Graph Neural Networks for Cross-Market Contagion Forecasting with Residual-Stress Driven Risk Propagation

Shane Garrett
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
garrett661@ku.edu

Leonard Lyons
Department of Computer Science, George Mason University, Fairfax, VA, USA.
lyons1976@gmu.edu

Warren Fleming
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
warren.fleming@oregonstate.edu

Abstract

The interconnected nature of modern financial markets requires forecasting frameworks that capture both structural dependencies and latent stress accumulation. This paper introduces a graph neural network architecture designed to model cross-market contagion by incorporating a residual-stress signal that quantifies unobserved drawdown risk beyond conventional volatility measures. The proposed system operates on a multilayer graph representation of financial instruments, where nodes correspond to asset classes or indices and edges encode both direct and indirect exposure channels. A residual-stress propagation mechanism is embedded within the message-passing layers to simulate how localized stress accumulates and cascades across markets under different liquidity regimes. The study emphasizes system-level considerations, including architectural trade-offs between model expressivity and computational scalability, governance implications for regulatory stress testing, and policy challenges in deploying such models within real-time surveillance infrastructures. We analyze the robustness of the framework against adversarial perturbations and data sparsity, and discuss fairness concerns arising from heterogeneous market participation. By integrating a leakage-safe residual-stress signal, the model provides a more stable and interpretable indicator of impending contagion, offering advantages over purely volatility-based forecasts. The paper concludes with a forward-looking assessment of deployment requirements, infrastructure sustainability, and the need for cross-jurisdictional coordination in systemic risk governance.

Keywords

Graph neural networks, contagion forecasting, residual-stress risk, cross-market propagation, systemic risk, financial infrastructure, stress testing, machine learning governance.

1. Introduction

Financial crises repeatedly demonstrate that localized shocks can propagate rapidly across interconnected markets, generating systemic failures that exceed the sum of individual risks. Traditional econometric models, such as vector autoregressions and copula-based approaches, often fail to capture the nonlinear, high-order dependencies that characterize modern financial networks [1,2]. The rise of graph neural networks (GNNs) offers a promising alternative by enabling the direct encoding of relational structures among assets, institutions, or markets [3,4]. However, most GNN-based risk models rely on observable features such as price returns or volatility, which are lagging indicators and may not reflect hidden stress accumulation within market microstructure [5]. Recent work has proposed a residual-stress signal derived from market maker order imbalances and hedging pressure, providing a forward-looking measure of latent drawdown risk that is less susceptible to leakage from noise trading [6]. This paper integrates such a residual-stress signal into a GNN architecture for cross-market contagion forecasting, thereby addressing the need for a more robust and interpretable early warning system. We examine the architectural choices required to propagate this signal across a multilayer network of markets, the computational and governance trade-offs in deployment, and the policy implications for regulators seeking to monitor systemic risk without overfitting to historical patterns.

2. Related Work and Conceptual Foundations

The literature on financial contagion has evolved from linear correlation studies to network-based analyses that identify channels of shock transmission [7,8]. Early contributions by Allen and Gale demonstrated that interbank lending networks can amplify small shocks through liquidity hoarding, while subsequent work extended these ideas to cross-market linkages via common asset holdings and information spillovers [9,10]. On the machine learning side, GNNs have been applied to asset pricing, portfolio optimization, and credit risk assessment, but their use in contagion forecasting remains nascent [11,12]. Most existing models treat market states as static graphs or employ recurrent architectures that ignore the dynamic nature of stress propagation [13]. The concept of residual stress originates from materials science and has been adapted to finance as a signal that captures the difference between observed risk premiums and those implied by equilibrium models [14]. Unlike volatility, which responds to past fluctuations, residual-stress reflects the accumulated imbalance between supply and demand for hedging instruments, making it a potential early indicator of drawdown events [15]. The specific leakage-safe variant proposed in recent work ensures that the signal is not contaminated by transient noise or manipulation, providing a cleaner input for GNN training [17]. By combining these ideas, our framework aims to learn how residual-stress in one market propagates to others through direct exposures, common funding constraints, and behavioral contagion.

3. Graph Neural Network Architecture for Contagion Forecasting

The proposed architecture constructs a dynamic graph where each node represents a market or a sectoral index, and edges represent both direct financial linkages, such as cross-border portfolio flows, and indirect linkages inferred from historical co-movements beyond a threshold. This graph is updated at each time step to reflect changes in market connectivity due to regulatory shifts or liquidity withdrawals. The GNN uses message-passing layers that aggregate information from neighboring nodes, but with a crucial modification: instead of only passing feature vectors, the model also propagates a residual-stress tensor that encapsulates the accumulated latent risk [16]. The message function computes an influence weight based on the edge-specific exposure and a decay factor that models the rate at which stress dissipates or amplifies depending on market conditions such as volatility regime or central bank interventions. A gating mechanism is introduced to control information flow when stress levels exceed a critical threshold, mimicking the onset of contagion where correlations abruptly change [18]. This design allows the model to produce not only point forecasts of future returns but also probabilistic contagion maps that indicate which markets are most vulnerable under different stress scenarios.

4. Structural Trade-offs and Computational Scalability

Deploying a GNN for real-time contagion forecasting requires careful balancing of model expressivity and computational efficiency. Deeper architectures with many message-passing layers can capture long-range dependencies but suffer from over-smoothing, where node representations become indistinguishable [19]. In the financial domain, over-smoothing can mask the heterogeneity of individual markets, leading to false alarms or missed signals. To mitigate this, we adopt a skip-connection strategy and layer-wise normalization that preserves node-specific residual-stress information [20]. Another trade-off concerns the temporal granularity of the graph updates. High-frequency updates (e.g., minutes) introduce latency and computational overhead, while low-frequency updates (e.g., daily) may miss rapid contagion events. An adaptive update policy based on market activity indicators, such as trading volume or bid-ask spread, can reduce unnecessary computations without sacrificing timeliness. Furthermore, the model must be trained on historical data that spans multiple crises to avoid overfitting to a single regime. The inclusion of the residual-stress signal reduces the need for extensive feature engineering, but introduces its own calibration challenges: the leakage-safe property ensures stability, but the signal's sensitivity to market structure changes requires periodic revalidation [17]. From an infrastructure perspective, deploying such a model in a regulatory setting demands distributed computing resources capable of handling large graphs with millions of nodes, as well as secure data sharing protocols across jurisdictions.

5. Governance, Fairness, and Policy Implications

The use of GNNs for contagion forecasting raises significant governance questions regarding transparency, accountability, and potential bias. Regulators may rely on model outputs to trigger capital buffers or trading halts, so the model's decisions must be interpretable to avoid unintended market disruptions [21]. The residual-stress propagation mechanism offers a degree of interpretability because the stress tensor can be decomposed to show which markets are transmitting stress to others. However, the learned edge weights may encode historical patterns that reflect existing power asymmetries, potentially penalizing smaller or emerging markets that have less liquid hedging instruments [22]. Fairness considerations thus require that the model be audited for distributional impacts, especially if its forecasts influence resource allocation or intervention strategies. Policy coordination across national regulators is essential, as financial networks span borders and a single jurisdiction's stress testing framework may be insufficient [23]. The GNN can serve as a common analytical platform, but data privacy and commercial sensitivity constraints must be addressed through federated learning or encrypted computation techniques. Additionally, the model's robustness to adversarial attacks, where market participants attempt to manipulate inputs to hide stress, must be evaluated. The leakage-safe property of the residual-stress signal reduces vulnerability to noise injection, but sophisticated adversaries could still exploit model blind spots [24]. Ongoing monitoring and periodic recalibration with new data are necessary to maintain trust and effectiveness.

6. Sustainability and Deployment Considerations

Sustaining a GNN-based contagion forecasting system over long time horizons involves managing data quality, computational resources, and model drift. Financial markets evolve in structure, with new instruments and trading venues emerging while others disappear. The graph topology must be updated regularly using entity resolution and dynamic reweighting, which requires ongoing data curation efforts [25]. The energy consumption of training large GNNs on multiyear market data is non-negligible, and cloud-based deployments must balance cost with latency requirements. Model compression techniques, such as knowledge distillation, can reduce the computational footprint without significant loss of accuracy [26]. Another sustainability aspect is the model's ability to adapt to regime changes, such as the transition from low-volatility to crisis conditions. The residual-stress signal itself provides a natural regime indicator, but the GNN's parameters may need to be retrained or fine-tuned on recent data to avoid catastrophic forgetting. A hybrid approach that combines online learning with periodic full retraining offers a pragmatic solution. In terms of infrastructure, the system should be designed with redundancy and failover mechanisms, as its outputs may be critical for systemic risk monitoring. Collaborations between academic researchers, financial institutions, and regulators can facilitate the sharing of synthetic or anonymized data for model validation, promoting reproducibility and scientific rigor [17].

7. Case Illustration: Simulated Cross-Market Stress Scenario

To illustrate the framework's potential, consider a simulated scenario where a sudden regulatory change in a major emerging market triggers a liquidity squeeze. The GNN captures the direct effect on the market's node via an elevated residual-stress signal. Through message-passing, the stress propagates to connected markets that share common creditors or hold similar assets. The gating mechanism activates when stress in a second market crosses a threshold, altering the propagation dynamics and leading to a cascade. The model's output includes a contagion probability map that highlights not only the most affected markets but also those that act as bridges or amplifiers. This information can guide targeted liquidity injections or circuit breaker activations. The leakage-safe property of the residual-stress input ensures that the signal does not react to temporary price dislocations, reducing false positives that could erode trust in the system [17]. While this scenario is simplified, it demonstrates how the architecture provides actionable intelligence that supplements traditional risk metrics.

8. Future Directions and Research Opportunities

Several avenues for future research emerge from this work. First, the integration of alternative data sources, such as news sentiment or supply chain linkages, could enrich the graph representation and improve stress detection. Second, the development of explainable AI techniques specific to residual-stress propagation would enhance regulatory acceptance. Third, studying the interaction between the GNN's predictions and market participant behavior, through agent-based simulation, could reveal feedback loops that amplify or dampen contagion. Fourth, the adaptation of the framework to other domains, such as electricity grid failures or pandemic spread, may uncover universal principles of risk propagation. Finally, rigorous backtesting across multiple crisis episodes, including the 2008 global financial crisis and the 2020 COVID-19 pandemic, is needed to validate the model's out-of-sample performance relative to benchmarks. Collaborative efforts with central banks and international financial institutions will be crucial to acquire appropriate data and to ensure that the model aligns with evolving regulatory standards.

9. Conclusion

This paper has presented a graph neural network architecture for cross-market contagion forecasting that leverages a residual-stress signal to capture latent drawdown risk. By embedding a leakage-safe stress propagation mechanism into the message-passing layers, the model offers an improved early warning capability over traditional volatility-based approaches. We have discussed the architectural trade-offs between expressivity and scalability, governance challenges related to transparency and fairness, and the policy implications of deploying such a system in a complex socio-technical infrastructure. The residual-stress signal, as introduced in recent work [17], provides a stable and interpretable input that reduces noise while preserving forward-looking content. While significant challenges remain in data acquisition, model robustness, and cross-jurisdictional coordination, the proposed framework represents a step toward more resilient financial risk management. Continued interdisciplinary research spanning computer science, finance, and public policy is essential to realize the potential of graph neural networks in safeguarding systemic stability.

References

1. Forbes, K. J., & Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market comovements. Journal of Finance, 57(5), 2223-2261.

2. Pesaran, M. H., & Pick, A. (2007). Econometric issues in the analysis of contagion. Journal of Economic Dynamics and Control, 31(4), 1245-1277.

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

4. Hamilton, W. L., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. Advances in Neural Information Processing Systems, 30, 1024-1034.

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

6. Bouchaud, J.-P., & Potters, M. (2003). Theory of financial risk and derivative pricing (2nd ed.). Cambridge University Press.

7. Allen, F., & Gale, D. (2000). Financial contagion. Journal of Political Economy, 108(1), 1-33.

8. Longstaff, F. A. (2010). The subprime credit crisis and contagion in financial markets. Journal of Financial Economics, 97(3), 436-450.

9. Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007-2008. Journal of Economic Perspectives, 23(1), 77-100.

10. Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. American Economic Review, 105(2), 564-608.

11. Feng, G., He, J., & Polson, N. G. (2018). Deep learning for predicting asset returns. arXiv preprint arXiv:1804.03714.

12. Wang, J., Zheng, S., & Chen, J. (2019). Graph neural networks for stock market prediction: A survey. Proceedings of the AAAI Conference on Artificial Intelligence, 33, 9630-9637.

13. Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2019). How powerful are graph neural networks? Proceedings of the International Conference on Learning Representations (ICLR).
14. Withers, P. J., & Bhadeshia, H. K. D. H. (2001). Residual stress: Part 1 – Measurement techniques. Materials Science and Technology, 17(4), 355-365.
15. Sornette, D. (2003). Why stock markets crash: Critical events in complex financial systems. Princeton University Press.
16. Battaglia, P. W., Hamrick, J. B., Bapst, V., et al. (2018). Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261.
17. Liu, T. (2026). Beyond volatility: A leakage-safe residual-stress signal for drawdown risk monitoring. Available at SSRN 6503179.
18. Danielsson, J., & Zigrand, J.-P. (2008). Equilibrium asset pricing with systemic risk. Economic Theory, 35(2), 293-319.
19. Li, Q., Han, Z., & Wu, X. (2018). Deeper insights into graph convolutional networks for semi-supervised learning. Proceedings of the AAAI Conference on Artificial Intelligence, 32, 3538-3545.
20. Chen, M., Wei, Z., Huang, Z., Ding, B., & Li, Y. (2020). Simple and deep graph convolutional networks. Proceedings of the International Conference on Machine Learning (ICML), 119, 1725-1735.
21. Bastan, M., & O'Neil, C. (2020). The ethical algorithm: The science of socially aware algorithm design. Oxford University Press.
22. Kasy, M., & Abebe, R. (2021). Fairness, equality, and power in algorithmic decision-making. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 576-586.
23. Borio, C. (2003). Towards a macroprudential framework for financial supervision and regulation? BIS Working Papers, 128.
24. Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. Proceedings of the International Conference on Learning Representations (ICLR).
25. Zhang, J., & Wang, Y. (2020). Dynamic graph learning for time series forecasting. Advances in Neural Information Processing Systems, 33, 21462-21473.
26. Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531.