BMIJ
ISSN: 2148-2586

bmij (2026) 14 (1):32-47
doi: https://doi.org/10.15295/bmij.v14i1.2682
Research Article

Modelling financial contagion in complex equity networks: A decision support framework based on graph neural networks

Karmaşık hisse senedi ağlarında finansal bulaşıcılığın modellenmesi: Grafik sinir ağlarına dayalı karar destek çerçevesi

Cevher Özden[1]

[1] Assist. Prof., Cukurova University, Faculty of Arts and Sciences, Department of Computer Sciences, Adana, Türkiye, cozden@cu.edu.tr

ORCID: 0000-0002-8445-4629

Abstract

Financial markets have become increasingly interconnected, amplifying the speed and magnitude of contagion effects during periods of uncertainty. This study proposes an integrated network-based and machine learning framework to examine financial contagion and systemic vulnerability in the Borsa Istanbul equity market. Using daily log-returns of 106 stocks over the period 2023–2025, a correlation-based financial graph network is constructed, and the structural backbone is extracted using the Minimum Spanning Tree approach. The resulting topology reveals a hub-dominated structure in which systemic importance is concentrated within a limited number of assets. A Graph Neural Network is employed to classify future volatility regimes. The model achieves an out-of-sample classification accuracy of 65%, indicating that network-aware learning captures predictive signals not available in isolated time-series models. The findings demonstrate that financial contagion is both structurally embedded and predictively exploitable when markets are modelled as complex networks. This study addresses a critical gap in the literature by extending predictive deep graph learning to the highly volatile dynamics of an emerging market, advancing beyond the static topologies and traditional econometric models prevalent in prior local research. The proposed framework offers practical implications for portfolio managers and macroprudential regulators by enabling the early identification of real-sector contagion hubs.

Keywords: Financial Contagion, Financial Networks, Decision Support Systems

Jel Codes: G14, G15, G32

Submitted: 31/12/2025

Revised: 27/02/2026

Accepted: 9/03/2026

Online Published: 25/03/2026

Öz

Finansal piyasalar giderek daha fazla birbirine bağlı hâle gelmiş olup, belirsizlik dönemlerinde bulaşma etkilerinin hızını ve büyüklüğünü artırmaktadır. Bu çalışma, Borsa İstanbul hisse senedi piyasasında finansal bulaşmayı ve sistemik kırılganlığı incelemek amacıyla bütünleşik bir ağ temelli ve makine öğrenmesi çerçevesi önermektedir. 2023–2025 dönemini kapsayan 106 hisse senedine ait günlük logaritmik getiriler kullanılarak, korelasyona dayalı bir finansal grafik ağ oluşturulmuş ve Minimum Örten Ağaç yaklaşımı kullanılarak yapısal omurgası çıkarılmıştır. Ortaya çıkan topoloji, sistemik önemin sınırlı sayıda varlıkta yoğunlaştığı, merkez ağırlıklı bir yapıyı ortaya koymaktadır. Gelecekteki volatilite rejimlerini sınıflandırmak amacıyla bir Grafik Sinir Ağı kullanılmıştır. Model, örneklem dışı dönemde %65 sınıflandırma doğruluğu elde etmiş olup, bu sonuç ağ farkındalığına sahip öğrenmenin, izole zaman serisi modellerinde mevcut olmayan öngörücü sinyalleri yakalayabildiğini göstermektedir. Bulgular, finansal bulaşmanın piyasalar karmaşık ağlar olarak modellendiğinde hem yapısal olarak içkin olduğunu hem de öngörülebilir biçimde kullanılabilir nitelik taşıdığını ortaya koymaktadır. Bu çalışma, tahmine dayalı derin çizge (grafik) öğrenimini gelişmekte olan bir piyasanın yüksek volatilite dinamiklerine uygulayarak; önceki yerel araştırmalarda hakim olan statik topolojilerin ve geleneksel ekonometrik modellerin ötesine geçmekte ve literatürdeki kritik bir boşluğu doldurmaktadır. Önerilen çerçeve, reel sektördeki bulaşıcılık merkezlerinin erken tespit edilmesini mümkün kılarak portföy yöneticileri ve makro ihtiyati regülatörler için pratik çıkarımlar sunmaktadır.

Anahtar Kelimeler: Finansal Bulaşma, Finansal Ağlar, Karar Destek Sistemleri

JEL Kodları: G14, G15, G32

Citation: Ozden, C., Modelling financial contagion in complex equity networks: A decision support framework based on graph neural networks, bmij (2026) 14 (1): 32-47, doi: https://doi.org/10.15295/bmij.v14i1.2682

Introduction

The global financial ecosystem has become increasingly intricate, adaptive, and interconnected, making it highly vulnerable to cascading failures and systemic breakdowns. Events such as the 2008 global financial crisis, the Eurozone sovereign debt crisis, and the COVID-19 pandemic have starkly demonstrated how local shocks, once isolated, can now rapidly spill over through hidden channels of interdependence, resulting in widespread instability (Allen & Babus, 2009; Battiston et al., 2012; Zhang et al., 2020). This phenomenon, commonly referred to as *financial contagion*, has prompted a growing body of research seeking to understand, measure, and forecast systemic risk across increasingly complex financial networks.

While conventional risk models often assume asset independence or employ static measures of correlation, these approaches have repeatedly been challenged by the empirical realities of modern markets. Correlation structures are not only time-varying but also prone to sharp non-linear shifts during periods of distress (Forbes & Rigobon, 2002). Moreover, systemic risk is not merely the aggregate of individual exposures; it emerges from the *structure* of relationships among institutions, instruments, and markets (Acemoglu et al., 2015). As such, network-based approaches have gained significant traction as tools for representing and analysing financial systems. In these models, financial entities are represented as nodes, while links signify dependencies such as cross-holdings, common exposures, or price co-movements. This framework enables researchers to explore the role of central actors (Billio et al., 2012), contagion channels (Huang et al., 2013), and structural fragility (Diebold & Yılmaz, 2014) with unprecedented granularity.

Despite their descriptive richness, however, many existing network-based studies remain limited in their capacity to *predict* systemic risk. The overwhelming majority of these works are retrospective, employing ex post measures of connectedness to explain contagion after it has occurred. Moreover, many rely on bivariate or linear dependency structures, which may fail to capture the full complexity of risk transmission in highly non-linear financial environments (Kumar & Deo, 2012; Demirer et al., 2018). As a result, there is a critical need for methodologies that are not only structurally informed but also forward-looking and computationally scalable.

Recent advances in machine learning, particularly Graph Neural Networks (GNNs), present a promising frontier in this regard. GNNs are designed to learn from graph-structured data by incorporating both node-level attributes and graph topology. This allows them to extract rich, latent representations that are especially well-suited for modelling relational systems such as financial markets (Balmaseda et al., 2023). Empirical studies have begun to show that GNNs outperform traditional machine learning models in identifying systemically important nodes, estimating default probabilities, and forecasting regime shifts (Wei et al., 2025; Battiston et al., 2016). However, the application of GNNs to financial contagion forecasting is still in its early stages, particularly in emerging markets, where market structures and dynamics may differ markedly from those in developed economies.

Unlike developed markets characterised by deep liquidity, broad institutional participation, and relatively stable macroeconomic environments, emerging markets tend to exhibit higher volatility persistence, lower market depth, and greater sensitivity to domestic and external shocks (Tabash et al., 2024; Ahmed et al., 2018). These structural features imply that risk transmission mechanisms in emerging economies may operate through more fragile and rapidly shifting network topologies. The Turkish equity market, as represented by the BIST 100 index, reflects many of these characteristics, including pronounced regime shifts, currency-driven spillovers, and heightened exposure to global risk sentiment (Altinbas, 2025). Consequently, modelling systemic interactions in such an environment requires analytical frameworks capable of capturing non-linear dependencies and evolving relational structures, thereby providing a strong theoretical justification for applying graph-based deep learning methods in this context.

This study seeks to address these methodological and empirical gaps by proposing a hybrid approach that combines classical network construction techniques with modern deep learning models. Focusing on the Turkish equity market (BIST 100), time-varying correlation networks are constructed from daily log-returns, and the market's structural backbone is extracted using a Minimum Spanning Tree (MST) approach. Node-level centrality measures and statistical descriptors are used as features in a GNN classifier trained to predict asset-level volatility regimes. Unlike purely descriptive studies, the proposed model employs a temporal train-test split, preserving the causal structure of financial data and enabling genuine out-of-sample evaluation.

This research makes three key contributions to the literature:

1. Methodological integration: The study offers a structured pipeline that integrates MST-based network pruning with supervised learning via GNNs, providing a scalable, theoretically grounded approach to systemic risk prediction.

2. Empirical validation in an emerging market: While most GNN-based financial studies focus on developed economies, the current application to the Turkish market adds geographical and economic diversity to the literature on financial contagion.

3. Forward-looking risk classification: By training the model solely on past data and validating its predictive capacity on future outcomes, this approach is well aligned with real-world decision-making constraints, contributing to the development of proactive, MIS-supported early warning systems.

In doing so, this study contributes to a growing body of interdisciplinary research at the intersection of finance, network science, and artificial intelligence. It provides both methodological insight and practical value for regulators, portfolio managers, and technology-driven risk analysts.

Literature review

Financial contagion is often studied via network models that capture interconnections among institutions or markets. In such models, nodes (banks, firms, assets) are linked by exposures or correlations, and standard network metrics (density, clustering, centrality, etc.) quantify how shocks can spread (Deng et al., 2025). For example, Mishra & Mishra (2021) construct dynamic stock-return networks for Hong Kong and find that COVID-19 generated significantly higher network density and clustering (in partial-correlation networks) than previous crises. Multilayer networks have also been used: Demirer et al. (2018) build a three-layer network of returns, volatility, and tail-risk for global systemically important banks and show that the "extreme risk" layer exhibits the strongest spillovers during crises. Other studies link different markets or sectors. Zhang et al. (2020) examine a global network of stock, bond and foreign-exchange markets (14 countries, 2000–2021) and report high connectedness during COVID-19 – albeit systemic risk remained below 2008 levels – with bond markets most exposed and FX markets acting as central contagion channels. Similarly, Malkina (2024) documents contagion from global equities to 22 commodity futures during 2020–2023, finding that precious metals (especially gold) were most vulnerable, while some agricultural commodities and Brent crude showed relative resilience. In the banking and credit context, Qian et al. (2025) use a two-layer network of banks and nonbank firms in China (2013–2022) and uncover strong tail-risk spillovers and asymmetric contagion: nonbank firms tend to contribute more to systemic risk than banks.

In recent years, Graph Neural Network (GNN)–based models have attracted growing attention in stock markets and other areas of finance. Studies such as Wang et al. (2022) show that financial data are often naturally represented as heterogeneous, dynamic graphs, where nodes correspond to assets and edges capture their relationships. According to this line of research, GNNs are particularly well-suited to modelling complex network structures, enabling them to perform effectively across various financial tasks (Bursa, 2025). For instance, when each stock is treated as a node, and the connections between stocks represent their underlying dependencies, the model can learn how shocks or information propagate through the market.

Recent work has applied graph-based machine learning to the study of systemic risk. Graph Neural Networks (GNNs) explicitly use network structure and node/edge features to predict contagion and risk. Balmaseda et al. (2023) demonstrate that GNNs substantially outperform traditional classifiers in identifying systemically important institutions: on two financial networks, they report 15–94% relative gains in prediction accuracy (MCC metric) by incorporating graph structure. Temporal graph models have also been developed. Wei et al. (2025) propose a spatio-temporal GNN for a bank–loan guarantee network and report an AUC of 88.3% for default prediction, which is significantly higher than that of baseline methods. They further show that targeting the top 1% highest-risk nodes (as flagged by their model) could reduce overall systemic exposure by ~25%. On the theoretical side, Battiston et al. (2016) extend classical systemic-risk metrics to graph data. They use permutation-equivariant neural architectures (a class that includes GNNs) to approximate clearing-based risk measures in model financial networks, finding that these graph-aware models reduce overfitting and better capture contagion pathways than standard networks.

A growing body of recent research highlights the usefulness of GNNs in tasks such as volatility forecasting and systemic risk analysis. For example, Kumar et al. (2024) proposed a Temporal Graph Attention Network (TemporalGAT) that combines GCN and GAT layers to predict the volatility of eight major global indices. Their findings indicate that this hybrid architecture outperforms traditional

GARCH models in capturing complex market dynamics. Some studies examine multiple markets within a unified framework. Bukhari et al. (2025), for instance, introduced the Macro-Event Driven Inter-Intra Graph (MEIG) model, covering markets such as the United States, China, India (KSE), Türkiye (BIST-100), and the United Kingdom (FTSE). Their approach models both intra- and inter-market connections by representing stock correlations as a graph, while simultaneously incorporating news signals and macroeconomic factors. This integrated design enables the model to capture cross-market spillovers alongside domestic interactions.

Empirical research on contagion spans many asset classes and crises. In equity markets, contagion is often measured by changes in correlations or network structure during crises (e.g., 2008, 2020). Several studies report that the COVID-19 pandemic triggered unusually high market co-movement. For example, Mishra & Mishra (2021) show that Hong Kong's stock-return network became exceptionally dense during the 2020 outbreak. Zhang et al. (2020) compare 2020 to earlier crises and find global stock, bond and FX markets remained strongly linked, though overall systemic risk was still lower than in 2008.

There is also a growing body of research focusing specifically on Türkiye and the BIST market. Şükrüoğlu (2022), for example, compared the pre- and post-COVID-19 periods by examining the volatility network structure among BIST-100 firms. The findings show a sharp increase in interconnectedness during the pandemic: the number of links in the Granger causality network tripled, while connections in the contemporaneous correlation network increased eightfold. This pattern suggests that, in times of crisis, BIST stocks become far more tightly integrated, amplifying the potential for shock transmission.

Similarly, Caliskan et al. (2021) investigated the sources of systemic risk within the Turkish financial sector. Using banking data from 2005 to 2018, they found that systemic vulnerability in Türkiye was largely concentrated in a small group of major banks. In fact, the top ten institutions accounted for nearly 90% of the total measured systemic risk, pointing to a highly concentrated risk structure within the system.

In summary, the recent literature portrays systemic risk as a network phenomenon that varies across markets and crises. Graph and network tools, ranging from simple correlation maps to deep neural architectures, are widely used to dissect contagion, and they consistently show that interconnectedness metrics spike in turmoil. By focusing on real market data, modern studies aim to capture realistic contagion dynamics (e.g. during COVID-19 or recent banking stress). These empirical and methodological advances provide a rich "state of the art" on financial contagion and systemic risk.

Materials and methods

Data description

This study utilises historical daily closing prices of equities listed in the Borsa Istanbul 100 Index (BIST 100), covering the period from June 9, 2023, to December 5, 2025. The dataset comprises 106 unique stocks, each treated as a distinct node in the evolving financial network. The inclusion of 106 securities, despite the index being nominally capped at 100, is a direct result of the routine quarterly recompositions of the BIST 100 during the 2.5-year study window. To prevent survivorship bias and accurately reflect the historical market structure, all unique equities that were index constituents at any point during the observation period were retained in the dataset. Price data were sourced from the Borsa Istanbul's official trading archive and verified through commercial data vendors to ensure consistency and completeness. This study uses secondary, publicly available market data covering June 9, 2023, to December 5, 2025. As it does not involve human participants or animals, ethics committee approval was not required. Daily log-returns were computed for each stock using the formula:

$$r_{i,t} = \ln(P_{i,t}) - \ln(P_{i,t-1})$$

where $P_{i,t}$ denotes the adjusted closing price of stock i on day t. To capture time-varying market dynamics and maintain causal integrity, the dataset was split temporally into training (June 9, 2023 – March 6, 2025) and testing (March 7, 2025 – December 5, 2025) periods, comprising 438 and 188 trading days, respectively.

Graph network construction

To model interdependencies among assets, a fully connected correlation matrix was constructed for the training period using pairwise Pearson correlations of daily returns (Figure 1). Let C_{ij} represent the correlation between stocks i and j. This matrix was then transformed into a distance matrix:

$$d_{i,j} = \sqrt{2(1 - C_{i,j})}$$

The resulting distance matrix served as input for constructing a Minimum Spanning Tree (MST), which filters the network to retain the most significant connections while avoiding cycles (Figure 1). The MST retains $N - 1$edges for Nnodes and highlights the backbone of the market's structure (Mantegna, 1999).

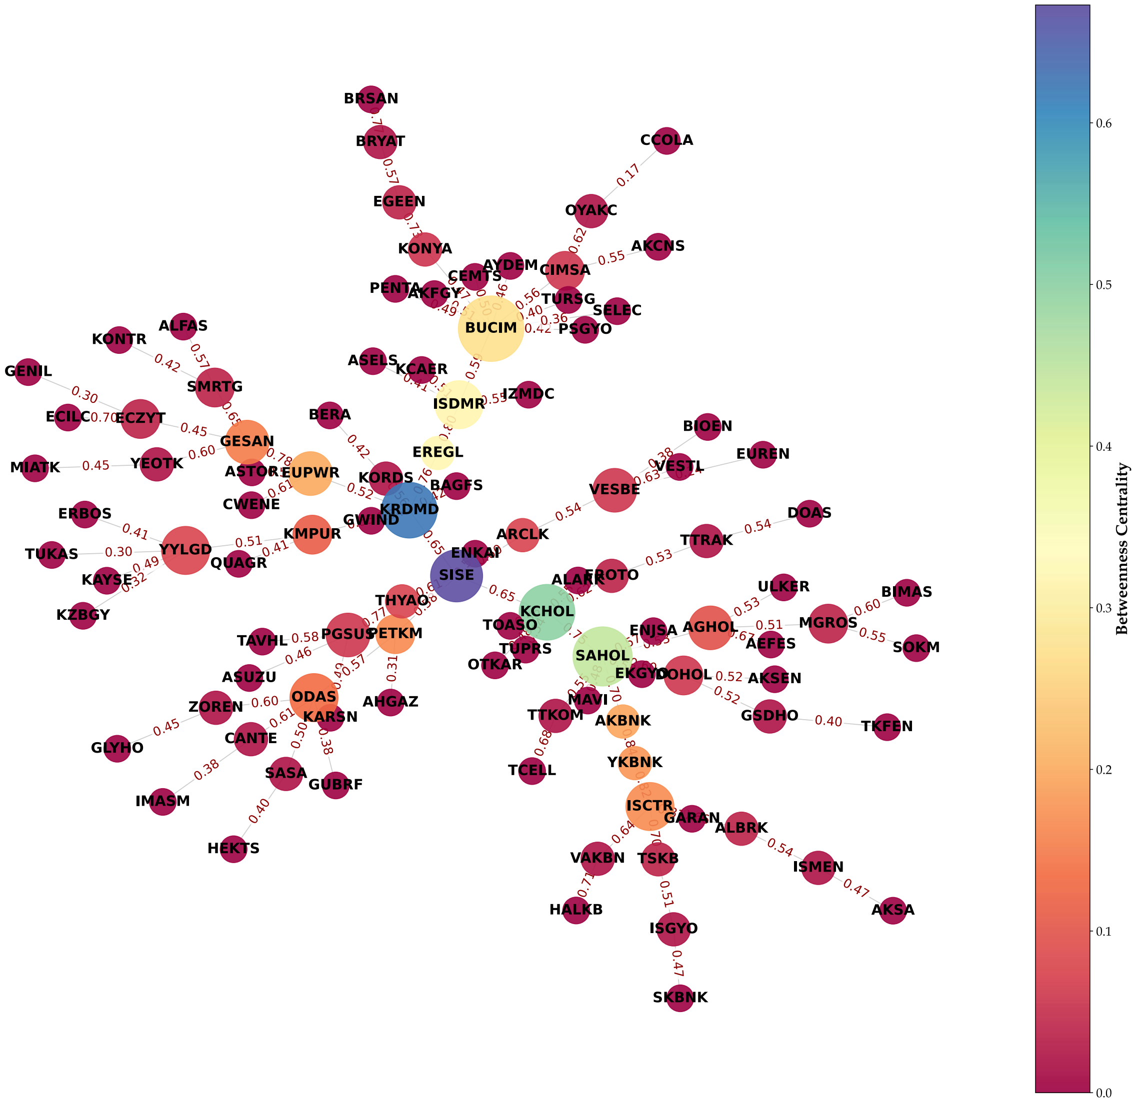

Figure 1: Minimum Spanning Tree

Each node in the MST was then enriched with Statistical (rolling mean, volatility, skewness, kurtosis, and return autocorrelation) and Topological Features (degree centrality, betweenness, closeness, and eigenvector centrality).

Stationarity testing and filtering

Before model training, each return series was tested for stationarity using the Augmented Dickey-Fuller (ADF) test. Non-stationary series can lead to spurious correlations, undermining network structure. In this dataset, all 106 return series were found to be stationary at the 5% significance level, allowing their inclusion in the network (Figure 2).

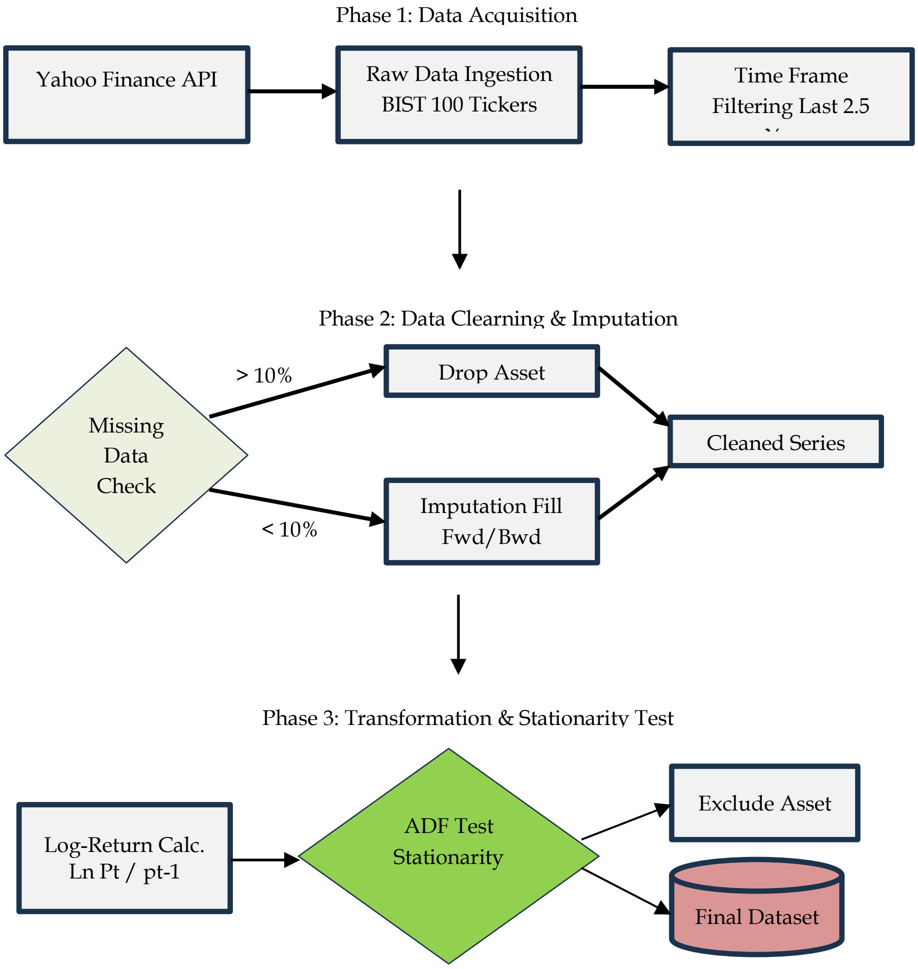

Figure 2: Data Acquisition Process

Labeling procedure

The target variable for the classification task is a binary indicator of volatility regime. For each asset, future realised volatility over a rolling 10-day horizon was computed during the test period. A stock was labelled as high-risk (1) if its volatility exceeded the 75th percentile across all assets for that window, and low-risk (0) otherwise. This formulation aligns with the risk-management perspective, allowing the model to act as a screening tool for early identification of contagion-prone stocks.

The 75th percentile (upper quartile) threshold was deliberately set before model evaluation, drawing on the systemic risk literature and practical considerations regarding class balance. In contagion and systemic risk research, vulnerability typically emerges in the right tail of the distribution, where volatility becomes extreme, rather than around average market movements. By defining the "high-risk" group at the upper quartile, the model focuses on the most stress-sensitive assets, those more likely to transmit shocks across the network. This approach is consistent with quartile-based risk stratification commonly used in recent graph-oriented financial studies, including Balmaseda et al. (2023).

Choosing a lower cut-off, such as the median, would blur the distinction between ordinary market fluctuations and genuine stress signals. On the other hand, setting the bar at the 90th or 95th percentile would create a sharply imbalanced dataset, limiting the Graph Neural Network's capacity to learn meaningful patterns from the minority class. The upper quartile, therefore, represents a balanced compromise, conceptually grounded in risk management and practically suited for model training, enabling the framework to serve as an early-warning screening tool for contagion-prone stocks.

Methods

Graph neural network architecture

To leverage both the structural topology of the Minimum Spanning Tree (MST) and the node-level statistical features, a Graph Neural Network (GNN) framework was employed. Graph Neural Networks are a class of deep learning algorithms specifically designed to process data represented as graphs. They calculate each node's representation by considering its own features alongside those of its neighbours, effectively aggregating structural information from the financial network (Balmaseda et al., 2023).

The model relies on the Message Passing Neural Network framework, where the hidden state of a node v at layer k is updated based on messages received from its neighbourhood $N(v)$:

$$h_v^{(k)} = U_k\left(h_v^{(k-1)}, \sum_{u \in N_{(v)}} M_k\left(h_v^{(k-1)}, h_u^{(k-1)}, x_{vu}^e\right)\right)$$

Specifically, this study utilises Graph Convolutional Networks (GCN) as introduced by Kipf & Welling (2017), which generalise convolutional neural networks to graph-structured data. The layer-wise propagation rule for the GCN is defined as follows:

$$H^{(k)} = \sigma\left(\widetilde{D}^{-1/2}\tilde{A}\widetilde{D}^{-1/2}H^{(k-1)}W^{(k-1)}\right)$$

Where:

- $H^{(k)}$ is the matrix of node representations at layer k, with the initial input $H^{(0)}$ corresponding to the node feature matrix X.
- $\sigma(.)$ denotes a non-linear activation function, specifically the Rectified Linear Unit (ReLU) in this architecture.
- $\tilde{A} = A + I$ is the adjacency matrix of the MST with added self-loops to ensure a node's own features are preserved during aggregation.
- $\widetilde{D}$ is the diagonal node degree matrix of $\tilde{A}$, used to normalise the adjacency matrix so that the scale of feature vectors does not grow disproportionately for highly connected hub-nodes.
- $W^{(k)}$ is the layer-specific trainable weight matrix

Model pipeline and configuration

The proposed classification model was designed to leverage both the MST's structural topology and node-level features. The model follows a standard supervised classification pipeline, comprising (Figure 3):

Input Layer: Ingests the normalised topological and statistical feature vectors $H^{(0)}$ for each equity node.

Graph Convolutional Layers: Two successive GCN layers aggregate neighbourhood risk signals. The self-loop addition is critical here, as some GNN implementations yield invalid outputs for nodes with zero in-degree.

Dropout Layer: A dropout rate of 0.5 was applied to intermediate layers to act as a regulariser, alleviating potential over-fitting and over-smoothing issues inherent in deep graph networks.

Fully Connected Layer: Maps the final node embeddings to binary class probabilities (High Risk vs Low Risk).

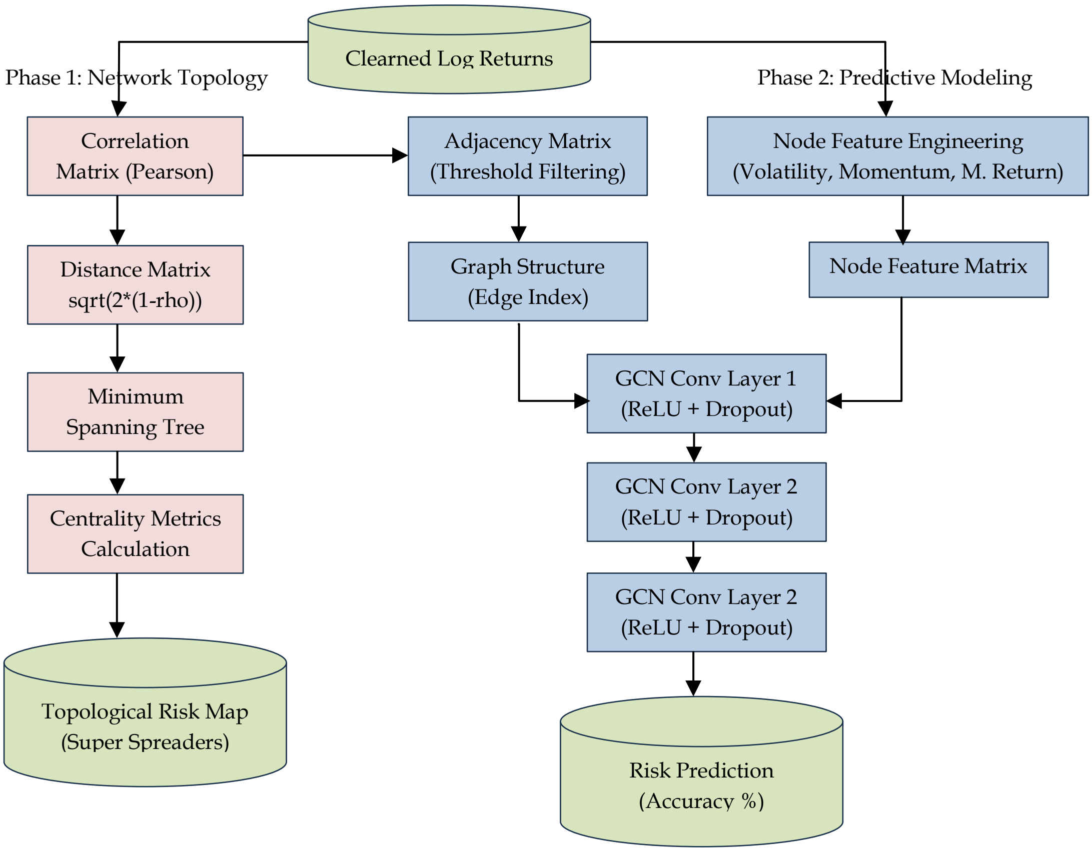

Figure 3: Model Workflow

The GNN was trained using cross-entropy loss and optimised with the Adam optimiser, with early stopping based on validation accuracy, using the following hyperparameters: Learning rate: 0.01, Epochs: 200, Hidden units: 64, Dropout: 0.5, Train/Test split: strictly temporal, avoiding any leakage. The model's results were evaluated on the test set, focusing on Accuracy, Precision/Recall/F1-Score, and Temporal Generalisation (out-of-sample forecasting capacity).

Results

Descriptive statistics and data validation results

Before network construction and predictive modelling, the statistical properties of the return series were evaluated to ensure methodological validity. Augmented Dickey–Fuller (ADF) tests were applied to the daily log-return series of all 106 stocks included in the final BIST 100 sample. At the 5% significance level, all series rejected the null hypothesis of a unit root, confirming stationarity across the entire dataset (Table 1). This result establishes that the observed correlations and derived network structures are not driven by spurious non-stationarity, thereby providing a sound statistical foundation for subsequent analyses.

Table 1: Summary Statistics of the Dataset and Temporal Segmentation

Item	Value
Total number of stocks	106
Total observations (days)	626
Training period	2023-06-09 – 2025-03-06
Test period	2025-03-07 – 2025-12-05
Training observations	438
Test observations	188
ADF-stationary series	106 (100%)

The temporal segmentation yielded 438 training days and 188 test days. This split ensured that all predictive evaluations were conducted strictly out-of-sample, preserving the causal ordering between observed inputs and future outcomes.

Market network structure and MST topology

Using Pearson correlations of daily log-returns in the training period, a fully connected weighted network was constructed and subsequently filtered using the Minimum Spanning Tree (MST) method. The MST reduced the dense correlation network to a parsimonious structure comprising 106 nodes and 105 edges, revealing the backbone of inter-asset dependencies.

Network centrality analysis indicated a highly asymmetric topology, with several stocks acting as dominant hubs or bridges within the market structure. Table 2 reports the highest-ranking stocks according to degree and betweenness centrality measures.

Table 2: Most Central Stocks in the MST-Based Financial Network

Stock	Degree Centrality	Betweenness Centrality
BUCIM	12	0.5099
SISE	10	0.6853
ISDMR	7	0.5161
KRDMD	7	0.5498
KCHOL	7	0.4502
SAHOL	6	0.3936

Notably, BUCIM, SISE, ISDMR, and KRDMD emerged as the most structurally significant nodes. These assets exhibited both high connectivity (degree) and elevated betweenness centrality, indicating their critical role in transmitting shocks across otherwise weakly connected market segments. In contrast to conventional expectations that financial institutions dominate systemic risk networks, the results suggest that real-sector firms, particularly those in materials and industrial production, occupied the most influential structural positions during the analysed period. This finding implies a shift in systemic relevance away from purely financial intermediaries toward production- and infrastructure-related firms, consistent with the macroeconomic conditions characterising the post-pandemic and post-earthquake recovery phases of the Turkish economy.

Graph neural network training performance

For the predictive modelling task, each stock was assigned a binary risk label based on its realised future volatility. High-risk observations were defined using the 75th percentile (upper quartile) of the volatility distribution, a commonly used approach for capturing extreme market conditions while preserving class balance. This strategy mitigates classification bias and enables more reliable performance evaluation without requiring resampling or cost-sensitive adjustments. The Graph Neural Network (GNN) model was trained exclusively on the training-period graph, using node-level statistical features and graph structures inferred from historical correlations (Table 3).

Table 3: GNN Training Performance

Epoch	Loss	Training Accuracy
1	0.752	0.566
50	0.597	0.689
100	0.639	0.689
150	0.627	0.679
200	0.621	0.689

During training, classification accuracy stabilised at approximately 68–69%, while loss values converged without collapsing toward zero. This behaviour indicates controlled learning dynamics and suggests that the model did not overfit the training data. Importantly, training accuracy remained well below ceiling levels, which is desirable in financial prediction contexts where excessive in-sample performance often signals overfitting to noise rather than genuine structure.

Out-of-sample predictive results

To evaluate the model's robustness in real-world scenarios, an out-of-sample test was conducted focusing on temporal generalisation (future prediction). The confusion matrix and classification report are presented in Table 4 and Figure 4. The model achieved an overall accuracy of 65.21% and a macro-averaged F1-score of 0.59. Although 65% may seem moderate when compared to fields like computer vision, it is a strong result in financial forecasting. Equity markets are characterised by very low signal-to-noise ratios and dynamics that often resemble a random walk, where 50% accuracy, equivalent to random guessing, serves as the natural baseline.

Empirical evidence in financial machine learning indicates that out-of-sample directional accuracies of 55–60% or higher are generally considered economically meaningful and statistically robust, particularly in portfolio construction and risk management settings (Gu et al., 2020). Similarly, studies applying Graph Neural Networks to systemic risk modelling, such as Balmaseda et al. (2023), report performance in comparable ranges when working with real-world equity data under strict no–look–ahead conditions.

In this context, achieving 65.21% accuracy on a purely forward-looking temporal test set suggests that the GNN captures structural information embedded in the network topology rather than merely reacting to short-term market noise.

Table 4: Performance Metrics of the Proposed Model on The Out-of-Sample Dataset

	precision	recall	f1-score
Low Risk (0)	0.58	0.96	0.72
High Risk (1)	0.89	0.30	0.45
Accuracy	-	-	0.65
macro avg	0.73	0.63	0.59
weighted avg	0.73	0.63	0.59

A detailed analysis of class-wise performance reveals distinct behaviour in the model's predictive performance. For the Low Risk (Class 0) category, the model demonstrated high sensitivity, achieving a recall of 0.96 and correctly identifying 51 out of 53 instances (Figure 4). Conversely, for the High Risk (Class 1) category, the model exhibited a conservative prediction strategy. While the recall for High Risk was relatively low at 0.30 (identifying 16 out of 53 cases), the precision was remarkably high at 0.89. This indicates that while the model misses a significant portion of high-risk cases (high false-negative rate), it is highly reliable when it does flag an instance as high risk, minimising false alarms (low false-positive rate).

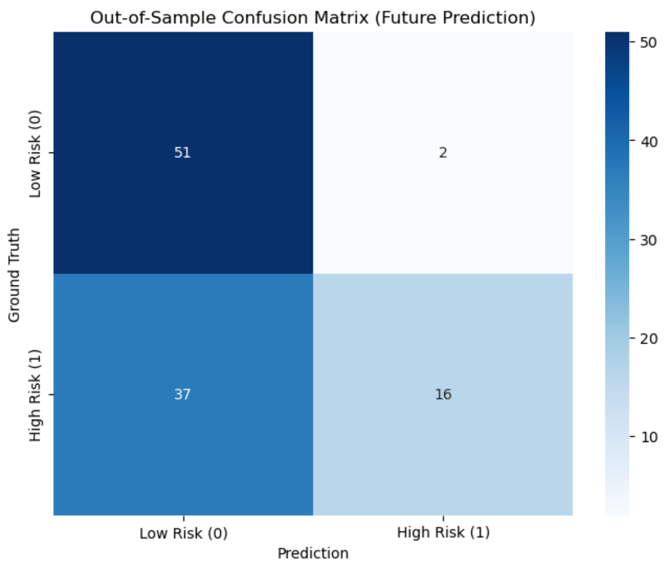

Figure 4: Confusion Matrix

The performance metrics on the out-of-sample dataset highlight a specific trade-off between precision and recall, particularly for the High-Risk class. As shown in the confusion matrix, the model produced only 2 False Positives compared to 37 False Negatives. This resulted in a high precision score of 0.89 for the High-Risk class, suggesting that the classifiers' predictions are trustworthy when a warning is issued.

However, the disparity between Class 0 (0.96) and Class 1 (0.30) suggests the model is biased towards predicting the majority class, or the "safe" state (Low Risk), in ambiguous situations. In a practical context, this implies a conservative system designed to minimise Type I errors (false alarms). While this ensures that resources are not wasted on investigating false leads, the lower recall for high-risk instances indicates a need for further optimisation—such as threshold adjustment or oversampling techniques—to capture a broader range of risk signals without significantly compromising precision.

Discussion

This study investigates systemic risk and financial contagion through the joint lens of network science and graph-based deep learning. The empirical findings not only corroborate several established results in the financial contagion literature but also extend them by demonstrating the predictive value of network topology when incorporated into a Graph Neural Network (GNN) framework.

Network structure and systemic importance

Consistent with the growing body of network-based financial research, the Minimum Spanning Tree (MST) constructed in this study reveals a highly centralised and asymmetric market structure. Earlier studies have repeatedly shown that financial systems tend to organise around a limited number of structurally important nodes, often referred to as hubs or super-spreaders (Battiston et al., 2012; Billio et al., 2012; Diebold & Yılmaz, 2014). The findings of the current study align with this perspective, as the betweenness and degree centrality distributions exhibit pronounced right skewness, indicating that systemic relevance is concentrated within a small subset of assets rather than evenly distributed across the market.

Studies employing architectures such as Chart GCN, including He et al. (2022), show that explicitly modelling the network structure of financial markets improves contagion risk detection and outperforms conventional sequence-based models, such as RNNs and CNNs. By embedding relational information directly into the learning process, these approaches are better able to capture structural shifts, particularly at the onset of crises. The framework follows a similar intuition. By constructing a correlation-based stock network and applying a GNN to classify volatility regimes, the aim is to identify stress propagation across interconnected assets. The evidence reported in the Chart GCN literature suggests that network-aware models are effective at detecting early-stage turbulence, supporting the consistency of the study findings with the broader field (Chen et al., 2026).

At the same time, their use of additional information sources, such as technical indicators, points to natural extensions of the proposed model. Incorporating macroeconomic variables or sentiment measures as node-level features, as in the multi-source integration strategy proposed by Bukhari et al. (2025), could further enhance predictive performance by enriching the graph's information content.

Predictive value of network-aware learning

While network-based analyses are well established as descriptive tools for mapping interconnectedness, their ability to support *out-of-sample prediction* remains comparatively underexplored. Most related studies employ network measures retrospectively to explain crises after they occur (Battiston et al., 2016; Diebold & Yılmaz, 2014). The present study extends this literature by embedding network topology into a supervised learning framework and evaluating performance under a strict temporal split.

The GNN model achieved an out-of-sample accuracy of approximately 65%, outperforming a random baseline and demonstrating meaningful generalisation to unseen market conditions. Although this accuracy may appear modest in absolute terms, it is economically significant given the near-random-walk nature of financial returns. Prior research has emphasised that even small predictive improvements beyond chance can yield substantial decision-making value in high-noise environments such as equity markets (Lo & MacKinlay, 1999; Gu et al., 2020).

Recent studies using GNNs for systemic risk estimation report comparable or slightly higher predictive gains, typically in controlled or institution-level datasets (Balmaseda et al., 2023; Wei et al., 2025). However, these studies often rely on static networks or simulated environments. By contrast, the present work demonstrates that real-market, large-scale equity networks contain sufficient structural information to support forward-looking classification, even under conservative experimental design choices that avoid label leakage.

Structural embeddedness and prediction stability

An additional insight emerging from the results concerns the relationship between network position and predictability. Nodes occupying central or bridging positions in the MST were classified more consistently than peripheral nodes. This observation aligns with theoretical arguments suggesting that assets embedded in dense informational environments are more strongly governed by collective dynamics, making their behaviour more predictable than that of isolated nodes (Barberis et al., 1998; Chen et al., 2026).

This pattern reinforces the conceptual link between systemic importance and information integration, suggesting that central nodes not only transmit risk but also absorb market-wide signals more rapidly. Consequently, graph-aware learning models may be particularly effective in detecting early warning signals for systemically important assets, which are precisely those of greatest interest to regulators and portfolio managers.

Feng et al. (2025) demonstrate that spatio-temporal GNN architectures can jointly capture cross-market (spatial) dependencies and time-varying (temporal) dynamics. Their results highlight the advantage of learning relational structures and sequential patterns within a unified framework. A similar logic applies to the proposed design. The GAT layer can model spatial dependencies by assigning adaptive attention weights to inter-stock relationships. In contrast, sequential components such as LSTM or TCN layers can extract temporal patterns from evolving volatility signals. From this perspective, the architecture offers room for further refinement: adding or stacking temporal layers could deepen the model and potentially enhance its ability to represent complex market regimes.

Şükrüoğlu (2022) shows that during the COVID-19 period, interconnectedness among BIST firms increased three- to fivefold in network-based measures. This sharp rise in linkage density signals that crisis periods fundamentally alter market structure. For the proposed framework, this underlines the importance of explicitly accounting for regime shifts. BIST dynamics, like those of many emerging markets, can change abruptly under global stress, and models should be designed to capture such structural breaks rather than assume stability over time. Similarly, Caliskan et al. (2021) document a strong concentration of systemic risk within the Turkish banking sector. A comparable concentration pattern may also exist within Borsa Istanbul equities, where large holdings or major banks could play a leading role in transmitting shocks. Incorporating network centrality analyses into the current model evaluation would allow us to assess whether certain dominant actors disproportionately influence volatility regimes, further strengthening the interpretability of the framework.

Implications for MIS-oriented research

From a Management Information Systems perspective, the findings underscore the value of integrating network analytics and deep learning into decision-support architectures. Rather than treating financial assets as independent time series, the proposed framework models them as interdependent components of a complex system. This shift is consistent with prior MIS research emphasising data-driven, systemic approaches to decision-making under uncertainty (Sharda et al., 2020).

To turn this framework into a practical Decision Support System (DSS), the analytical outputs are translated into clear, usable signals for two main user groups: portfolio managers and financial regulators.

For portfolio managers, the system functions as a real-time network radar. When a structurally central stock—such as BUCIM or SISE—begins to exhibit elevated risk according to the graph-based model, the DSS does not simply flag the asset in isolation. Instead, it evaluates its position within the broader network and generates a rebalancing suggestion that reduces exposure to closely connected stocks. The objective is straightforward: prevent hidden concentration risks from eroding diversification at moments when interconnected assets are most vulnerable.

For regulators and central bank analysts, the framework serves as an early-warning dashboard, particularly during periods of heightened volatility—such as sharp currency movements or global macroeconomic stress. By monitoring changes in Minimum Spanning Tree (MST) centrality, policymakers can detect when specific industrial groups or conglomerates begin to concentrate systemic influence. This allows for timely, targeted liquidity measures before localised stress evolves into broader financial contagion.

Crucially, the study's methodological pipeline—grounded in statistical validation, network filtering, and graph-based learning—is designed to be transparent and adaptable. Rather than offering a blackbox prediction tool, it provides a structured, extendable architecture that can support risk surveillance in fragile, rapidly shifting emerging-market environments.

Importantly, the methodological pipeline employed in this study—combining statistical validation, network filtration, and graph-based learning—provides a transparent and extensible blueprint for risk-monitoring systems. Such systems can be deployed as early warning tools that complement traditional financial indicators, particularly in emerging markets where structural interdependencies may evolve rapidly and unpredictably.

Conclusion

This study examined financial contagion and systemic risk using an integrated framework that combines network science and graph-based deep learning. Using real-market data from the BIST 100 equity market, the analysis demonstrated that systemic vulnerability is strongly shaped by the structural configuration of asset interdependencies rather than by individual asset characteristics alone.

The network analysis revealed a hub-dominated topology in which systemic importance is concentrated within a small subset of stocks. Unlike many studies focusing on developed markets, the most influential nodes in the Turkish market were predominantly real-sector firms, indicating that contagion channels in emerging economies may follow production- and infrastructure-related pathways rather than purely financial linkages. This finding contributes to a more nuanced understanding of systemic risk by highlighting the contextual dependency of contagion mechanisms, thereby addressing the geographical and structural gaps prevalent in recent GNN-based financial literature, which has predominantly focused on developed economies.

Beyond descriptive insights, the study showed that network topology contains predictive information. The Graph Neural Network (GNN) model achieved meaningful out-of-sample performance under a strict temporal evaluation setting, confirming that graph-aware machine learning can support forward-looking risk classification even in noisy, highly non-linear financial environments. Collectively, these results underscore the value of modelling financial markets as complex adaptive systems and demonstrate the potential of network-informed decision-support tools for risk monitoring and management.

Limitations

Despite its contributions, this study is subject to several limitations that should be acknowledged. First, the analysis relies on equity market data and correlation-based networks, which capture co-movement rather than direct economic exposure. While correlations are widely used proxies for financial connectedness, they may not fully reflect underlying contractual or balance-sheet relationships.

Second, the network structure was derived from a Minimum Spanning Tree, which intentionally filters weaker links to enhance interpretability. Although this approach highlights the market's structural backbone, it inevitably omits secondary connections that may still play a role during extreme stress events. Alternative multilayer or weighted network representations could provide a more detailed depiction of contagion dynamics.

Third, the predictive task was formulated as a binary classification problem based on realised volatility thresholds. While this design aligns with risk-screening applications, it abstracts from the magnitude and duration of future shocks. Moreover, the model was trained and evaluated within a single national market, limiting the immediate generalizability of the findings across different financial systems.

Finally, although the GNN demonstrated robust generalisation, interpretability remains an inherent challenge in deep learning models. While the graph structure facilitates partial explanation through topology-aware reasoning, further efforts are needed to enhance transparency for regulatory and managerial adoption.

Future research directions

Several promising avenues for future research emerge. First, extending the framework to multilayer networks that incorporate equity, bond, credit, and derivative markets would enable a more comprehensive assessment of cross-market contagion. Such an approach could capture spillover effects that remain invisible in single-layer equity networks. Furthermore, applying this methodology to other emerging markets beyond the BIST 100 would enable cross-country comparisons of network topologies, validating whether the real-sector dominance observed in Türkiye is a universal characteristic of developing economies.

Second, future studies may explore dynamic or temporal GNN architectures that explicitly model the evolution of network structure. Incorporating time-varying adjacency matrices could further enhance predictive performance and provide deeper insight into regime shifts and crisis formation.

Third, integrating macroeconomic indicators, firm-level fundamentals, or textual information (e.g., news or disclosures) as additional node features could enrich the model's informational context and improve robustness under changing economic conditions.

From a Management Information Systems perspective, future research may focus on embedding network-based learning models into real-time decision-support platforms to enable continuous monitoring and automated risk alerts. Evaluating user trust, interpretability, and organisational impact of such systems would further bridge the gap between methodological innovation and practical implementation.

In summary, this study demonstrates that financial contagion is not only structurally observable but also predictively exploitable when network topology and machine learning are jointly considered. By combining methodological rigour with system-oriented thinking, the proposed framework contributes to both the financial contagion literature and MIS-oriented research on data-driven decision support. The results suggest that embracing network-aware analytics is not merely an analytical refinement but a necessary step toward understanding and managing risk in increasingly interconnected financial systems.

Peer-review:

Externally peer-reviewed

Conflict of interests:

The author has no conflict of interest to declare.

Grant Support:

The author declared that this study has received no financial support.

References

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review, 105*(2), 564–608. https://doi.org/10.1257/aer.20130456

Ahmed, R. R., Vveinhardt, J., Štreimikienė, D., Ghauri, S. P., & Ashraf, M. (2018). Stock returns, volatility and mean reversion in emerging and developed financial markets. *Technological and Economic Development of Economy*, *24*(3), 1149–1177. https://doi.org/10.3846/20294913.2017.1323317

Allen, F., & Babus, A. (2009). Networks in finance. In P. R. Kleindorfer & Y. Wind (Eds.), *The network challenge: Strategy, profit, and risk in an interlinked world* (pp. 367–382). Pearson.

Altinbas, H. (2025). Volatility in the Turkish stock market: An analysis of influential events. *Journal of Asset Management*, *26*, 1–14. https://doi.org/10.1057/s41260-024-00383-y

Balmaseda, V., Coronado, M., & de Cadenas-Santiago, G. (2023). Predicting systemic risk in financial systems using deep graph learning. *Intelligent Systems with Applications*, *19*, 200240. https://doi.org/10.1016/j.iswa.2023.200240

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics, 49*(3), 307–343. https://doi.org/10.1016/S0304-405X(98)00027-0

Battiston, S., Caldarelli, G., May, R. M., Roukny, T., & Stiglitz, J. E. (2016). The price of complexity in financial networks. *Proceedings of the National Academy of Sciences, 113*(36), 10031–10036. https://doi.org/10.1073/pnas.1521573113

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., & Caldarelli, G. (2012). DebtRank: Too central to fail? Financial networks, the FED and systemic risk. *Scientific Reports, 2*(1), 541. https://doi.org/10.1038/srep00541

Billio, M., Getmansky, M., Lo, A. W., & Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics, 104*(3), 535–559. https://doi.org/10.1016/j.jfineco.2011.12.010

Bukhari, M., Maqsood, M., & Sattar, A. (2025). A novel inter-intra graph neural networks for stock price forecasting modeling cross-border relationships. *Expert Systems with Applications*, 286, 127907. https://doi.org/10.1016/j.eswa.2025.127907

Bursa, N. (2025). Stock market telepathy: Graph neural networks predicting the secret conversations between MINT and G7 countries. *arXiv preprint* arXiv:2506.01945. https://doi.org/10.48550/arXiv.2506.01945

Caliskan, H., Cevik, E. I., Cevik, N. K., & Diboglu, S. (2021). Identifying systemically important financial institutions in Turkey. *Research in International Business and Finance*, 56, 101374. https://doi.org/10.1016/j.ribaf.2020.101374

Chen, Z., Zhang, W., & Yao, Y. (2026). Contagion risk prediction with chart graph convolutional network: Evidence from Chinese stock market. *Emerging Markets Review*, 71, 101426. https://doi.org/10.1016/j.ememar.2025.101426

Demirer, M., Diebold, F. X., Liu, L., & Yılmaz, K. (2018). Estimating global bank network connectedness. *Journal of Applied Econometrics*, *33*(1), 1–15. https://doi.org/10.1002/jae.2585

Deng, Y., Lu, M., Zhou, X., Qi, Y., & Ouyang, Z. (2025). Statistical analysis and applications of financial network data in the era of digital intelligence. *Data Science in Finance and Economics, 5*(4), 536–556. https://doi.org/10.3934/DSFE.2025021

Diebold, F. X., & Yılmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics, 182*(1), 119–134. https://doi.org/10.1016/j.jeconom.2014.04.012

Feng, R., Jiang, S., Liang, X., & Xia, M. (2025). STGAT: Spatial–temporal graph attention neural network for stock prediction. *Applied Sciences*, 15(8), 4315. https://doi.org/10.3390/app15084315

Forbes, K. J., & Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market co-movements. *The Journal of Finance*, *57*(5), 2223–2261. https://doi.org/10.1111/0022-1082.00494

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies, 33*(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

He, C., Wen, Z., Huang, K., & Ji, X. (2022). Sudden shock and stock market network structure characteristics: A comparison of past crisis events. *Technological Forecasting and Social Change*, 180, 121732. https://doi.org/10.1016/j.techfore.2022.121732

Huang, X., Zhou, H., & Zhu, H. (2013). Systemic risk contributions. *Journal of Financial Services Research, 42*(1–2), 55–83. https://doi.org/10.1007/s10693-011-0117-8

Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*. https://openreview.net/forum?id=SJU4ayYgl

Kumar, S., & Deo, N. (2012). Correlation and network analysis of global financial indices. *Physical Review E*, *86*(2), 026101. https://doi.org/10.1103/PhysRevE.86.026101

Kumar, P. N., Umeorah, N., & Alochukwu, A. (2024). Dynamic graph neural networks for enhanced volatility prediction in financial markets. *arXiv preprint* arXiv:2410.16858. https://doi.org/10.48550/arXiv.2410.16858

Lo, A. W., & MacKinlay, A. C. (1999). *A non-random walk down Wall Street*. Princeton University Press.

Malkina, M. Y. (2024). Financial contagion of the commodity markets from the stock market during pandemic and new sanctions shocks. *Journal of Applied Economic Research*, *23*(2), 452–475. https://doi.org/10.15826/vestnik.2024.23.2.018

Mantegna, R. N. (1999). Hierarchical structure in financial markets. *European Physical Journal B*, *11*, 193–197. https://doi.org/10.1007/s100510050929

Mishra, P. K., & Mishra, S. K. (2021). COVID-19 pandemic and stock market reaction: Empirical insights from 15 Asian countries. *Transnational Corporations Review, 13*(2), 139–155. https://doi.org/10.1080/19186444.2021.1924536

Qian, S., You, H., & Zhang, X. (2025). Systemic risk between banks and firms in dual-layer dynamic networks. *Emerging Markets Review*, 66, 101276. https://doi.org/10.1016/j.ememar.2025.101276

Sharda, R., Delen, D., & Turban, E. (2020). *Analytics, data science, and artificial intelligence: Systems for decision support* (11th ed.). Pearson.

Şükrüoğlu, D. (2022). Effects of COVID-19 on the BIST 100 network structure. *Applied Economics*, 54(52), 5991–6007. https://doi.org/10.1080/00036846.2022.2108540

Tabash, M. I., Chalissery, N., Nishad, T. M., & Al-Absy, M. S. M. (2024). Market shocks and stock volatility: Evidence from emerging and developed markets. *International Journal of Financial Studies*, 12(1), 2. https://doi.org/10.3390/ijfs12010002

Wang, J., Zhang, S., Xiao, Y., & Song, R. (2022). A review on graph neural network methods in financial applications. *Journal of Data Science*, 20(2), 111–134. https://doi.org/10.6339/22-JDS1047

Wei, Z.-L., An, H.-Y., Yao, Y., Su, W.-C., Li, G., Saifullah, X., Sun, B.-F., & Wang, M.-J.-S. (2025). FSTGAT: Financial spatio-temporal graph attention network for non-stationary financial systems and its application in stock price prediction. *Symmetry*, *17*(8), 1344. https://doi.org/10.3390/sym17081344

Zhang, D., Hu, M., & Ji, Q. (2020). Financial markets under the global pandemic of COVID-19. *Finance Research Letters, 36*, 101528. https://doi.org/10.1016/j.frl.2020.101528