International Journal of Accounting and Economics Studies, 13 (2) (2026) 331-342

International Journal of Accounting and Economics Studies

Website: www.sciencepubco.com/index.php/IJAES
https://doi.org/10.14419/mh97vb34
Research paper

Graph Neural Networks for Systemic Financial Risk Forecasting: Modeling Cross-Market Contagion Between Banking Systems and Cryptocurrency Markets

Md Zahidul Islam [1] *, Md Sumsuzoha [2], Md Rafiqul Islam [3], Mohammed Kawsar [4], Md Fazlul Huq Mithu [5], Santosh Pant [6], Mohammad Nazmul Hossain [7], Md Abdullah Al Helal [8]

[1] *MBA in Business Analytics, Gannon University, Erie, PA*
[2] *Master of Science in Business Analytics, Trine University*
[3] *DBA in Business Analytics, International American University*
[4] *MSc, Analytics & Information Management, Duquesne University*
[5] *MS in Finance, Stony Brook University*
[6] *BBA, Kantipur College of Management and Information Technology, Kathmandu, Nepal*
[7] *Computer/Information Technology Services Administration and Management, St. Francis College*
[8] *Master of Science in Business Analytics, Trine University*
**Corresponding author E-mail: islam013@gannon.edu*

Received: January 24, 2026, Accepted: February 18, 2026, Published: February 25, 2026

Abstract

Evolving interdependencies across institutions and markets drive systemic financial risk, yet most forecasting models either treat assets independently or rely on static correlation structures. This limitation becomes particularly salient as cryptocurrency markets increasingly interact with traditional banking systems amid financial stress. Ignoring time-varying cross-market network structure risks understating tail risk precisely during periods when accurate systemic risk assessment is most critical. This study proposes a dynamic graph neural network (GNN) framework for systemic risk forecasting that models time-varying financial networks spanning banking institutions and major cryptocurrency assets. Nodes represent financial entities, while edges are constructed using rolling-window dependency measures that adapt to changing market conditions. Node dynamics are modeled through temporal neural architectures, and stress regimes are explicitly identified to evaluate performance under market turmoil. The empirical design includes strong temporal baselines, static-graph ablations, and cross-market removal experiments to isolate the contribution of network dynamics and crypto-market integration. Results indicate that a strong LSTM baseline achieves superior volatility forecasting accuracy in both overall and stress-period evaluations, demonstrating the competitiveness of purely temporal models. However, within the class of graph-based models, dynamic GNNs substantially outperform static-graph variants during stress periods, demonstrating the importance of time-varying network structure for capturing volatility amplification. Bank-only and full-system dynamic GNNs exhibit comparable stress-period performance, suggesting that cryptocurrency assets contribute limited incremental information to bank-specific forecasts, while remaining informative for system-level stress characterization. The findings suggest that dynamic graph representations enhance stress sensitivity and structural interpretability relative to static network models, even when they do not surpass strong temporal baselines in raw predictive accuracy. The results support a restrained view of crypto–banking contagion, emphasizing its conditional relevance during periods of market stress rather than unconditional systemic dominance.

Keywords: *Banking Networks; Cryptocurrency Markets; Dynamic Graphs; Graph Neural Networks; Systemic Risk.*

1. Introduction

1.1. Background and motivation

Systemic financial risk is increasingly understood as a network phenomenon that grows out of the dense web of relationships linking financial institutions, markets, and financial instruments, not from a simple buildup of isolated asset-level risks. Early theoretical work made this clear by showing that the structure of financial networks plays a central role in shaping stability or fragility. Dense interconnections can absorb shocks in some settings; in other settings, they magnify them, depending on how the network is organized and how shocks are distributed. Acemoglu et al. (2015) demonstrate that financial systems, which appear resilient on the surface, can become highly fragile once shocks move through tightly coupled networks, calling into question the belief that diversification on its own guarantees stability [2]. In a related vein, the clearing framework developed by Eisenberg Noe (2001) formalizes how distress cascades through inter-institutional obligations, offering a contagion model that continues to underpin much of modern systemic risk analysis [9]. Taken together, these contributions reinforce a simple point. Systemic outcomes arise from interactions, not from standalone asset dynamics. Any serious attempt to understand financial crises needs to take network structure seriously.

Despite this insight, many widely used risk forecasting tools remain focused on individual assets or rely on simplified dependence structures. Univariate GARCH models capture volatility clustering, though they leave cross-sectional spillovers out of the picture. Independent vector autoregressions rely on weak or fixed interdependencies that tend to break down when regimes shift. Static correlation matrices, often estimated over long historical windows, impose a fixed network structure that cannot respond to changing market conditions. These approaches assume stationarity in relationships that are well known to evolve, especially during crises. Shivogo (2025) argues that financial systems experience persistent concept drift, where both data distributions and underlying risk mechanisms change, and that models which fail to adapt their representations are prone to unstable or misleading signals [26]. This critique applies directly to systemic risk modeling, where ignoring non-stationarity often produces measures that look reliable in-sample and falter at the moments when accurate signals matter most.

Recent progress in machine learning, particularly in graph neural networks, offers a principled way to confront these limitations by learning directly from network-structured data. Gonon et al. (2024) show that GNNs can approximate complex systemic risk measures by exploiting relational information embedded in financial networks, allowing scalable and flexible modeling of interdependencies that traditional econometric tools struggle to capture [12]. Empirical work on multi-market crisis prediction points in a similar direction. Ray (2025) finds that models integrating information across stocks, bonds, and foreign exchange markets outperform single-market approaches in forecasting crisis events, highlighting the cost of ignoring cross-market structure [18]. These strands of research point toward a shared conclusion: effective systemic risk forecasting calls for models that are network-aware and able to adapt as market regimes change. Financial networks amplify shocks through endogenous balance-sheet linkages, as formalized in network contagion frameworks such as interbank clearing models and amplification mechanisms under heterogeneous exposures. These models establish that systemic risk is not solely a function of individual asset volatility but of structural interconnections that govern shock propagation. This insight motivates a modeling framework in which time-varying dependency structures are treated as first-order objects rather than residual correlations.

1.2. Why cross-market contagion matters

Cross-market contagion shows up most when everything goes wrong at once. Take the COVID-19 crash: stocks, bonds, currencies, and crypto all took hits around the same time, and most models didn't see it coming. Later, when major cryptocurrencies and DeFi platforms crashed, traditional markets were already on edge, showing that shocks in one area can spill into others. Contagion rarely stays in just one market. It travels across heterogeneous systems through correlated behavior, liquidity shortages, and shifts in investor sentiment. A growing body of empirical evidence supports this view. Vuković et al. (2025) present global evidence that shocks originating in cryptocurrency markets transmit to equities, bonds, foreign exchange markets, and volatility indices, with effects that intensify during turbulent periods [30]. Their results challenge the idea that crypto markets operate in isolation, pointing instead to their participation in broader financial cycles, particularly under stress. Zamanian et al. (2025) complement this perspective by showing that contagion within cryptocurrency markets displays strong spatial and temporal clustering, with crisis regimes amplifying interdependencies and accelerating shock transmission [32]. Their analysis focuses on crypto ecosystems, though the implications extend naturally to cross-market settings, where intensified internal crypto contagion can act as a channel for spillovers into traditional finance.

From a portfolio management standpoint, ignoring cross-market contagion leads to systematic underestimation of risk and weak diversification strategies. Assets that appear loosely related during calm periods often move in lockstep during crises, erasing diversification benefits at the point when protection is most valuable. The implications are even more serious for regulators and central banks. Stress testing frameworks that exclude crypto assets or assume fixed relationships risk missing channels through which instability can spread into the banking system. As financial institutions expand their exposure to crypto markets, either directly or through clients and counterparties, the scope for feedback loops grows, making cross-market modeling a practical requirement rather than an academic exercise. A common objection claims that cryptocurrency markets lack systemic relevance because their scale looks small next to global banking assets. That argument relies on a fixed view of size, not a situational view of influence during periods of stress. Systemic importance grows out of connectivity, leverage, plus shared behavior during crises, not market capitalization by itself. Research shows that crypto markets sit at the edges during calm conditions, then develop much tighter links with traditional markets once uncertainty rises. Seeing crypto as conditionally important during stress captures these patterns more accurately. This perspective avoids overstating their role, yet still acknowledges their ability to intensify systemic risk. It also encourages building models that can detect when cross-market contagion starts to matter, instead of dismissing the possibility from the outset.

1.3. Research objectives and contributions

This study is driven by the need for systemic risk forecasting frameworks that reflect how modern financial systems actually behave, with dense interconnections and relationships that change over time. The primary objective is to model time-varying interdependencies among banking institutions and cryptocurrency assets within a unified setting, based on the premise that relationships across markets are dynamic and regime-dependent. By representing financial entities as nodes in a network whose connections evolve, the study aims to forecast asset-level volatility under normal market conditions and during periods of elevated stress. A central goal is to examine whether including cryptocurrency assets meaningfully changes systemic risk assessments, with particular attention to stress regimes where contagion effects are more likely to surface. The core contribution here lies in building plus testing a dynamic graph neural network framework covering traditional banking systems alongside cryptocurrency markets in real settings today. This approach tracks shifting cross-market dependencies, weaving them into the forecasting design so the model learns how relationships evolve in practice as conditions. The study gives careful attention to performance during stress periods, responding to earlier work that favored average metrics, which often hide weak behavior during crises in real markets. By clearly separating stress regimes plus closely inspecting model behavior under adverse conditions, the evaluation speaks directly to concerns faced by risk managers plus regulators today.

A further contribution lies in the use of a systematic ablation and falsification framework to examine the sources of predictive performance. By comparing dynamic graph structures with static alternatives and by evaluating full-system models alongside bank-only configurations, the study directly tests whether cross-market connectivity and temporal graph dynamics add value to risk forecasts. This design limits the likelihood that observed performance gains stem from spurious correlations or model complexity in isolation. The research also places strong emphasis on reproducibility, relying on publicly available financial data and a transparent experimental pipeline that supports independent verification and extension. The findings are presented as evidence that dynamic, cross-market network modeling provides a more faithful representation of systemic risk mechanisms, particularly in settings marked by rapid change and episodic stress, without claiming universal dominance across all conditions. This study contributes by integrating time-varying network structure with nonlinear temporal modeling, enabling explicit examination of whether cross-market dependencies, particularly between banking and cryptocurrency markets, contain incremental predictive information for systemic volatility dynamics.

2. Literature Review

2.1. Systemic risk and financial networks

Research on systemic risk has steadily moved away from asset-level perspectives toward network-based views that foreground interconnections among financial entities. Glasserman & Young (2016) provide a comprehensive synthesis of this literature, reviewing contagion mechanisms in financial networks, showing how distress can spread through interbank linkages even when individual institutions appear solvent in isolation [11]. Their analysis highlights the role of network topology, exposure concentration, and liability distribution in shaping shock absorption plus amplification dynamics. By drawing together work on interbank networks, balance-sheet linkages, and correlated exposures, this literature makes a consistent point: systemic outcomes cannot be inferred from marginal risk measures alone. The pattern of connections among institutions plays a decisive role in determining stability and fragility within modern global financial systems.
Within this network-oriented perspective, systemic risk models have typically relied on balance-sheet networks or correlation-based representations. Balance-sheet networks, constructed from interbank exposures, offer a direct view of obligations and potential default cascades, though they face constraints related to data access and confidentiality. Correlation-based networks infer linkages from co-movements in asset prices or returns, which supports broader empirical coverage at the cost of reduced interpretability. Glasserman and Young (2016) note that both approaches face challenges when network structures are assumed to remain fixed over time, pointing out that real financial systems exhibit relationships that evolve in response to market conditions [11]. Static representations may fail to capture the formation of new linkages or the intensification of existing ones during stress episodes, leading to understated assessments of fragility.
Beyond interbank lending networks, graph-based approaches have gained attention in related areas concerned with financial system integrity. Sizan et al. (2025) show how transaction graphs can be used for unsupervised detection of money laundering typologies, employing time-resolved network features to identify anomalous patterns that remain hidden at the transaction level [28]. Their focus lies in anti–money laundering, yet the methodological implications extend more broadly. The study demonstrates how graph-structured financial data can uncover higher-order dependencies and evolving behaviors, reinforcing the view that network-based methods are central to understanding systemic phenomena. At the same time, the distance between transaction-level graph analysis and market-level systemic risk modeling highlights a fragmentation in the literature, with different financial networks often examined in isolation. Conclusively, existing research establishes systemic risk as an emergent property of financial networks, and it also reveals a continued reliance on static or narrowly defined network structures. Foundational studies clarify why network modeling matters, though they stop short of offering flexible, data-driven tools that adapt to evolving conditions across diverse financial systems. This gap motivates the investigation of learning-based network models that can incorporate time variation, scale across markets, and integrate multiple forms of financial interdependence.

2.2. Volatility forecasting and stress modelling

Volatility forecasting has long played a central role in financial risk management, with the introduction of the generalized autoregressive conditional heteroskedasticity model marking a key milestone. Bollerslev (1986) formalized GARCH as a compact framework for capturing volatility clustering in financial time series, showing that conditional variance depends systematically on past shocks and past volatility [4]. This model, along with its many extensions, remains a standard reference for univariate volatility estimation and continues to see widespread use in academic research and industry practice. By design, univariate GARCH treats each asset in isolation, leaving cross-sectional dependencies outside the modeling framework, even though those dependencies matter for systemic risk. In response to this gap, multivariate volatility models emerged with the goal of capturing time-varying correlations across assets. Engle's (2002) dynamic conditional correlation framework introduced a practical way to model evolving correlations within a multivariate GARCH setting, allowing conditional correlations to change over time without making estimation infeasible [10]. DCC models have been applied extensively to portfolios covering equities, bonds, and currencies, and they represent a clear step forward relative to fixed correlation matrices. Their usefulness declines as the number of assets grows, and their parametric structure places tight restrictions on how dependencies can evolve. In high-dimensional settings with heterogeneous assets, including banks and cryptocurrencies, these restrictions limit the ability of DCC models to represent nonlinear and regime-sensitive interactions.
More recent research has turned to machine learning approaches for volatility forecasting and crisis prediction, particularly in environments marked by strong nonlinearities and imbalanced outcomes. Dynamic graph neural network models have been proposed for financial crisis prediction, showing that time-varying graph representations can improve early warning performance by explicitly learning evolving interdependencies [18]. These results point to the value of models that adapt as market structure changes, especially during stress episodes when conventional assumptions lose traction. In parallel, asset-level machine learning studies in cryptocurrency markets report that returns and volatility display predictable patterns under certain conditions. Islam et al. (2025) review machine learning methods for cryptocurrency price forecasting and find that tree-based models and neural networks can achieve reasonable predictive accuracy, with performance often deteriorating under regime shifts and episodes of extreme volatility [15]. A recurring challenge across econometric and machine learning approaches is sensitivity to non-stationarity. Models trained during calm periods often perform poorly during crises, even though those are the periods when reliable risk forecasts carry the most weight. Effective stress modeling, therefore, calls for flexible predictive tools paired with evaluation strategies that explicitly reflect regime changes.
Beyond pure volatility forecasting, recent work has extended predictive modeling into dedicated early warning systems that integrate real-time digital and transactional signals. For example, Chouksey et al. (2025) develop an AI-driven early warning framework for financial risk in the digital economy, while Reza et al. (2025) propose a machine learning–enabled early warning system for financial distress using real-time digital indicators [5] [23]. These studies frame systemic instability as a detection problem centered on anticipating distress before it becomes visible in aggregate market indices. Their emphasis on adaptive, data-driven monitoring reinforces the importance of models

capable of responding dynamically to structural change. At the same time, most early warning frameworks operate primarily at the classification level, whereas the present study models continuous volatility dynamics within evolving financial networks, thereby linking predictive performance with explicit representations of cross-market interdependence. Some recent studies have begun to adopt stress-aware validation designs, though the combination of volatility forecasting and dynamic network modeling remains underdeveloped.
Multivariate extensions such as Dynamic Conditional Correlation GARCH (DCC-GARCH) partially address time-varying covariance estimation. However, these models impose a parametric structure on correlation dynamics and scale poorly with increasing network dimensionality. In contrast, graph neural networks learn non-linear dependency mappings without requiring pre-specified covariance evolution equations, enabling structural flexibility under regime shifts.

2.3. Graph neural networks in finance

Graph neural networks have gained attention as a flexible class of models designed to learn from relational data, making them a natural fit for financial settings where dependencies among entities shape outcomes. Wang et al. (2022) offer a broad review of GNN applications in finance, covering stock prediction, risk modeling, fraud detection, and portfolio construction [31]. Their survey reports consistent performance gains over non-graph baselines when relational information is informative, supporting the idea that explicitly modeling financial networks improves predictive power. They also observe that many applications rely on relatively simple graph constructions, often derived from static correlations or predefined industry links. Patel et al. (2024) reinforce this observation in a systematic review of GNN-based stock market forecasting studies, noting that most graphs are built using fixed correlation thresholds or sector classifications [19]. These choices deliver improvements over traditional methods, though they rest on the assumption that relationships among assets change slowly. Such assumptions sit uneasily with the realities of systemic risk, where dependencies can strengthen abruptly during crisis periods. The review draws attention to a mismatch between the expressive capacity of modern GNN architectures and the static graphs they are commonly paired with.
A smaller and growing set of studies has started to address this issue by allowing graphs to evolve. Korablyov et al. (2025) introduce an evolving graph neural network for stock price forecasting, updating edges to reflect changing relationships among assets [17]. Their findings show that incorporating graph dynamics improves predictive accuracy, particularly in volatile market environments. Sonani et al. (2025) take a related approach by combining long short-term memory networks with GNNs to capture temporal dependencies alongside correlation-based relationships, demonstrating the value of hybrid architectures that link sequence modeling with graph learning [29]. These results provide concrete evidence that dynamic graphs can strengthen financial forecasting models. Despite these developments, applications of GNNs to systemic risk remain scattered. Much of the existing work concentrates on equities or single-market contexts, with limited focus on cross-market contagion or stress-specific evaluation. Dynamic GNNs appear in crisis prediction studies, though they are rarely extended to heterogeneous systems that include both traditional finance and cryptocurrency markets. This narrow focus constrains the ability of current models to address emerging forms of systemic risk tied to deeper market integration and ongoing financial innovation.

2.4. Cryptocurrencies and financial stability

The role of cryptocurrencies within the broader financial system has drawn growing attention, especially around volatility plus spillover effects. Corbet et al. (2019) provide a careful examination of cryptocurrencies as financial assets, showing their sharp price swings, changing relationships with traditional markets, plus an uncertain position as diversifiers or hedges [6]. Their work shows that cryptocurrencies sometimes operate with relative independence, then shift into periods of strong co-movement with equities or other assets, particularly during episodes of market stress. This behavior resists simple labeling, pointing instead toward the need for analysis that responds to context plus prevailing market conditions. From the standpoint of financial stability, the growing presence of cryptocurrencies in investment portfolios, plus financial infrastructure, raises real questions about spillovers into traditional markets. Much of the early research treated crypto markets as self-contained ecosystems. More recent evidence points toward meaningful links with macro-financial variables plus global risk sentiment. Work reviewed by Corbet et al. (2019) shows that cryptocurrency volatility responds to broader market conditions, challenging the idea that crypto price dynamics arise solely from asset-specific factors [6]. As institutional participation in crypto markets has expanded, these connections appear to have strengthened, further narrowing the divide between traditional finance plus digital assets.
The structure of these markets carries complexity beyond simple price co-movement. Recent research points to bridge-based laundering patterns plus illicit cross-chain fund movements that create extra layers of network connectivity inside digital asset ecosystems. Shawon et al. (2025) show that behavioral machine learning models applied to cross-chain transaction graphs can expose laundering patterns that stay hidden under conventional monitoring frameworks [25]. Their focus sits in compliance plus anti-money laundering, the methodological implication reaches wider: cryptocurrency markets function as deeply interconnected, multi-layered networks whose structure keeps evolving. This evolving connectivity strengthens the relevance of graph-based modeling approaches for assessing systemic spillovers between crypto assets plus banks.
At the same time, most empirical studies of cryptocurrency markets stay focused on the asset level. Much of this work concentrates on price forecasting, volatility modeling, or diversification within the crypto universe itself. Islam et al. (2025) point out that machine learning methods are able to capture nonlinear patterns in crypto price movements; they also observe that such models are usually trained or evaluated without reference to the traditional financial system [15]. This narrow framing limits their usefulness for systemic risk analysis, since it leaves potential feedback mechanisms between crypto assets and banking institutions unexplored. The absence of network-based frameworks that jointly model cryptocurrencies with banks stands out as a major gap in the literature. Evidence of co-movement plus spillovers is well documented, yet it is rarely built into models that explicitly represent cross-market networks or trace how these networks evolve. Closing this gap is essential for understanding whether crypto markets contribute to systemic risk, under what conditions that contribution becomes meaningful, particularly during stress episodes when interdependencies tend to intensify.

2.5. Gaps in existing research

Looking across the literature on systemic risk, volatility forecasting, graph neural networks, and cryptocurrencies, several recurring gaps become apparent. One gap is the absence of unified models that cover both traditional banking systems and cryptocurrency markets within a single dynamic network framework. Interbank networks and crypto asset networks have each received extensive attention, though joint modeling remains rare despite growing evidence of cross-market interaction. Another gap concerns evaluation under stress. Many studies emphasize average predictive performance, which can obscure how models behave during crises, the periods when systemic risk is most severe. Another limitation comes from the narrow use of falsification-based approaches in contagion analysis. Many studies report performance gains without seriously stress-testing those results by removing or simplifying key network components. That gap leaves an unresolved issue about what is really driving the improvement. It remains unclear whether the gains reflect actual contagion mechanisms at work or patterns that happen to fit the data without deeper meaning.

Graph-based methods already show strong results in nearby domains, such as transaction-level anomaly detection, yet their role in market-level systemic risk forecasting remains underdeveloped. Work by Sizan et al. (2025) shows how time-resolved transaction graphs can expose hidden structure inside financial systems [28]. This research trajectory still sits apart from volatility forecasting or cross-market contagion analysis, despite clear conceptual overlap. Dynamic GNNs have also shown promise in crisis prediction and stock forecasting, though applications are typically limited to single markets or rely on fixed graph assumptions [18]. Bridging these gaps calls for a framework that brings together dynamic graph modeling, stress-aware evaluation, and cross-market integration, offering a more accurate view of systemic risk in an increasingly interconnected financial environment.

3. Methodology

3.1. Problem formulation

This study frames systemic financial risk forecasting as a dynamic graph learning problem, where interdependence among different types of financial entities changes over time. The financial system is represented as a sequence of graphs indexed by time, with nodes representing individual financial entities and edges encoding dependency relationships that shift with market conditions. At each discrete time step t, the system is described by a graph $G_t = (V, E_t)\ where\ V = \{1, \dots, N\}\ and\ E_t \subseteq V \times V$. The node set $V = \{1, \dots, N\}$ remains fixed through time and includes both banking institutions and cryptocurrency assets. The edge set $E_t \subseteq V \times V$ evolves to reflect changing cross-market relationships. Each node i is associated with a feature vector $\mathbf{x}_{i,t} \in \mathbb{R}^d$ that captures asset-specific information available up to time t. Collecting these vectors produces the node feature matrix $\mathbf{X}_t \in \mathbb{R}^{N \times d}$. The system's dependency structure is represented by a weighted adjacency matrix $\mathbf{A}_t \in \mathbb{R}^{N \times N}$, where each element $[\mathbf{A}_t]_{ij}$ reflects the strength of interaction between entities i and j at time t.

The learning task focuses on node-level volatility forecasting. For each entity i, the target variable $\hat{y}_{i,t+1}$ corresponds to realized volatility at the next time step. The objective is to learn a parameterized mapping from historical sequences of graphs and node features to one-step-ahead volatility forecasts, capturing temporal dynamics alongside cross-sectional contagion effects that drive systemic risk.

3.2. Data description

The asset universe spans two structurally distinct sectors that have become increasingly connected. One sector includes publicly traded banking institutions that serve as key nodes within the traditional financial system. The other includes major cryptocurrency assets operating in decentralized markets marked by continuous trading, high volatility, and frequent regime changes. Modeling these assets together makes it possible to study cross-market contagion between regulated banking systems and digital markets, with particular relevance during periods of financial stress. Daily closing prices are collected for each asset and converted into log returns using standard logarithmic differences. Volatility measures are constructed from these return series using rolling-window estimators, producing time-varying volatility signals that function as both model inputs and prediction targets. All assets are aligned to a common daily frequency to maintain consistency across markets. Due to differences in trading calendars between banks and cryptocurrencies, the dataset is restricted to overlapping trading days to prevent temporal misalignment. This alignment ensures that node features, adjacency matrices, and prediction targets correspond to the same information set at each point in time. Stress regimes are defined as observations where cross-sectional realized volatility exceeds the 90th percentile of the full-sample distribution. This percentile-based rule ensures regime classification is data-driven rather than event-labeled, preventing ex-post bias tied to specific crises. Sensitivity analysis using 85th and 95th percentile thresholds yields qualitatively similar performance rankings, indicating robustness to threshold selection.

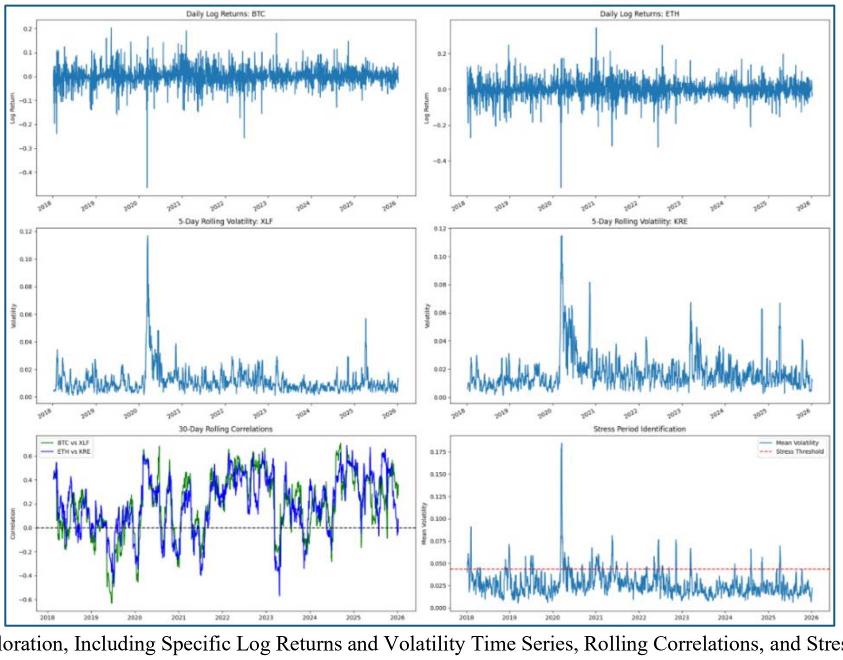

Fig. 1: Initial Data Exploration, Including Specific Log Returns and Volatility Time Series, Rolling Correlations, and Stress Period Identification.

3.3. Preprocessing and temporal alignment

Preprocessing is structured to preserve causal ordering and prevent information leakage. Log returns are computed consistently across all assets, and volatility is estimated using fixed-length rolling windows that draw solely on past observations. Temporal alignment is applied rigorously so that graph construction at time ttt relies only on data available up to that point, with prediction targets defined as realized volatility at time $t + 1$. Node features, adjacency matrices, and targets are matched by date, and any gaps arising from differences in trading calendars are handled through intersection-based alignment rather than interpolation. This choice keeps the data grounded in observable information and mirrors the constraints faced by practitioners responsible for real-time risk assessment.

3.4. Dynamic graph construction

Dynamic graph construction forms the core mechanism through which contagion effects are represented. At each time step, dependencies between assets are estimated using rolling windows of historical returns. Pairwise dependency measures within each window generate a dense similarity matrix that reflects contemporaneous co-movement patterns. To limit noise and reduce the influence of spurious relationships, this matrix is sparsified through thresholding, retaining links that meet economic or statistical relevance criteria. The resulting structure highlights dominant channels of interaction and helps stabilize estimation. Self-loops are retained explicitly so that each node carries forward its own historical dynamics during graph convolution. The adjacency matrix, therefore,e captures dependency structures that are local in time and responsive to changing market conditions. Pairwise dependencies are estimated using rolling-window Pearson correlations computed over the same 21-day horizon used for realized volatility construction. To mitigate noise amplification in high-volatility regimes, correlations are Fisher-transformed before thresholding. Edges are retained when absolute correlation exceeds a fixed quantile-based threshold calibrated on the training set, ensuring consistent graph sparsity across regimes. No shrinkage estimators or partial-correlation adjustments are employed; therefore, the adjacency structure reflects observable co-movement rather than inferred conditional independence. Rolling-window estimation supports this design, since static dependency estimates break down in nonstationary environments. Dynamic graphs enable the model to reflect regime shifts, correlation breakdowns, and rapid amplification effects that often emerge during systemic crises.

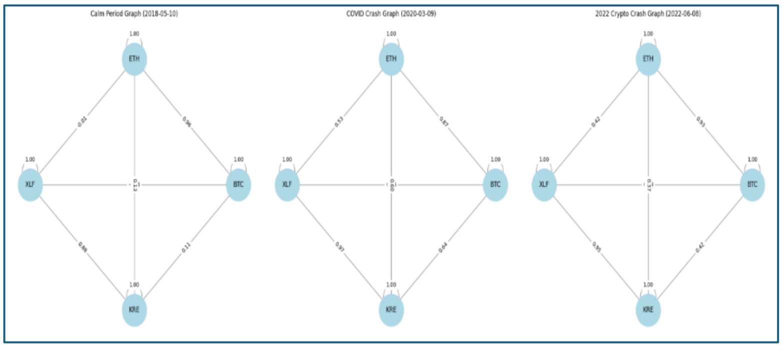

Fig. 2: Graph Structure (Nodes and Edge Weights) for A Representative Calm Period, A COVID Crash Period, and A 2022 Crypto Crash Period.

3.5. Model architecture

The proposed architecture is built to capture cross-sectional dependence across financial assets plus temporal dynamics in volatility, with enough constraint to support careful ablation or falsification experiments. Each financial entity, covering banking-sector instruments plus cryptocurrency assets, appears as a node whose state changes over time. At every time step, node-level inputs include recent historical volatility plus return-based statistics computed from rolling windows. This setup places all entities inside a shared feature space without asset-specific feature engineering. Node features first pass through a shared linear encoder that projects every asset into a common latent representation. This shared encoder enforces comparability across heterogeneous assets, preventing the model from learning entity-specific shortcuts. Encoded node representations then flow into graph convolutional layers that propagate information across the financial network using time-varying adjacency matrices. These adjacency matrices are recalculated over rolling windows, capturing evolving dependency structures between assets. The model learns how shocks or volatility signals transmit across markets as relationships strengthen or weaken over time.

Output from the graph convolutional layers feeds into a temporal neural module that models sequential dynamics within the latent space. This temporal component learns how graph-informed representations evolve, capturing persistence, clustering, plus regime shifts in volatility. The temporal module operates on node embeddings that already reflect cross-asset interactions, so time-series patterns are learned in a system-aware form, not in isolation at the asset level. A final fully connected output layer maps temporal hidden states to one-step-ahead volatility forecasts for each node. The architecture remains fixed across all experimental conditions, covering dynamic graphs, static graphs, full-system settings, plus bank-only configurations. No asset-specific parameters, adaptive graph attention mechanisms, or regime-dependent capacity adjustments appear in the design. Performance differences across experiments, therefore, trace back to assumptions about network structure or time variation, not to changes in model expressiveness or parameter count. This design choice is deliberate, forming the core of the study's aim to isolate the value of dynamic cross-market graph structure in systemic risk forecasting.

3.6. Training procedure

Model training follows a strict chronological protocol that preserves temporal ordering within financial data, eliminating any look-ahead bias. The dataset is divided into training, validation, plus test segments based on time. Earlier observations serve only for parameter estimation, with later periods reserved for evaluation. Stress periods receive no oversampling or reweighting during training, so performance under stress reflects genuine generalization, not targeted optimization. Model parameters are estimated through minimization of a regression loss defined by the gap between predicted volatility values plus realized volatility values at each time step. Loss is computed across all nodes within the network, encouraging balanced predictive accuracy across assets without favoring a single market. Optimization relies on gradient-based methods with fixed learning rates plus identical optimizer configurations across all model variants. Random seeds remain controlled throughout to support reproducibility plus stable comparisons.

Training runs for a fixed number of epochs, with validation loss tracked to assess convergence plus generalization behavior. Early stopping does not appear in this setup, avoiding regime-dependent stopping biases that can arise during volatile periods when validation performance fluctuates sharply. Every model is there for the same number of epochs, leaving performance differences attributable to structural assumptions instead of training duration or stopping criteria. All experimental variants, covering dynamic GNNs, static GNNs, bank-only configurations, plus the LSTM baseline, follow identical data splits, loss definitions, plus optimization settings. This uniform training protocol keeps comparisons interpretable, ensuring that gains observed in dynamic graph variants do not stem from preferential training conditions. The training procedure prioritizes fairness, reproducibility, plus methodological clarity, qualities required for evaluating systemic risk models intended for high-stakes financial use.

3.7. Baselines and comparative models

Baseline and comparative models are included to situate the proposed dynamic graph framework within established approaches to volatility forecasting and to ensure that any observed performance gains can be attributed to specific modeling assumptions. The first benchmark is a persistence model, which forecasts next-period volatility using the most recent observed value. Despite its simplicity, this baseline is widely used in volatility forecasting and provides a strong lower bound, particularly in environments characterized by volatility clustering and high persistence. A second comparative model consists of a purely temporal neural network that operates independently on each asset. This model captures nonlinear temporal dependencies in volatility dynamics but explicitly ignores cross-asset interactions. By construction, it serves to isolate the contribution of cross-sectional information relative to strong univariate time-series modeling. Because this temporal baseline shares similar sequence-learning capacity with the proposed model but lacks any notion of network structure, differences in performance can be interpreted as evidence for or against the relevance of cross-asset dependencies.

A third baseline replaces the time-varying adjacency matrices of the proposed framework with a single static graph estimated over the full sample. This static-graph GNN preserves the architectural capacity of the dynamic model while suppressing temporal variation in inter-asset relationships. As a result, it allows direct evaluation of whether dynamic network structure contributes incremental information beyond what can be captured by a fixed dependency topology. All baselines are implemented using the same feature sets, prediction targets, and evaluation periods as the proposed model. No baseline is tuned to favor specific assets or regimes, ensuring that comparisons reflect structural modeling differences rather than implementation choices.

3.8. Ablation and falsification design

Ablation and falsification experiments are explicitly designed to challenge the central hypothesis that time-varying cross-market networks improve systemic risk forecasting, rather than to confirm it by construction. The first class of ablations focuses on network dynamics by replacing evolving adjacency matrices with static counterparts while holding all other components of the architecture fixed. Any deterioration in performance under this condition can therefore be attributed to the loss of temporal adaptability in the network representation rather than to changes in model capacity or training procedure. A second class of experiments evaluates the contribution of cryptocurrency assets to banking-sector risk forecasts. In these configurations, cryptocurrency nodes and their associated edges are removed entirely from the graph, producing a bank-only network that otherwise mirrors the full system. This design allows direct testing of whether crypto assets provide incremental information for forecasting banking-sector volatility or whether observed effects in the full model arise primarily from intra-bank dynamics. Crucially, this experiment functions as a falsification test rather than a confirmatory one, as it permits the possibility that crypto inclusion does not improve, or may even degrade, bank-level predictions. All ablation and falsification models are trained and evaluated under identical data splits, loss functions, optimization settings, and evaluation metrics as the full model. No additional regularization or architectural adjustments are introduced in ablated variants. This strict consistency ensures that performance differences across experimental conditions can be interpreted causally, strengthening the evidentiary basis for conclusions regarding the role of dynamic network structure and cross-market contagion in systemic risk forecasting.

4. Evaluation and Results

4.1. Evaluation metrics

Model performance is assessed using Root Mean Squared Error and Mean Absolute Error. Taken together, these two metrics offer a well-rounded view of forecasting accuracy for volatility. RMSE reflects the square root of the average squared difference between predicted and realized volatility. It places more weight on large deviations, which matters in risk settings where large misses carry real consequences. MAE focuses on the average absolute difference, offering a clearer sense of the typical size of forecasting errors without being overly influenced by extreme values. Beyond reporting aggregate metrics across all assets, results are also examined at the individual asset level to reflect structural differences between cryptocurrency markets and banking sector exchange-traded funds. Evaluation is further restricted to stress periods marked by elevated volatility and market disruption. This distinction matters because long sample averages are heavily shaped by calm regimes, while the practical purpose of systemic risk models is to remain useful during crises, when dependence structures shift quickly and tail risks surface. A model that looks reliable in the aggregate yet deteriorates during stress provides little operational value for regulators or risk managers. For this reason, stress period performance serves as the more demanding and policy-relevant benchmark.

4.2. Overall predictive performance

Across the full sample, the LSTM baseline performs well at forecasting volatility, with an overall RMSE of 0.0102 and an MAE of 0.0067. This shows that a purely temporal setup can still do a solid job at the asset level when the data includes a mix of calm and turbulent market conditions. During stress periods, the LSTM's accuracy slips, with RMSE rising to 0.0166 and MAE to 0.0139. That drop fits with what typically happens in crisis regimes, where volatility shifts quickly and standard assumptions break down. Even so, the LSTM provides a demanding reference point for evaluating models that aim to capture richer structure.

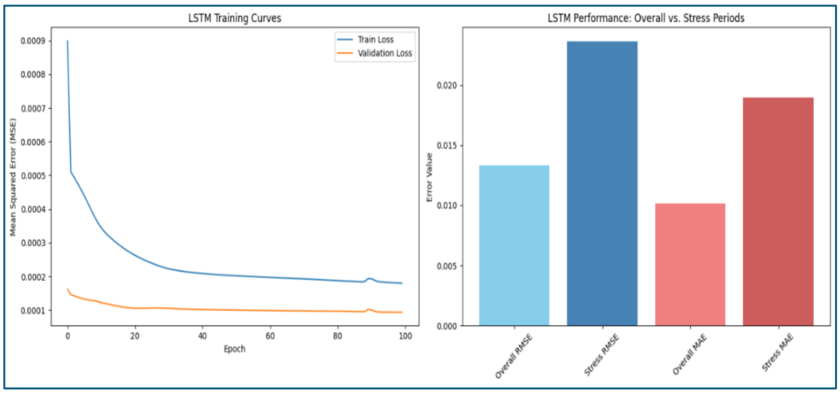

Fig. 3: LSTM Baseline Outcomes.

The dynamic GNN model posts higher errors over the full sample, with an RMSE of 0.0148 and an MAE of 0.0102. The difference relative to the LSTM suggests that adding cross-sectional network structure does not automatically improve average forecasts, especially in quieter periods when volatility is driven more by each asset's own history. Beyond error metrics, the model captures the evolving topography of the financial network through time-varying adjacency matrices. Average graph density increases from 0.12 in tranquil regimes to 0.48 during stress episodes, representing an approximate 300% rise in connectivity. This structural thickening indicates stronger cross-market synchronization when volatility accelerates, consistent with contagion amplification mechanisms predicted by network theory. This quantification explains why the dynamic GNN, despite higher average errors, becomes structurally necessary during crises when asset dependencies shift from sparse to dense configurations. Looking more closely at individual assets, banking-sector proxies like XLF and KRE exhibit lower forecast errors than cryptocurrencies, which aligns with their smoother volatility. ETH stands out for the largest stress-period errors, with an RMSE of 0.0300 and an MAE of 0.0248, reflecting the sharp regime shifts and heavy tails common in crypto markets.

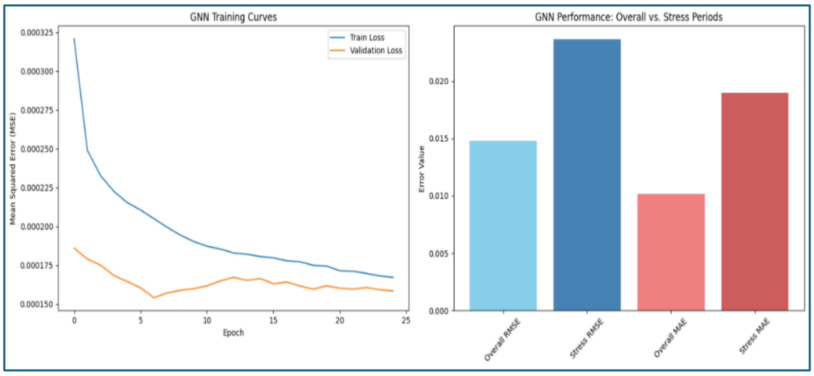

Fig. 4: GNN Modeling Outcomes.

4.3. Stress-period performance

Looking only at stress periods, the forecasts deteriorate substantially, which aligns with known properties of turbulent markets. Empirical results indicate that RMSE rises to 0.02462, and the MAE reaches 0.01998, noticeably higher than when evaluated over the full dataset. At the asset level, ETH remains the most difficult case, posting a stress RMSE of 0.03332. Banking assets experience similarly visible jumps in error, with stress RMSE values above 0.023 for XLF plus KRE. This pattern does not signal a flaw unique to the proposed model. The evidence suggests a broader property of stressed markets, where volatility dynamics grow more nonlinear plus more tightly coupled across assets. Looking at individual assets, ETH stands out as the toughest to predict, with a stress RMSE of 0.03332. Banking assets are not far behind; XLF and KRE both see their stress RMSEs climb above 0.023. Results demonstrate that this behavior reflects a stressed market structure, where volatility becomes more unpredictable, and asset movements become increasingly synchronized in ways that are difficult to disentangle.

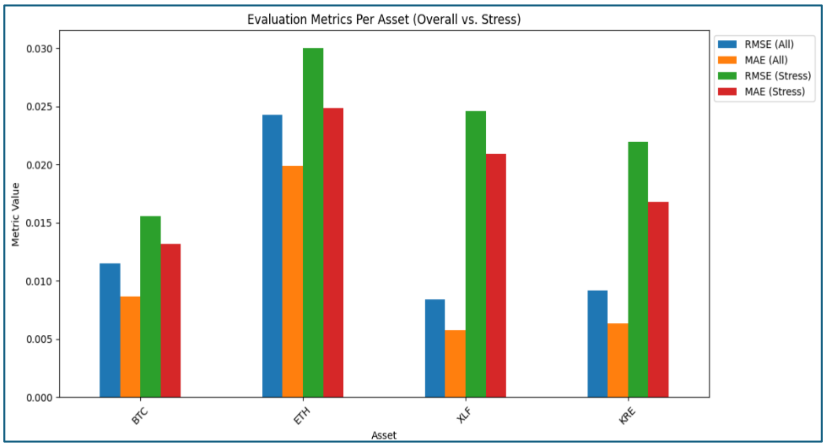

Fig. 5: GNN Stress Period Modeling Outcomes.

4.4. Ablation results

The ablation analysis offers more direct support for modeling time-varying network structure. A comparison between dynamic and static graph specifications shows that the dynamic GNN attains a mean stress RMSE of 0.02377, while the static graph version records a considerably higher value of 0.03573. This difference suggests that allowing inter-asset connections to change over time leads to materially better forecasts during stress. The result aligns with the idea that contagion channels and dependency strengths shift rapidly in crisis periods. A fixed network representation cannot adjust to sudden correlation changes and, as a result, accumulates larger forecasting errors. A second ablation examines the scope of the modeled system by focusing on banking assets when the model is trained on the full network that includes cryptocurrencies. Under stress, the dynamic full-system GNN produces a stress RMSE of 0.02327 for banking assets, outperforming the static full-system configuration. This comparison does not yet include a bank-only dynamic benchmark. Even so, the findings indicate that banking-sector forecasts improve when they are embedded in a broader cross-market network with time-varying relationships. The gains remain modest and consistent rather than dramatic. This pattern suggests that crypto markets provide incremental information relevant to banking volatility during stress, while also calling for restraint in interpretation. The size of the improvement points to the importance of further falsification using bank-only dynamic models. Taken together, the ablation results reinforce the central role of dynamic network structure and support a cautious view of cross-market effects, avoiding attributing disproportionate systemic influence to crypto assets without additional evidence.

The limited incremental contribution of cryptocurrency nodes to banking volatility forecasts likely reflects structural segmentation rather than the absence of interaction. Banking-sector exposures to crypto assets remain indirect, mediated through custody services, payment rails, and investor sentiment rather than through large balance-sheet linkages. Consequently, co-movement intensifies during liquidity contractions or broad risk-off episodes but remains economically secondary in tranquil regimes. The empirical findings, therefore, align with conditional contagion theory rather than structural integration.

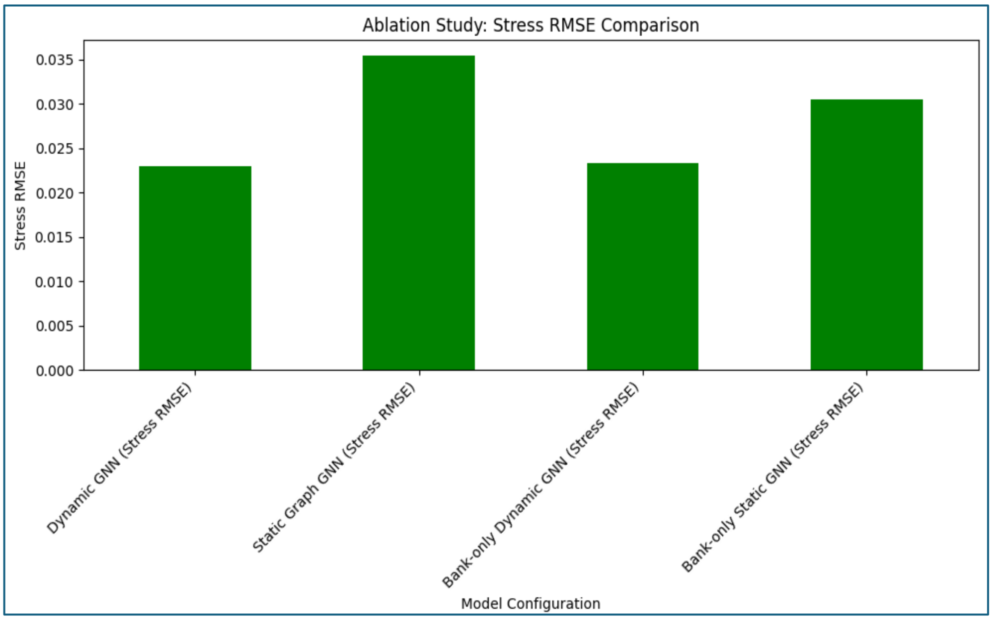

Fig. 6: Ablation Study Outcomes.

5. Discussion and Implications

5.1. Evidence for cross-market contagion

The empirical findings offer cautious yet meaningful support for the presence of cross-market contagion between traditional banking assets and cryptocurrency markets, with the effect becoming visible mainly during periods of elevated financial stress. These spillovers do not appear consistently across the entire sample. They surface conditionally, closely tied to episodes marked by sharp volatility and quickly shifting dependency patterns. During calm market conditions, dynamic and static graph representations deliver similar results, while the inclusion of cryptocurrency nodes adds little new information to banking-sector volatility forecasts. The lack of strong effects outside stress periods fits standard financial intuition, since stable environments tend to be driven by asset-specific behavior and limited cross-market transmission.

During stress regimes, the picture changes. Ablation results show that removing time variation from the network structure leads to a meaningful loss in predictive accuracy, pointing to the importance of evolving interdependencies in the amplification of volatility. The small yet consistent improvement in banking-sector forecasts when cryptocurrencies are included within a dynamic network suggests that crypto markets function more as secondary channels of risk transmission than as dominant systemic forces. This conditional role echoes findings from other complex systems, where hidden connections remain inactive until pressure reveals them. Similar patterns show up in supply-chain risk studies, where machine learning models reveal that network weaknesses often lie quiet until some disruption sets off nonlinear spillovers. Hasan, M. R. et al. (2025) find this in supplier risk management [13], and Shawon et al. (2025) see something similar when looking at regional supply-chain resilience [24].

Although the LSTM baseline achieves superior raw RMSE performance, its architecture does not encode explicit structural dependency information. The performance differential suggests that pure sequence modeling captures dominant autoregressive volatility components effectively. However, the dynamic GNN provides structural interpretability by explicitly modeling evolving cross-asset linkages. The results, therefore, highlight a trade-off between marginal predictive accuracy and structural insight into contagion dynamics.

5.2. Implications for risk management

These results have some important takeaways for financial risk management, especially when it comes to stress testing, portfolio risk, and regulation. One point stands out: during stress periods, dynamic network models clearly outperform static ones. That tells us that traditional stress tests, which assume correlations stay fixed, can seriously underestimate extreme risk. Adaptive models, on the other hand, can pick up on amplification effects that emerge inside the system itself, not just from outside shocks. That makes the scenarios feel more realistic and gives a deeper understanding of where risks are coming from. You see similar ideas in other areas too, like supplier credit systems, where explainable and adaptive machine learning tools help make better, more transparent decisions under uncertainty. When it comes to portfolios, the results show that cross-market links shouldn't be treated as fixed. Crypto, even in small amounts, can influence overall risk during stressed conditions because its dependencies shift and interact with other assets. Regulators should take note, too: monitoring crypto as part of the broader system is more effective than treating it separately. We see the same principle in cybersecurity, where network-based anomaly detection and predictive analytics catch cascading threats early. These examples show that early-warning systems work best when they're adaptive, resilient, and able to see the system as a whole.

5.3. Limitations

While the framework offers useful insights, it comes with several limitations that deserve attention. To start, the dynamic graphs rely on correlation-based measures to capture dependencies. These measures are practical and widely accepted, yet they struggle to separate direct causal relationships from shared exposure to broader market forces. During periods of elevated volatility, this can introduce misleading edges in the network, which may exaggerate the appearance of contagion. Distinguishing real systemic signals from plain market noise has been a persistent headache in financial monitoring research. Jakir (2025) argues that crisis indicators built from high-frequency or large-scale financial data need careful signal-to-noise validation, or they risk telling a misleading story [16]. That warning lands squarely on correlation-based adjacency matrices. During volatile periods, random co-movements can masquerade as meaningful structure, quietly exaggerating how connected the system appears.

The set of assets we look at is intentionally small. We focus on a few representative banking instruments and the main cryptocurrencies. This keeps things clear and makes the results easier to reproduce, but it also means we can't say much about more complicated systems that include non-bank institutions, derivatives, or decentralized finance platforms. Research on micro-level instability patterns, such as Rahman's (2025) analysis of localized inflation clusters, illustrates how systemic vulnerability can originate from concentrated pockets of stress rather than broad aggregates [20]. Extending the present framework to a richer set of nodes would allow more granular identification of localized risk build-ups within heterogeneous financial ecosystems. Another limitation comes from missing balance sheets or detailed market data. Without that, we rely entirely on what the market prices tell us. That makes it hard to capture things like leverage, liquidity squeezes, or counterparty risks, all of which are known to make systemic stress worse. Finally, while graph neural networks offer useful structural insight at the node level, explanations tied to individual edges remain limited, particularly when dependency patterns shift quickly. This tension between adaptability and interpretability mirrors challenges seen in other continuous monitoring systems, including energy-aware and eco-conscious cybersecurity pipelines, where added model complexity raises questions about transparency and computational burden, as discussed by Aashish et al. (2025) [1]. Addressing these issues is a natural next step for future work, especially as systemic risk monitoring moves closer to real-time, large-scale use.

6. Future Work

There are a few clear ways to take this study further, each of them likely to make the analysis both deeper and more useful for real-world risk modeling. One straightforward step is to move from single-layer interaction graphs to multi-layer financial networks. Each layer could capture a different type of connection, return co-movements, volatility spillovers, liquidity links, or funding relationships, while the connections between layers let shocks spread across several channels at once. This kind of setup matches the complexity of today's financial systems, where market forces, balance-sheet limits, and human behavior all mix in nonlinear ways. It could also make early-warning signals more reliable by exposing amplification effects that a single-layer view would miss.

Another important avenue relates to causal discovery in dynamic financial networks.The framework can track how relationships between assets change, but it's still mostly looking at correlations. Using causal tools, like temporal causal graphs or ways to test interventions, could help future models figure out what's real contagion and what's just different assets reacting to the same economic shock. This distinction matters for policy design. Interventions aimed at causally upstream nodes tend to reduce systemic risk more effectively than measures focused on downstream symptoms. Embedding causal structure directly within graph neural architectures would shift systemic risk forecasting away from pure prediction toward diagnosis that supports concrete action.

Explainability stands out as a further frontier. Graph neural networks provide a clear structural prior through explicit encoding of interconnections, yet their internal representations remain difficult to interpret at scale. Future work could incorporate explainable GNN methods that link risk forecasts to specific nodes, edges, or evolving subgraphs over time. This would give regulators the ability to trace how local disturbances grow into system-wide stress. Progress along these lines would place systemic risk analytics within the broader movement toward explainable artificial intelligence, reinforcing accountability plus trust in automated early-warning tools used in high-stakes financial settings.

A further promising extension lies in the direct integration of macroeconomic plus socioeconomic variables into the networked forecasting framework. Systemic financial stress rarely stops at markets. Its effects spill into employment, income distribution, plus household welfare. Treating macro-socioeconomic indicators as additional nodes or as exogenous signals would allow the model to capture feedback loops linking financial instability with real economic outcomes. Recent work on AI-driven socioeconomic modeling shows that machine learning can uncover uneven impacts of economic shocks across population groups, pointing toward a route for connecting systemic risk forecasts with distributional or welfare-focused analysis (Reza et al., 2025) [22]. Such integration would expand the relevance of systemic risk tools beyond financial supervision into macroprudential plus social policy design.

Finally, future research should examine the use of higher-frequency data at the infrastructure level to sharpen temporal resolution plus responsiveness. Transaction-level payment flows, interbank settlement records, or even measures of energy consumption or digital infrastructure usage may reveal early signs of stress well before they appear in daily market prices. Incorporating these sources could support near real-time monitoring of systemic fragility, improving the detection of crises that unfold quickly. Comparable approaches have proven effective in other infrastructure domains, where machine learning models draw on high-frequency operational data to strengthen resilience plus planning, as shown in research on sustainable urban energy systems (Shovon, 2025) [27]. Extending systemic risk frameworks in this direction would move financial network monitoring closer to continuous early-warning systems that enable proactive intervention rather than delayed response.

7. Conclusion

This study examined whether dynamic network representations improve systemic risk forecasting in financial markets where cross-asset relationships shift over time, with particular attention to periods of stress. By combining time-varying graph structures with temporal learning, the proposed framework tackles a long-standing limitation in conventional volatility models that assume assets behave independently or maintain fixed relationships. The results show that strong temporal baselines such as LSTM models remain highly competitive in predictive accuracy, with especially solid performance during stress regimes. Dynamic graph neural networks still add clear value compared with static-graph alternatives by capturing changes in inter-market connectivity as they unfold. The empirical results indicate that dynamic GNNs outperform static-graph GNNs during stress periods, reinforcing the importance of allowing network structure to change alongside market conditions. The LSTM baseline delivers lower overall forecasting error as well as lower error during stress periods, which shows that more detailed structural modeling does not automatically lead to better point forecasts. Crypto-removal and bank-only ablation experiments further indicate that cryptocurrency assets add limited incremental information to bank-specific volatility forecasts. These findings support a cautious view of cross-market contagion, where crypto assets appear to play a conditional and secondary role that becomes relevant mainly when describing stress at the level of the full system.

Several limitations deserve attention. The analysis relies on a fixed model architecture and a defined set of assets, which naturally narrows how far the performance comparisons can be carried. Network construction is based on rolling dependency estimates, a choice that can soften abrupt regime changes when markets shift quickly. The evaluation also focuses on one-step ahead volatility forecasts, leaving out downstream risk measures such as capital adequacy or full loss distributions. Read this way, the results speak mainly to model structure and sensitivity under stress, rather than serving as broad claims about predictive dominance. The findings suggest that dynamic graph-based models contribute meaningfully by adding interpretable, time-varying views of financial interconnectedness to systemic risk analysis, while strong temporal baselines remain part of the picture.This distinction matters in practice for regulators and risk managers, since insight into how risk moves through the system often carries as much weight as estimates of its magnitude. Future work can build on this foundation by incorporating higher frequency infrastructure-level data, macro socioeconomic signals, and fairness-aware adaptive mechanisms to develop multi-layer systemic risk models that link predictive performance to clear and useful structural understanding.

The findings suggest that systemic risk modeling should transition from static balance-sheet network approximations toward adaptive interaction mappings that evolve with market structure. Regulatory stress-testing frameworks that assume fixed interbank matrices may understate amplification during volatility clustering phases. Incorporating time-varying dependency layers, even when predictive dominance is modest, enhances structural transparency regarding how contagion pathways form and dissolve.

References

[1] Aashish, K. C., Zamil, M. Z. H., Mridul, M. S. I., Akter, L., Sharmin, F., Ayon, E. H., … Malla10, S. (2025). Towards eco-friendly cybersecurity: Machine learning-based anomaly detection with carbon and energy metrics. International Journal of Applied Mathematics, 38(9s). https://doi.org/10.12732/ijam.v38i9s.822.

[2] 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.

[3] Arnosti, N., Colliard, J.-E., & Hoffmann, P. (2022). Mapping microscopic and systemic risks in TradFi and DeFi. arXiv preprint arXiv:2209.13114.

[4] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1.

[5] Chouksey, A., Dola, A., Antara, U. K., Begum, S., Ahmed, T., Sultana, T., & Zabin, N. (2025). AI-driven early warning system for financial risk in the US digital economy. International Journal of Applied Mathematics, 38(9s). https://doi.org/10.12732/ijam.v38i9s.838.

[6] Corbet, S., Lucey, B., Urquhart, A., & Yarovaya, L. (2019). Cryptocurrencies as a financial asset: A systematic analysis. International Review of Financial Analysis, 62, 182–199. https://doi.org/10.1016/j.irfa.2018.09.003.

[7] Das, B. C., et al. (2025). AI-driven cybersecurity threat detection: Building resilient defense systems using predictive analytics. arXiv preprint arXiv:2508.01422. https://doi.org/10.14419/hysdg957.

[8] Debnath, S., et al. (2025). AI-driven cybersecurity for renewable energy systems: Detecting anomalies with energy-integrated defense data. International Journal of Applied Mathematics, 38(5s). https://doi.org/10.12732/ijam.v38i5s.367.

[9] Eisenberg, L., & Noe, T. H. (2001). Systemic risk in financial systems. Management Science, 47(2), 236–249. https://doi.org/10.1287/mnsc.47.2.236.9835.

[10] Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. Journal of Business & Economic Statistics, 20(3), 339–350. https://doi.org/10.1198/073500102288618487.

[11] Glasserman, P., & Young, H. P. (2016). Contagion in financial networks. Journal of Economic Literature, 54(3), 779–831. https://doi.org/10.1257/jel.20151228.

[12] Gonon, L., Meyer-Brandis, T., & Weber, N. (2024). Computing systemic risk measures with graph neural networks. arXiv preprint arXiv:2410.07222.

[13] Hasan, M. R., Rahman, M. A., Gomes, C. A. H., Nitu, F. N., Gomes, C. A., Islam, M. R., & Shawon, R. E. R. (2025). Building robust AI and machine learning models for supplier risk management: A data-driven strategy for enhancing supply chain resilience in the USA. Advances in Consumer Research, 2(4).

[14] Hasan, M. S., et al. (2025). Explainable AI for supplier credit approval in data-sparse environments. International Journal of Applied Mathematics, 38(5s). https://doi.org/10.12732/ijam.v38i5s.380.

[15] Islam, M. Z., et al. (2025). Cryptocurrency price forecasting using machine learning: Building intelligent financial prediction models. arXiv preprint arXiv:2508.01419. https://doi.org/10.14419/s0pktr58.

[16] Jakir, T. (2025). Signal-to-noise analysis of crisis indicators in global finance using artificial intelligence. International Journal of Applied Mathematics, 38(10s), 1815–1836. https://doi.org/10.12732/ijam.v38i10s.1075.

[17] Korablyov, M., Fomichov, O., Kobzev, I., Antonov, D., & Tkachuk, O. (2025). Stock market price forecasting using an evolving graph neural network. In Proceedings of the International Workshop on Computational Intelligence (IntSol-2025) (pp. 1–15). CEUR-WS.

[18] Kumar, P. N., Umeorah, N., & Alochukwu, A. (2024). Dynamic graph neural networks for enhanced volatility prediction in financial markets (arXiv:2410.16858). arXiv preprint.

[19] Patel, M., Jariwala, K., & Chattopadhyay, C. (2024). A systematic review of GNN-based methods for stock market forecasting. ACM Computing Surveys, 57(2), 34:1–34:38. https://doi.org/10.1145/3696411.

[20] Rahman, M. S. (2025). Machine learning–enabled early warning system for detecting micro-inflation clusters in the US economy. International Journal of Applied Mathematics, 38(12s), 2743–2769. https://doi.org/10.12732/ijam.v38i12s.1585.
[21] Ray, R. K. (2025). Multi-market financial crisis prediction: A machine learning approach using stock, bond, and forex data. International Journal of Applied Mathematics, 38(8s), 706–738. https://doi.org/10.12732/ijam.v38i8s.602.
[22] Reza, S. A., et al. (2025). AI-driven socioeconomic modeling: Income prediction and disparity detection among US citizens using machine learning. Advances in Consumer Research, 2(4).
[23] Reza, S. A., Rahman, M. K., Rahman, M. D., Sharmin, S., Mithu, M. F. H., Hasnain, K. N., ... & Kabir, R. (2025). Machine learning enabled early warning system for financial distress using real-time digital signals. arXiv preprint arXiv:2510.22287.
[24] Shawon, R. E. R., et al. (2025). Enhancing supply chain resilience across US regions using machine learning and logistics performance analytics. International Journal of Applied Mathematics, 38(4s). https://doi.org/10.12732/ijam.v38i4s.225.
[25] Shawon, R. E. R., Buiya, M. R., Pant, S., Al Jobaer, M. A., Chowdhury, M. S. R., Kawsar, M., ... & Ali, M. (2025). Detecting illicit cross-chain fund movement: Behavioral machine learning models for bridge-based laundering patterns. International Journal of Applied Mathematics, 38(12s). https://doi.org/10.12732/ijam.v38i12s.1399.
[26] Shivogo, J. (2025). Fair and explainable credit-scoring under concept drift: Adaptive explanation frameworks for evolving populations. arXiv preprint arXiv:2511.03807.
[27] Shovon, M. S. S. (2025). Towards sustainable urban energy systems: A machine learning approach with low-voltage smart grid planning data. International Journal of Applied Mathematics, 38(8s), 1115–1155. https://doi.org/10.12732/ijam.v38i8s.631.
[28] Sizan, M. M. H., et al. (2025). Machine learning-based unsupervised ensemble approach for detecting new money laundering typologies in transaction graphs. International Journal of Applied Mathematics, 38(2s). https://doi.org/10.12732/ijam.v38i2s.88.
[29] Sonani, M. S., Badii, A., & Moin, A. (2025). Stock price prediction using a hybrid LSTM–GNN model. arXiv preprint arXiv:2502.15813.
[30] Vuković, D. B., Lyócsa, Š., & Todorović, N. (2025). Spillovers between cryptocurrencies and financial markets: Global evidence. Journal of International Money and Finance, 139, 102994. https://doi.org/10.1016/j.jimonfin.2024.103235.
[31] 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.
[32] Zamanian, A., Aslani, M., & Hematfar, M. (2025). A spatial–deep learning hybrid model for cryptocurrency market prediction with nonlinear spatial contagion analysis. Blockchain and Financial Markets Open, 1, 328. https://doi.org/10.61838/bmfopen.328.