbarnejad, A., & Günnemann, S. (2018). Adversarial attacks on neural networks for graph data. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 2847–2856).

27. Bollerslev, T., & Zhou, H. (2006). Volatility puzzles: A simple framework for gauging return-volatility regression. Journal of Econometrics, 131(1–2), 123–150.