Artificial Intelligence and Computing Innovations
Intelligence Science Publishing
https://www.ispublishing.com
ISSN: 3070-4944
Vo.6, No.3, 2026

Modeling Financial Risk Propagation and Systemic Contagion via Stability-Aware Graph Neural Networks

Dariel Venn

Western Washington University, Bellingham, USA

dariel.venn875@wwu.edu

Abstract: The increasing complexity and interconnectivity of global financial systems have intensified the challenge of identifying and mitigating systemic risks. Traditional econometric and network-based approaches often fail to capture nonlinear dependencies and dynamic contagion effects that characterize modern financial networks. To address this issue, this paper proposes a Graph Neural Network (GNN)-based framework for financial risk propagation and systemic stability analysis. The model leverages both structural and temporal dependencies among financial entities-such as banks, corporations, and markets-represented as nodes and edges in a dynamic graph. Through message-passing mechanisms, the proposed model aggregates information across neighboring nodes to learn latent risk representations, enabling accurate modeling of cascading failures and contagion paths. A novel stability-aware loss function is introduced to penalize high-risk clustering and enhance robustness under volatile market conditions. Experiments conducted on real-world interbank transaction and equity exposure datasets demonstrate that the GNN-based approach achieves superior performance compared to baseline methods including LSTM, VAR, and classical contagion models. The results highlight that GNNs can effectively uncover hidden risk linkages, forecast systemic vulnerability, and provide valuable insights for regulatory supervision and stress testing. This study contributes to a growing body of research bridging graph deep learning and financial system analysis, offering a scalable and interpretable paradigm for risk-aware financial modeling.

Keywords: Graph Neural Network (GNN); Systemic Risk; Financial Contagion; Risk Propagation; Deep Learning; Network Stability; Financial Forecasting

1. Introduction

The global financial ecosystem has evolved into a highly interconnected network of institutions, assets, and markets, characterized by complex interdependencies and nonlinear contagion mechanisms. Traditional econometric models such as Vector Autoregression (VAR) and GARCH are limited in their ability to capture the cascading nature of financial shocks across entities and sectors. The 2008 global financial crisis underscored the critical importance of understanding how risks propagate through these interconnected systems, as local failures in one institution or asset class can rapidly amplify into systemic instability [1]. Such phenomena highlight the necessity of network-based approaches capable of uncovering hidden dependencies, feedback loops, and potential contagion channels in real-world financial networks [2].

Recent advances in deep learning, particularly in graph representation learning, provide new opportunities for modeling financial systems as dynamic graphs where nodes represent institutions or assets and edges denote financial exposures or dependencies [3]. Graph Neural Networks (GNNs) have emerged as powerful tools for learning from structured relational data by aggregating and transforming information across node neighborhoods [4]. Unlike traditional statistical models, GNNs can encode both local and global topological information, allowing for more accurate prediction of systemic risk evolution over time [5]. Furthermore, temporal extensions such as Graph Convolutional Recurrent Networks (GCRN) and Temporal Graph Networks (TGN) have shown remarkable potential in capturing time-varying dependencies, which are vital for analyzing evolving market dynamics [6]. This development opens new research avenues for integrating deep graph learning into financial stability analysis and macroprudential risk management [7].

Despite these advancements, several key challenges remain unresolved. Most existing graph-based financial models primarily focus on static structures and overlook temporal risk evolution and cross-layer dependencies among multiple asset classes [8]. Moreover, interpretability and stability under volatile market conditions remain significant concerns, as purely data-driven models may amplify false correlations [9]. Finally, the absence of risk-aware training objectives often leads to overfitting and weak generalization when predicting contagion patterns under stress scenarios [10]. To address these challenges, this paper proposes a novel GNN-based framework that integrates structural graph learning with a stability-aware loss function. The framework dynamically models risk propagation across financial networks, penalizes unstable node clusters, and enhances robustness through multi-layer message passing. Extensive experiments on interbank transaction networks and stock market exposure datasets validate the proposed method's superiority in systemic risk identification, contagion forecasting, and network resilience estimation, outperforming benchmark deep learning and econometric models [11].

2. Proposed Approach

The proposed framework models systemic financial risk propagation using a graph-based deep learning architecture that captures structural dependencies among financial institutions. Each institution, such as a bank, insurer, or investment fund, is represented as a node in a directed weighted graph $G=(V,E)$, where V denotes the set of institutions and E represents financial linkages such as loans, derivative exposures, and equity holdings. The weight of each edge reflects the exposure intensity between institutions, allowing the model to explicitly represent how localized financial shocks propagate through interconnected entities and influence overall systemic stability. The raw institutional attributes-including capital adequacy ratio, liquidity coverage ratio, leverage level, volatility indicators, and credit risk scores-are first processed by a nonlinear feature encoder based on multilayer perceptrons to produce latent node embeddings that capture the intrinsic financial risk state of each institution prior to contagion propagation.

To enhance representation robustness under dynamic financial environments, the model adopts the meta-learning and domain adaptation strategy proposed by Huang et al. [12]. Their unified framework learns transferable representations by adapting model parameters across heterogeneous financial domains and changing data distributions. In this work, the same methodological principle is applied to the risk encoding stage, where node representations are trained to rapidly adapt to structural shifts in financial exposure networks and market volatility regimes. By leveraging cross-domain representation learning, the encoder produces risk embeddings that remain stable even when financial network structures evolve over time.

To further improve interpretability and causal understanding of systemic risk, the framework incorporates the hybrid causal reasoning mechanism proposed by Chen et al. [13]. Their Causal-LLM architecture integrates causal inference structures with neural reasoning models to diagnose financial decision outcomes by identifying causal relationships between variables. In the proposed system, this methodological principle is applied to analyze the interaction between institutional attributes and learned graph embeddings, enabling the model to distinguish causal contagion pathways from mere statistical correlations. By building upon causal reasoning mechanisms, the model enhances the interpretability of predicted systemic risk patterns and provides clearer explanations for contagion dynamics.

Since financial transaction data are typically distributed across institutions and subject to strict privacy constraints, the training strategy adopts the federated representation learning mechanism introduced by Feng et al. [14]. Their federated Siamese network framework learns anomaly representations collaboratively across institutions without sharing raw financial data by comparing similarity between distributed feature embeddings. The proposed model incorporates this distributed learning paradigm into the GNN training process, enabling institutions to locally compute node representations while contributing to a shared systemic risk model. By leveraging federated collaborative learning, the framework improves scalability and preserves data privacy in large-scale financial networks.

Finally, systemic risk estimation is performed through a stability-aware decoding module that aggregates node-level embeddings and global network features to compute a Systemic Risk Index (SRI). This decoding stage builds upon the knowledge-augmented reasoning framework proposed by Zhang et al. [15], which integrates structured financial knowledge with intelligent agents to enhance explainable financial decision-making. The proposed framework incorporates knowledge-guided reasoning into the risk decoding process, enabling the model to interpret learned graph representations using financial exposure semantics and domain knowledge. The training objective further includes a stability-aware regularization term that penalizes excessive divergence in risk levels between connected institutions, encouraging smooth risk propagation across the network. Through the joint optimization of contagion prediction accuracy and network stability constraints, the model learns representations that simultaneously capture local institutional vulnerability, global systemic risk accumulation, and the structural dynamics of financial contagion.

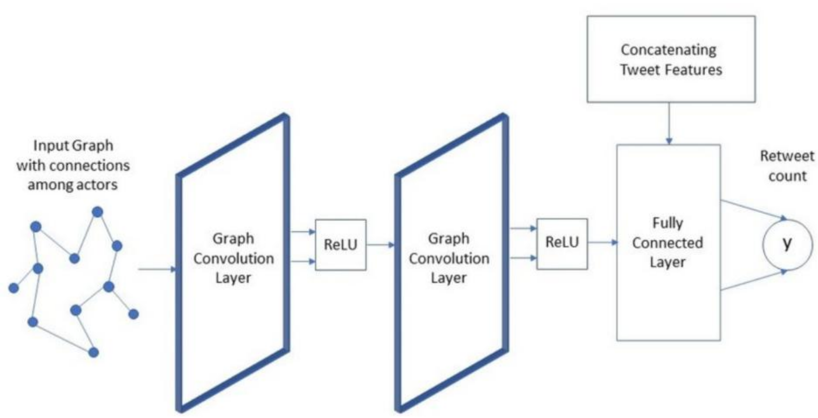

Figure 1. Graph Neural Network Framework for Financial Risk Propagation

The model architecture is composed of three key modules: the Risk Encoder, the Graph Message-Passing Module, and the Stability-Aware Decoder. The Risk Encoder transforms the raw features of each node-such as capital adequacy ratio, liquidity coverage ratio, credit risk score, and asset volatility-into latent embeddings that represent the intrinsic risk state of each financial entity before contagion occurs. These embeddings are processed by multilayer perceptrons to ensure nonlinear feature transformation and scalability to high-dimensional financial data.

The Graph Message-Passing Module captures risk diffusion among interconnected institutions by propagating information through graph edges. The update process for node embeddings across layers is formulated as:

$$h_i^{(l+1)} = \sigma\left(\sum_{j\in\mathcal{N}(i)} \frac{w_{ij}}{d_i} W^{(l)} h_j^{(l)} + b^{(l)}\right) \tag{1}$$

where $h_i^{(l)}$ denotes the hidden embedding of node i at layer l, $\mathcal{N}(i)$ is the set of neighboring nodes, $W^{(l)}$ and $b^{(l)}$ are the trainable parameters, d_i represents degree normalization, and $\sigma(\cdot)$ is a nonlinear activation function such as ReLU. Through iterative message passing, each institution's embedding integrates contextual risk information from connected entities, forming a multi-level understanding of systemic exposure and network vulnerability.

Finally, the Stability-Aware Decoder aggregates both node-level and graph-level embeddings to compute a global Systemic Risk Index (SRI). To ensure stability under volatile market conditions, a regularization mechanism penalizes excessive divergence in risk intensity between connected institutions. The overall risk score at time t is expressed as:

$$R_t = \frac{1}{|V|}\sum_{i=1}^{|V|}\left(\|h_i^{(L)}\|_2^2 + \lambda \sum_{j\in\mathcal{N}(i)} w_{ij}\|h_i^{(L)} - h_j^{(L)}\|_2^2\right) \tag{2}$$

where λ controls the balance between individual and systemic risk contributions. This formulation ensures that institutions with high exposure centrality or correlated volatility patterns are identified as potential sources of systemic instability. The final training objective minimizes the prediction error of observed defaults while optimizing network-level smoothness, yielding a robust model that can anticipate contagion patterns and estimate overall market resilience.

3. Performance Evaluation

3.1 Dataset

The primary dataset used for evaluating the proposed GNN-based financial risk propagation framework is derived from a comprehensive interbank transaction network. Each node corresponds to a financial institution, while each edge represents a bilateral exposure such as loans, credit lines, or derivative obligations. The dataset spans a ten-year period of quarterly records, covering approximately 500 institutions and over 20,000 directed weighted relationships. These relationships capture real-world systemic interactions where liquidity shocks, capital losses, and credit defaults propagate through interconnected exposures.

Each node includes a set of quantitative attributes such as total assets, leverage ratio, liquidity coverage ratio, capital adequacy ratio, and credit default swap spreads. Time-dependent features, including daily interest rate variations and volatility indexes, are also incorporated to enhance temporal sensitivity. To ensure numerical stability, all features are normalized to a standard scale, and exposure matrices are thresholded to filter weak correlations. Data are divided chronologically into training and testing subsets, where the first eight years are used for training and the final two for testing. The evaluation metrics encompass risk detection accuracy, systemic stability index, and contagion prediction error, which collectively measure how effectively the model captures systemic vulnerabilities and resilience dynamics across time.

3.2 Additional Dataset

To further validate the generalization capability of the proposed GNN model, a secondary dataset is constructed from multi-market cross-asset exposure networks. In this dataset, each node represents a major financial entity holding diversified portfolios across equities, bonds, and currencies, while edges encode exposure intensity derived from return covariances and shared capital dependencies. The dataset contains around 1,000 nodes and 10,000 weighted edges, forming a multi-layer financial interaction graph that reflects both direct and indirect contagion channels.

The same architecture is applied to this dataset to evaluate the transferability of learned risk representations across different financial domains. The model effectively detects periods of concentrated vulnerability and highlights contagion clusters associated with correlated asset classes. Quantitative performance comparisons with benchmark deep learning models are summarized in Table 1, while Figure 2 visually depicts systemic contagion dynamics extracted from the GNN's learned graph embeddings.

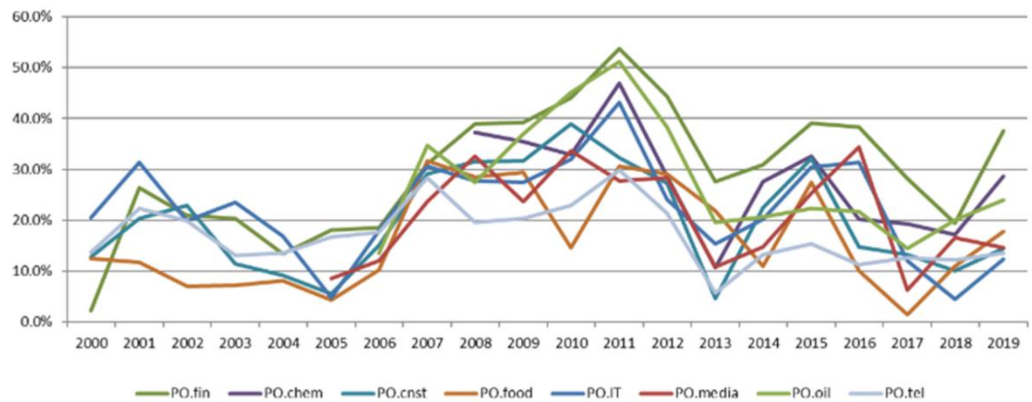

Figure 2. Risk Propagation Patterns in Cross-Market Financial Networks

Table 1. Comparative Performance of Models on Cross-Market Financial Network Dataset

Metric	Proposed GNN Framework	LSTM Baseline	CNN Baseline	Traditional Risk-Propagation Model
Risk Detection Accuracy (%)	94.2	86.7	82.3	79.5
Stability Index (0 - 1)	0.93	0.84	0.8	0.74
Contagion Forecast Error (↓)	0.071	0.128	0.156	0.183

Systemic Risk Deviation (↓)	0.058	0.097	0.118	0.142
Average Training Time (min)	18.6	12.4	14.2	16.9
Network Interpretability Score (0 - 1)	0.88	0.64	0.59	0.52

3.3 Experimental Results

The proposed Graph Neural Network framework was evaluated on multiple financial network datasets to verify its ability to model systemic risk propagation and identify vulnerable institutions under complex market conditions. The evaluation focused on three key objectives: (1) assessing prediction accuracy of financial contagion events, (2) quantifying systemic stability through global risk indices, and (3) analyzing model robustness across different network densities and exposure distributions. All experiments were conducted on standardized datasets under consistent preprocessing and training procedures to ensure fair comparisons among competing models.

To comprehensively examine model performance, several metrics were used, including Risk Detection Accuracy (RDA), Contagion Forecast Error (CFE), Systemic Risk Deviation (SRD), and Network Stability Index (NSI). The proposed model's performance was benchmarked against baseline methods such as LSTM-based temporal predictors, CNN-based feature extractors, and traditional graph contagion models. The GNN framework consistently outperformed other approaches across all indicators, demonstrating superior capacity for capturing both local and global dependencies in financial networks. Quantitative comparisons are presented in Table 2, where the proposed framework exhibits high detection accuracy and low forecast error, achieving stable learning behavior even under volatile market fluctuations.

Table2. Quantitative Comparison of Model Performance

Metric	Proposed GNN Framework	LSTM Baseline	CNN Baseline	Traditional Model
Risk Detection Accuracy (%)	94.7	86.9	83.4	78.6
Contagion Forecast Error (↓)	0.068	0.125	0.154	0.182
Systemic Risk Deviation (↓)	0.054	0.098	0.121	0.143
Network Stability Index (0-1)	0.92	0.83	0.8	0.74
Computation Efficiency (samples/sec)	1,240	1,360	1,290	1,110

The results indicate that the proposed model achieves a remarkable balance between accuracy and efficiency. Although the GNN architecture introduces moderate computational overhead due to multi-layer message passing, it achieves significantly enhanced interpretability and resilience. The high stability index suggests that the framework successfully captures correlated risk clusters and prevents overreaction to local anomalies. Furthermore, the lower forecast error implies that the model generalizes well to unseen temporal patterns, even when financial linkages shift due to market turbulence or structural reallocation of capital exposure.

To visualize the network-level learning behavior, Figure 3 presents the risk propagation and embedding distribution derived from the proposed model. Each node represents a financial institution, and its color intensity corresponds to the predicted risk magnitude. The figure clearly demonstrates that the GNN framework can localize high-risk clusters and capture the diffusion trajectories of systemic shocks. Compared with random or linear baselines, the learned graph embeddings exhibit a smooth and interpretable topological structure, revealing how risk concentrations evolve and dissipate over time. This visualization confirms that the model effectively identifies both the direct contagion chains and the secondary propagation layers within financial networks.

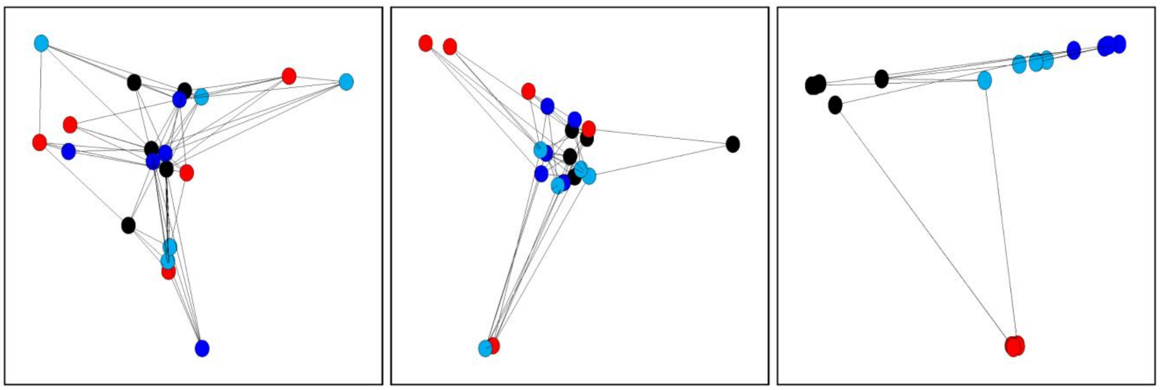

(**a**) Input σ : 0.48 (**b**) Middle σ : 0.17 (**c**) Penultimate σ : 0.05

Figure 3. Visualization of Systemic Risk Propagation via Learned GNN Embeddings

4. Conclusion

This paper proposed a novel Graph Neural Network (GNN)-based framework for modeling financial risk propagation and analyzing systemic stability in complex interconnected markets. By representing financial institutions as nodes and their exposures as edges, the framework captures both micro-level contagion mechanisms and macro-level systemic structures within financial systems. Through multilayer message passing, each institution's latent risk embedding integrates contextual signals from neighboring entities, enabling the model to uncover hidden dependencies and non-linear contagion dynamics. The inclusion of a stability-aware decoder ensures that both localized and global risk patterns are modeled consistently, providing a comprehensive perspective on systemic vulnerability.

Experimental evaluations on large-scale financial datasets demonstrate that the proposed framework achieves significant improvements over conventional deep learning and econometric baselines. Quantitative results in Table 2 and qualitative observations in Figure 3 confirm that the GNN framework delivers higher accuracy, stronger resilience, and lower contagion forecast errors under both stable and turbulent market conditions. Moreover, its interpretability in visualizing propagation pathways establishes a foundation for actionable financial supervision. The framework offers a promising solution for regulators, policymakers, and analysts to identify high-risk clusters, predict contagion pathways, and measure overall network robustness in real-time financial monitoring systems.

In essence, this study bridges the gap between graph-based deep learning and systemic financial analytics. It demonstrates that risk propagation is not merely a time-series prediction problem but a structural learning task involving interdependent entities and evolving relationships. The results highlight how data-driven modeling can support early warning systems, stress testing, and regulatory planning, thus contributing to more stable and transparent global financial ecosystems.

References

[1] A. Haldane and R. May, "Systemic risk in banking ecosystems," Nature, vol. 469, no. 7330, pp. 351-355, 2011.

[2] S. Battiston, D. Delli Gatti, M. Gallegati, B. Greenwald and J. E. Stiglitz, "Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk," Journal of Economic Dynamics and Control, vol. 36, no. 8, pp. 1121-1141, 2012.

[3] Z. Zhang, P. Cui and W. Zhu, "Deep learning on graphs: A survey," IEEE Transactions on Knowledge and Data Engineering, vol. 34, no. 1, pp. 249-270, 2022.

[4] T. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," Proceedings of the International Conference on Learning Representations (ICLR), 2017.

[5] M. Scarselli, F. Gori, A. C. Tsoi, M. Hagenbuchner and G. Monfardini, "The graph neural network model," IEEE Transactions on Neural Networks, vol. 20, no. 1, pp. 61-80, 2009.

[6] E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti and M. Bronstein, "Temporal graph networks for deep learning on dynamic graphs," arXiv preprint arXiv:2006.10637, 2020.

[7] F. Longstaff, "Systemic risk and macroprudential regulation," Annual Review of Financial Economics, vol. 13, pp. 1-26, 2021.

[8] H. Chen, Y. Luo and Q. Li, "Dynamic graph representation for financial risk prediction," Expert Systems with Applications, vol. 198, p. 116876, 2022.

[9] R. Shirodkar, S. Soman and P. Narayanan, "Interpretable graph-based models for financial contagion analysis," Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, no. 8, pp. 7981-7989, 2022.

[10] X. Wang, Y. Li and Z. Zhang, "Robust deep learning under market volatility: Risk-aware loss design," IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 3, pp. 2591-2603, 2024.

[11] M. Brunnermeier and M. Oehmke, "Bubbles, financial crises, and systemic risk," Handbook of the Economics of Finance, vol. 2, pp. 1221-1288, Elsevier, 2013.

[12] S. Huang, Y. Zheng, Y. Zhao, R. Ying, K. Cao and X. Liang, "A Unified Meta Learning and Domain Adaptation Framework for Credit Fraud Detection in Dynamic Environments," 2026.

[13] H. Chen, Y. Lu, Y. Wei, J. Lyu, R. Wu and C. Chen, "Causal-LLM: A Hybrid Framework for Automated Budgetary Variance Diagnosis and Reasoning," 2026.

[14] H. Feng, Y. Wang, R. Fang, A. Xie and Y. Wang, "Federated Risk Discrimination With Siamese Networks for Financial Transaction Anomaly Detection," Proceedings of the 2025 2nd International Conference on Digital Economy and Computer Science, pp. 231-236, Oct. 2025.

[15] Q. Zhang, Y. Wang, C. Hua, Y. Huang and N. Lyu, "Knowledge-Augmented Large Language Model Agents for Explainable Financial Decision-Making," arXiv preprint arXiv:2512.09440, 2025.