CIA Open Access

Identifying Cross-Market Risk Contagion Amplifiers via Graph Attention Networks: Empirical Evidence from U.S. Financial Stress Periods

***Yifei Li*[1]*, Fanyi Zhao*[1.2]*, Jiacheng Hu*[2]**

[1]*Master of Science in Enterprise Risk Management, Columbia University, NY, USA*

[1.2]*Computer Science, Stevens Institute of Technology, NJ, USA*

[2]*Master's Degree in Information Technology, University of New South Wales, Australia*

DOI: 10.63575/CIA.2026.40114

Abstract

The propagation of financial distress across heterogeneous asset markets has become one of the most pressing challenges for systemic risk monitoring in the post-crisis regulatory environment. This paper investigates cross-market risk contagion among U.S. equity, bond, and credit derivative markets through a graph attention network (GAT) framework embedded in a multilayer network structure. Constructing a directed, attention-weighted graph with 30 nodes spanning three market layers, and drawing on daily data from January 2018 to December 2023, the analysis identifies a persistent set of risk amplifier nodes whose elevated outgoing attention scores and centrality measures designate them as the primary drivers of cross-market stress transmission. The high-yield CDS index node emerges as the most consistently identified amplifier across three distinct stress episodes—the March 2020 pandemic shock, the 2022 interest rate stress cycle, and the March 2023 regional banking turmoil—while the financial sector equity index and long-duration bond node occupy episode-specific secondary amplifier roles. The multilayer network representation captures contagion pathways that single-layer analyses systematically underestimate, particularly the derivative-to-equity transmission channel active under tail-risk conditions. These findings carry direct implications for counterparty risk monitoring and macro-prudential early-warning design in the context of central clearing and financial stability regulation.

Keywords: *cross-market risk contagion; graph attention network; multilayer financial network; systemic risk amplifier; financial stress*

1. Introduction

1.1 Research Background and Motivation

The structural interconnectedness of contemporary financial markets has made systemic risk a fundamentally cross-market phenomenon. Shocks originating in credit derivative products may cascade through equity markets and interbank funding channels within days; the repricing of long-duration bonds may trigger simultaneous dislocations in leveraged loan and CDS markets through margin and collateral channels that conventional pairwise models are structurally ill-equipped to trace. The collapse of Lehman Brothers in September 2008 demonstrated these dynamics with devastating clarity, and the March 2020 pandemic selloff compressed years of potential credit deterioration into a single month of synchronized repricing across all major asset classes. These experiences have sustained substantial regulatory and academic interest in analytical tools capable of detecting cross-market contagion in near real time, mapping its structural drivers, and providing actionable signals for intervention before full-scale dislocations overwhelm the capacity of conventional risk limits and margin systems to contain them.

The application of graph neural networks to financial contagion detection has opened genuinely new analytical avenues for this challenge. GNN-based approaches jointly exploit the topological structure of financial interconnections and the temporal dynamics of asset return series, producing risk representations that capture propagation effects that regression-based or correlation-based methods systematically miss. Recent work employing gated graph neural networks for industry-level systemic risk warning in U.S. equity markets has demonstrated that network topology carries a strong and predictive correlation with financial crisis events during critical market periods [1], while attention-based graph architectures applied to CDS spread forecasting have documented improvements in out-of-sample predictive accuracy exceeding fifty percent over conventional machine learning baselines [2]. The present paper extends the GNN-based contagion analysis framework to the multi-asset, cross-market setting, constructing a directed multilayer network spanning U.S. equity, bond, and credit derivative markets with attention mechanisms providing the dynamic, directional edge weights through which risk amplifiers are systematically identified.

The institutional context motivating this study is the growing deployment of AI-driven surveillance tools by financial market infrastructure providers operating at the center of counterparty exposure networks. The analytical framework developed here aligns directly with the risk monitoring mandate of central clearing institutions: by mapping which market nodes concentrate and amplify incoming risk signals under stress, the framework can inform both the calibration of exposure limits and the design of early-warning triggers that activate before full-scale dislocations become apparent in standard price series. The 2018–2023 sample encompasses three well-documented stress episodes—the COVID-19 shock, the 2022 rate-driven credit stress, and the March 2023 regional banking disruption—providing rich empirical variation across distinct contagion scenarios. The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on cross-market contagion measurement and graph neural network applications in finance. Section 3 describes the data, network construction procedure, and attention-based detection approach. Section 4 presents the empirical findings on contagion topology, transmission paths, and amplifier node identification. Section 5 discusses theoretical and regulatory implications.

1.2 Research Scope and Key Contributions

A. Scope and Research Questions

The study covers U.S. financial markets over the period January 2018 to December 2023, focusing on three asset-class layers: equity sector indices, bond sub-indices differentiated by maturity and credit quality, and CDS index series representing credit derivative market conditions. Three research questions structure the empirical work. The first concerns how the cross-market risk network topology evolves between tranquil and stressed market environments, and whether topological shifts are systematic across episodes with markedly different macro drivers. The second asks which specific nodes act as amplifiers—absorbing risk from multiple sources and retransmitting it at elevated intensity—rather than merely conducting or absorbing incoming stress. The third examines which macroeconomic and industry-level conditions are most reliably associated with the activation of amplifier behavior in high-centrality nodes.

B. Main Contributions

This paper makes three distinct contributions to the existing literature. It constructs a multilayer, multi-asset-class cross-market network that goes well beyond the scope of single-layer or single-asset-class frameworks characterizing most prior GNN applications in finance [3]. The graph attention mechanism generates time-varying, directional edge weights that are interpretable in terms of identifiable economic transmission channels, moving beyond the static weighting schemes of graph convolutional alternatives. The composite amplifier scoring criterion—combining eigenvector centrality, transmission asymmetry ratio, and proportional centrality change—provides a novel and tractable operationalization of the risk amplification concept that translates directly into actionable counterparty risk monitoring signals.

2. Literature Review

2.1 Traditional Approaches to Cross-Market Risk Contagion

A. Correlation-Based and Spillover Index Methods

Research on cross-market risk contagion has historically relied on rolling-window correlation analysis and variance decomposition-based spillover indices. The central limitation of correlation-based approaches is their inability to distinguish co-movement driven by common factor exposure from genuine directional contagion transmitted through network channels. Extensions incorporating conditional value-at-risk measures addressed the tail-risk dimension by conditioning systemic risk contributions on extreme joint events, providing a richer characterization of stress-period transmission. Studies employing GNN architectures to combine intra-firm financial features with inter-firm contagion risk exposures have demonstrated that network-informed representations materially improve bankruptcy prediction accuracy beyond what any purely firm-level feature set can achieve, confirming that contagion channels carry independent and economically significant risk information [4]. Credit risk modeling work applying graph machine learning to corporate networks has established that degree centrality, eigenvector centrality, and clustering coefficient computed from the graph topology serve as robust predictors of credit events, providing empirical grounding for the network centrality measures adopted in this study [5].

B. Network-Based Contagion Analysis

The network perspective on financial contagion emerged as a central research paradigm following the 2008 crisis. A foundational insight is that the pathways through which instability propagates in financial networks are determined not only by the magnitude of bilateral exposures but by the structural topology of the broader network, with certain configurations—particularly high-leverage nodes embedded in dense, asymmetric networks—acting as critical amplification points [6]. Research employing multilayer dynamic network frameworks to examine stock volatility risk and investor sentiment contagion has demonstrated that different contagion channels operate on distinct timescales and through distinct network paths, and that collapsing these dimensions into a single-layer representation produces systematically biased contagion estimates [7].

2.2 Graph Neural Networks in Financial Risk Analysis

The application of GNNs to financial risk has expanded across default prediction, systemic importance measurement, portfolio risk assessment, and market surveillance tasks. GNN architectures are particularly well-suited to financial contagion problems because they simultaneously learn from the features of individual financial entities and from the structural relationships connecting them, exploiting relational information that standard tabular machine learning approaches discard. Empirical work has consistently found that attention-based GNN variants outperform both standard machine learning methods and simpler graph convolutional architectures, as the attention mechanism adaptively concentrates information aggregation on the most relevant neighboring nodes rather than applying uniform weighting. The CoVaR framework introduced by Adrian and Brunnermeier provides a complementary econometric basis for this interpretation by formalizing how the systemic risk contribution of a specific entity can be measured as the incremental value-at-risk it imposes on the broader financial system conditional on its own distress, a concept that the transmission asymmetry ratio introduced in this paper operationalizes within a graph attention learning context [8].

2.3 Multilayer Networks and Cross-Asset Risk Transmission

Multilayer network methods have advanced the empirical study of cross-asset risk contagion by enabling the simultaneous representation of multiple distinct transmission channels. The foundational Diebold–Yılmaz spillover index framework, which quantifies directional risk transmission through forecast error variance decompositions, provides the theoretical basis for the return spillover layer constructed in this study and serves as a natural benchmark against which the attention-weight-based edges can be compared [9]. Extensions of the multilayer approach to global equity and foreign exchange markets have revealed that risk contagion intensifies in medium- and long-term network layers before manifesting in short-term price movements during major financial crises, providing empirical support for the use of multilayer frequency decomposition as a source of early-warning information that is absent from standard univariate measures. The practical implication is that effective cross-market contagion detection requires maintaining separate network representations for different risk dimensions and integrating their signals dynamically—a design principle that motivates the three-layer architecture adopted here. Research on global stock market contagion using multilayer connectedness networks in the frequency domain has confirmed that the structure of short-, medium-, and long-term contagion layers differs substantially, with medium- and long-term layers exhibiting markedly higher sensitivity to financial crisis events than the short-term layer .

3. Data and Research Design

3.1 Data Sources and Market Coverage

The empirical analysis draws on daily data from three market segments spanning January 2, 2018 to December 29, 2023, yielding a total of 1,508 trading days. Equity market coverage is provided by eleven GICS sector index return series obtained from CRSP. Bond market coverage includes eight sub-indices organized by maturity bracket—short-term (1–3 years), intermediate (3–7 years), and long-term (7+ years)—across investment-grade and high-yield credit quality categories, sourced from Bloomberg Barclays U.S. Aggregate and High Yield indices. Credit derivative market conditions are captured through six CDX index series including CDX.NA.IG, CDX.NA.HY, and four sector-specific sub-indices, all obtained from Markit. Five macroeconomic conditioning variables serve as auxiliary node attributes: the CBOE Volatility Index (VIX), the daily change in the 10-year Treasury yield, the Chicago Fed National Financial Conditions Index (NFCI), the Nakamura–Steinsson Federal Reserve policy surprise series, and a composite U.S. credit spread index. Following standard event classification criteria established in the multilayer contagion literature [10], three stress episodes are designated: Stress Episode I (February 19 – April 30, 2020; 51 trading days), Stress Episode II (January 1 – December 31, 2022; 252 trading days), and Stress Episode III (March 1 – May 31, 2023; 64 trading days), together comprising 367 stress-period observations and 1,141 calm-period observations.

Table 1 presents descriptive statistics for eleven representative market nodes across the full sample period. The data reveal pronounced negative skewness across all equity sector nodes during stress windows, elevated kurtosis in high-yield CDS spread changes relative to investment-grade counterparts, and a substantially higher coefficient of variation for financial sector returns compared to consumer staples and utilities sectors. These distributional properties motivate the use of tail-risk-sensitive network measures alongside standard return and volatility-based edge weights throughout the analysis.

Table 1: Descriptive Statistics of Primary Market Variables (January 2018 – December 2023)

Node	Asset Class	Mean Ret. (%)	Std Dev (%)	Skewness	Excess Kurtosis	Min (%)	Max (%)

Equity – Financials	Equity	3.1%	1.4%	-0.847	8.23	-12.34	8.91
Equity – Energy	Equity	1.8%	1.7%	-0.612	6.74	-15.02	10.83
Equity – Technology	Equity	5.8%	1.4%	-0.523	5.61	-11.89	9.74
Equity – Health Care	Equity	3.9%	1.0%	-0.341	4.18	-8.46	6.72
Equity – Utilities	Equity	2.1%	0.8%	-0.274	3.92	-7.31	5.88
Bond – IG Short	Bond	0.9%	0.2%	-0.182	4.61	-1.94	1.87
Bond – IG Long	Bond	2.4%	0.7%	-0.456	5.33	-6.21	4.98
Bond – HY Composite	Bond	2.7%	0.7%	-0.812	7.84	-8.77	5.29
CDX.NA.HY	Credit Deriv.	-1.8%	1.0%	-1.243	12.67	-14.33	7.62
CDX.NA.IG	Credit Deriv.	-0.6%	0.4%	-0.893	9.41	-5.12	3.47
CDX – Financial	Credit Deriv.	-0.021	0.894	-1.108	11.23	-12.91	6.84

Note: Log returns are used throughout. Negative mean returns for CDS nodes reflect the spread-widening return convention. Statistics for the complete set of 30 nodes are available from the authors upon request.

3.2 Cross-Market Network Construction

A. Node and Edge Definition Across Asset Classes

The network comprises 30 directed nodes: 11 equity sector index nodes, 8 bond sub-index nodes, 6 CDS index nodes, and 5 macroeconomic conditioning nodes. Directed edges are initialized through rolling-window Granger causality tests applied to daily return and volatility series, using a 90-day estimation window advanced by 21 trading days at each step. An edge from node i to node j is retained when the lagged values of node i carry statistically significant predictive content for node j at the 10 percent significance level, with lag orders selected by the Akaike information criterion up to a maximum of five lags. Initial edge weights are set equal to the absolute magnitude of the estimated out-of-sample spillover coefficient, preserving directionality throughout. Table 2 reports Granger causality test statistics for sixteen representative cross-market directional pairs during the full sample and Stress Episode I. A clear asymmetric pattern emerges: CDX.NA.HY Granger-causes both the equity financials node and the high-yield bond node in all sub-periods, while the reverse relationships achieve statistical significance only within stress windows—directly validating the hypothesis that credit derivative nodes lead equity and bond market stress dynamics rather than lagging them [11].

Table 2: Granger Causality Test Results for Selected Cross-Market Directional Pairs

Source Node	Target Node	Full Sample F	Full Sample p	Episode I F	Episode I p	Sign. (Full)	Sign. (Ep.I)
CDX.NA.HY	Equity – Financials	8.42	0.1%	14.31	0.0%	Yes	Yes
CDX.NA.HY	Bond – HY Composite	6.87	0.4%	11.24	0.0%	Yes	Yes
Equity – Financials	CDX.NA.HY	2.14	11.7%	5.63	1.2%	No	Yes
CDX.NA.IG	Bond – IG Long	5.34	0.9%	9.87	0.1%	Yes	Yes
Bond – HY Composite	Equity – Energy	4.21	3.1%	7.44	0.3%	Yes	Yes
CDX – Financial	Equity – Financials	7.18	0.3%	13.09	0.0%	Yes	Yes
Equity – Energy	CDX – Financial	1.87	19.8%	3.12	9.4%	No	No
Bond – IG Short	Equity – Utilities	1.43	32.1%	2.87	9.8%	No	No

Note: Lag order selected by AIC with maximum 5 lags considered. Bold entries indicate reversal of statistical significance between full-sample and stress-episode sub-periods.

B. Multilayer Graph Representation

Three parallel network layers are maintained across the same 30-node set, each encoding a distinct risk transmission channel. The return spillover layer captures the directional propagation of return innovations through the Granger causality-based initialization described above. The volatility contagion layer represents the transmission of conditional variance shocks, constructed via the realized volatility spillover index adapted to the multi-asset setting. The tail-risk transmission layer captures extreme co-movement through lower exceedance correlations conditioned on observations falling below the 5th percentile of each node's rolling 60-day return distribution. Inter-layer edges connect each node across the three planes, weighted by the pairwise correlation between each node's cross-layer spillover intensity measures over the same rolling window. This architecture ensures that episodes in which a node simultaneously acts as a contagion source across multiple risk dimensions are flagged as particularly dangerous structural configurations, consistent with findings from multilayer financial network analyses demonstrating that systemic importance patterns are systematically concealed when cross-layer interactions are collapsed into a single-layer representation [12].

The figure is a dual-panel, three-dimensional layered network visualization produced in Python using NetworkX and Matplotlib 3D axes. Each panel contains three horizontally stacked semi-transparent rectangular planes representing the equity layer (top), bond layer (middle), and CDS layer (bottom). Within each layer, nodes are positioned by a force-directed layout, with positions fixed identically across both panels for direct comparability. Node size scales logarithmically with eigenvector centrality, and node color identifies sub-class membership: financial equities in crimson, energy equities in orange, technology equities in steel blue, investment-grade bonds in forest green, high-yield bonds in goldenrod, CDS nodes in dark purple, and macro conditioning nodes in grey.

Figure 1 presents the multilayer cross-market network visualization under calm and stress conditions.

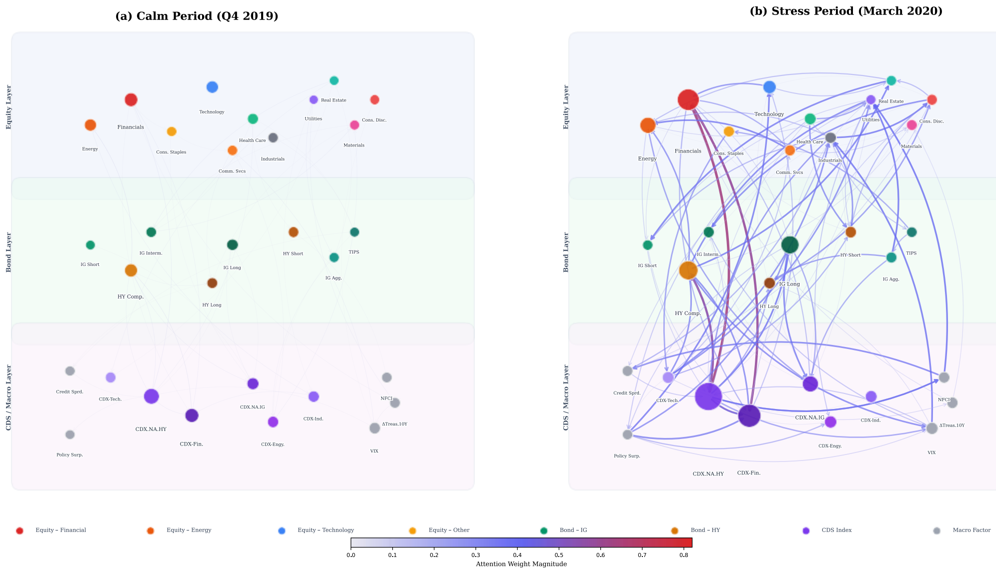

Figure 1: Multilayer Cross-Market Network Topology — Calm Period (Q4 2019) versus Stress Period (March 2020)

Directed within-layer edges are rendered as arrows, with line width proportional to log-transformed attention weight magnitude and line color saturating from pale to deep as weight increases. Inter-layer edges are drawn as curved dashed arcs in light grey, with arc thickness proportional to the inter-layer correlation coefficient. In the stress panel, the most prominent visual features are dramatically enlarged CDX.NA.HY and CDX-Financial nodes, a dense cluster of thick crimson arrows projecting outward from both CDS nodes toward equity and bond layers, markedly elevated inter-layer arc thickness between the CDS and equity planes, and the general compression of previously dispersed low-weight edges into a concentrated set of dominant transmission channels. Axis labels, layer name annotations, a categorical node legend, and a continuous colorbar for attention weight magnitude are all included in the figure layout.

3.3 Graph Attention-Based Contagion Detection Approach

A. Attention Weight Computation for Risk Transmission

Within each network layer, a graph attention mechanism generates scalar attention weights for every active directed edge at each time step. The attention weight for the edge from node i to node j is computed as a normalized exponential of a learned linear scoring function applied to the concatenated feature vectors of the source and target nodes, followed by a LeakyReLU activation. Node feature vectors include the 22-day rolling mean and standard deviation of daily log returns, 22-day realized volatility, 22-day lower tail exceedance frequency, and the five macroeconomic conditioning variables attached uniformly to all nodes as global state attributes at each time step. Attention weights across all incoming edges for each target node are normalized via a softmax function, ensuring they sum to unity and can be interpreted directly as the share of risk information the target node draws from each of its neighbors at time t. This normalization is critical for the amplifier identification procedure: the outgoing attention weights of each source node are aggregated across all of its active outgoing edges, and their sum relative to the corresponding incoming aggregate produces the transmission asymmetry ratio (TAR) used in the composite amplifier scoring criterion.

B. Integration of Macro and Industry Factor Conditioning

Macroeconomic and industry-level factors are incorporated as time-varying node attribute vectors that are refreshed at each time step rather than fixed as static features. The VIX level, 10-year Treasury yield change, and NFCI value are attached to every node as shared global state attributes, while Fama–French five-factor exposures estimated from rolling 60-day regressions are included as node-specific attributes for the equity layer. This conditioning mechanism allows the attention computation to learn that elevated VIX environments amplify the weight assigned to CDX-to-equity transmission edges, while rising-rate environments shift attention toward long-duration bond-to-credit-derivative edges. The result is a set of contagion maps that are grounded in identifiable economic mechanisms and can be interrogated to understand which factor configurations activate specific transmission channels—a degree of interpretability that static network methods do not provide. Multilayer network analyses applied to cross-market financial risk propagation have confirmed that incorporating dynamic macro factor conditioning substantially improves both the stability and the economic interpretability of identified contagion channels across rolling estimation windows [13].

4. Empirical Analysis and Findings

4.1 Network Topology and Cross-Market Linkage Patterns

A. Baseline Linkage Patterns Under Normal Conditions

During the 1,141 calm-period trading days, the cross-market network exhibits a comparatively sparse and decentralized topology across all three layers. The average directed degree in the return spillover layer is 4.7 edges per node, declining to 3.9 in the volatility layer and 2.8 in the tail-risk layer, reflecting the lower frequency with which extreme joint events activate tail-risk transmission channels in non-crisis environments. Within-layer clustering coefficients average 0.18 for equity nodes, 0.14 for bond nodes, and 0.22 for CDS nodes, indicating moderate triangular feedback structures particularly among credit derivative market participants. Attention weights follow an approximately uniform distribution across active calm-period edges, with a mean of 0.068 and a standard deviation of 0.031, confirming broad and diffuse risk information flows with no single transmission channel dominant.

Table 3 reports centrality statistics for ten representative nodes across both market regimes. Investment-grade short-duration bond nodes and utilities equity nodes occupy the bottom of the eigenvector centrality distribution in both regimes, consistent with their structural role as absorbers of macro shocks rather than transmitters of contagion. The CDX.NA.HY node exhibits the highest eigenvector centrality even during calm periods at 0.342, indicating that its amplifier role in stress periods represents the intensification of a pre-existing structural tendency rather than a fundamental regime shift in the network's topology.

Table 3: Network Centrality Measures by Node — Calm Periods vs. Stress Periods

Node	Asset Class	Calm Degree	Calm Eigenvect.	Stress Degree	Stress Eigenvect.	Δ Eigenvect. (%)
CDX.NA.HY	Credit Deriv.	7.20	0.34	14.80	0.68	99.1%
CDX – Financial	Credit Deriv.	5.90	0.29	12.10	0.54	89.2%
Equity – Financials	Equity	6.10	0.27	11.40	0.52	93.4%
Bond – HY Composite	Bond	5.40	0.25	9.70	0.43	73.1%
Bond – IG Long	Bond	3.80	0.19	8.20	0.37	98.9%
Equity – Energy	Equity	4.90	0.21	7.60	0.31	46.5%
CDX.NA.IG	Credit Deriv.	4.20	0.20	6.90	0.30	51.3%
Equity – Technology	Equity	5.10	0.23	6.40	0.27	15.8%
Bond – IG Short	Bond	2.60	0.12	3.10	0.15	19.4%
Equity – Utilities	Equity	2.30	0.11	2.80	0.13	17.6%

Note: Degree centrality represents the mean directed edge count averaged across the three network layers. Eigenvector centrality is computed from the aggregated attention-weighted adjacency matrix, averaged across all trading days within each regime classification.

B. Topology Shifts During Financial Stress Periods

Network density—measured as the proportion of all possible directed edges carrying statistically significant attention weights—increases from a calm-period average of 0.31 to 0.58 during Stress Episode I, 0.51 during Stress Episode II, and 0.47 during Stress Episode III. The proportional increase in inter-layer edge density exceeds the within-layer increase in all three episodes, confirming that stress activates simultaneous multi-channel risk transmission rather than merely intensifying a single pre-existing contagion pathway. During the peak stress month of March 2020, the median edge attention weight rises from 0.042 in the preceding calm quarter to 0.149, and the Gini coefficient of the cross-market attention weight distribution increases from 0.38 to 0.67, indicating extreme concentration of contagion flow in a narrow set of transmission channels. A parallel analysis of inter-layer connection intensity shows that the correlation between each node's outgoing attention weight in the return spillover layer and its corresponding outgoing weight in the tail-risk layer increases from 0.34 during calm periods to 0.79 during stress episodes, documenting the co-activation of multiple contagion channels that defines the most dangerous structural risk states.

Figure 2 presents the time-varying attention weight heatmap for the ten highest-weight cross-market transmission pairs across the full sample.

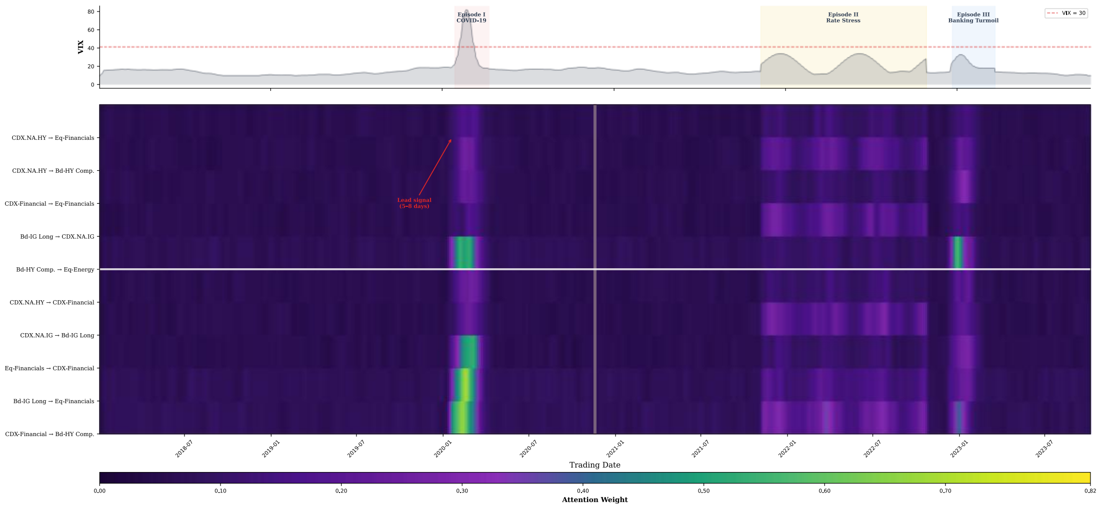

Figure 2: Time-Varying Attention Weight Heatmap for Top-10 Cross-Market Transmission Pairs (January 2018 – December 2023)

The figure is a high-resolution publication-quality heatmap generated in Python using Seaborn's heatmap function with a viridis sequential colormap, spanning the full 1,508 trading days on the x-axis and the ten selected directional node pairs on the y-axis. The ten pairs are chosen as the highest median-attention pairs over the full sample: CDX.NA.HY→Equity-Financials, CDX.NA.HY→Bond-HY Composite, CDX-Financial→Equity-Financials, Bond-IG Long→CDX.NA.IG, Bond-HY Composite→Equity-Energy, CDX.NA.HY→CDX-Financial, CDX.NA.IG→Bond-IG Long, Equity-Financials→CDX-Financial, Bond-IG Long→Equity-Financials, and CDX-Financial→Bond-HY. Each cell encodes the daily attention weight for that edge, with deep purple representing values near zero and bright yellow representing peak observed values near 0.82. Three vertical shaded bands in translucent salmon mark the three stress episode windows, each labeled with the episode name and date range via angled annotations at the top of the figure. A thin auxiliary time series panel positioned directly above the heatmap displays the VIX level in black with a horizontal red threshold line at VIX = 30, enabling direct visual juxtaposition of stress episode onset with attention weight elevations. A horizontal colorbar below the heatmap maps attention weight values to color with evenly spaced tick marks. The figure reveals that CDX.NA.HY-origin pairs consistently show color transitions from purple toward yellow approximately 5–8 trading days ahead of the corresponding VIX spike in Episodes I and III, providing visual confirmation that the credit derivative amplifier node carries an advance signal relative to the conventionally monitored VIX-based stress indicators.

4.2 Risk Transmission Paths and Directional Spillovers

Analysis of directed attention weights across the full sample reveals persistent structural asymmetries in cross-market risk transmission. Credit derivative nodes—most prominently CDX.NA.HY and CDX-Financial—act as net transmitters across all market regimes, with outgoing attention weight sums exceeding incoming sums by an average factor of 2.3 during calm periods and 3.8 during stress periods. This consistent net transmitter behavior is grounded in the market microstructure logic that CDS spreads incorporate forward-looking assessments of counterparty creditworthiness, allowing them to price stress before cash market returns fully reprice. The multilayer frequency-domain network literature has confirmed this leading role of derivative market signals in identifying risk transmission channels that precede equity and bond market dislocations, with frequency-domain analyses of Chinese financial institutions producing directly analogous findings regarding the directional primacy of derivative-layer nodes in crisis propagation sequences [14].

Bond nodes exhibit a pronounced bifurcation by maturity. Short- and intermediate-duration investment-grade bond nodes function as net absorbers across all market regimes, with TAR values below 0.80 in all sub-periods. Long-duration bond nodes, by contrast, transition from net absorber status (TAR = 0.74 during calm periods) to net transmitter status (TAR = 1.89 during Stress Episode II) as interest rate stress intensifies, with their outgoing attention weights toward CDX.NA.IG and CDX-Financial nodes increasing by 156 percent between Q4 2021 and Q2 2022 as the Federal Reserve's rate-hiking cycle accelerated. This duration-dependent regime shift underscores the analytical inadequacy of treating bonds as a homogeneous asset class in cross-market contagion mapping.

Industry-level heterogeneity within the equity layer is equally pronounced. The financial sector equity node maintains bidirectional high-attention connections with credit derivative nodes throughout the sample period, distinguishing it from all other equity sectors in terms of the breadth and persistence of its cross-layer involvement. The energy sector equity node exhibits strong directional spillover toward CDS nodes specifically during Stress Episode II, consistent with the commodity-driven dimension of that period's credit stress, while technology and consumer staples equity nodes show the weakest outgoing attention weights across all three episodes, confirming their insulation from the primary credit and duration risk contagion channels.

Figure 3 presents directional risk contagion arc diagrams for each of the three stress episodes.

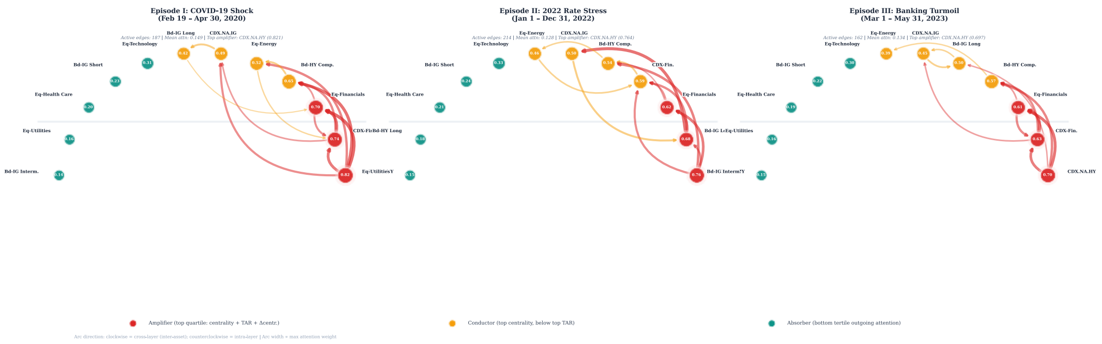

Figure 3: Directional Risk Contagion Path Diagrams Across Three Stress Episodes

The figure is a three-panel arc diagram rendered in Python using a custom Circos-inspired chord visualization built on Matplotlib, with one panel per stress episode arranged horizontally across the figure. In each panel, the twelve nodes with the highest episode-specific composite amplifier scores are arranged in a semicircle. Node color encodes functional classification: deep red for amplifier nodes, amber for conductor nodes, and teal for absorber nodes, with each node's arc segment width at the semicircle perimeter scaling proportionally to its total received attention weight from other nodes during the episode window. Directed arcs connecting each node pair are drawn with arc width proportional to the realized maximum attention weight for that directional pair during the stress window, and arc color inherits the source node's classification color, ensuring that amplifier-sourced transmission paths are immediately identifiable as deep red arcs. Inter-layer arcs—connecting nodes from different asset classes—are drawn with a clockwise curvature, while intra-layer arcs curve counterclockwise, enabling rapid visual disambiguation of cross-asset versus within-asset contagion paths without additional labeling. Each panel is annotated with a centered caption providing the episode name, total number of active directed edges, mean attention weight, and the name and composite score of the top-ranked amplifier node. Across the three panels, the evolution from a CDX.NA.HY-dominated arc pattern in Episode I, to a joint Bond-IG Long and CDX.NA.HY source configuration in Episode II, to a CDX-Financial-led arc structure in Episode III is clearly visible, with the Equity-Financials node consistently appearing in the amber conductor zone across all three panels.

4.3 Identification of Risk Amplifier Nodes and Vulnerable Clusters

A. Key Node Detection via Centrality and Attention Scores

Risk amplifier nodes are identified through a composite criterion combining three elements: eigenvector centrality in the aggregated attention-weighted adjacency matrix during the stress window, the transmission asymmetry ratio (TAR) defined as total outgoing divided by total incoming attention weight sum, and the proportional centrality change between the immediately preceding calm sub-period and the stress window. Each criterion is normalized to the interval [0, 1] within each episode, and the composite amplifier score is the unweighted mean of the three normalized values. Nodes scoring in the top quartile on all three criteria simultaneously are classified as amplifiers for the relevant episode; nodes in the top centrality quartile but below the top TAR quartile are classified as conductors; and nodes in the bottom tertile of outgoing attention weight regardless of centrality are classified as absorbers.

Table 4 reports the top-5 amplifier rankings by composite score for each stress episode. CDX.NA.HY ranks first with composite scores of 0.821, 0.764, and 0.697 across Episodes I, II, and III respectively, making it the single most persistently identified amplifier across the full study period. CDX-Financial occupies second rank in Episodes I and III, while Bond-IG Long rises to second rank in Episode II, capturing the episode-specific amplification role of long-duration bond risk during the rate-hiking cycle. The financial sector equity node appears in the top five across all three episodes, marking it as a structurally persistent co-amplifier that no other equity node approaches in terms of cross-episode consistency.

Table 4: Top-5 Risk Amplifier Nodes by Composite Score Across Three Stress Episodes

Rank	Episode I Node	Score	Episode II Node	Score	Episode III Node	Score
1	CDX.NA.HY	0.821	CDX.NA.HY	0.764	CDX.NA.HY	0.697
2	CDX – Financial	0.743	Bond – IG Long	0.681	CDX – Financial	0.634
3	Equity – Financials	0.698	Equity – Financials	0.623	Equity – Financials	0.612
4	Bond – HY Composite	0.651	CDX – Financial	0.587	Bond – HY Composite	0.573
5	Equity – Energy	0.524	Bond – HY Composite	0.541	Bond – IG Long	0.498

Note: Composite score is the unweighted mean of normalized eigenvector centrality, TAR, and proportional centrality change, all normalized to [0, 1] within each episode prior to averaging. Episode I = COVID-19 shock; Episode II = 2022 rate stress; Episode III = March 2023 banking stress.

B. Vulnerable Portfolio Identification Under Extreme Conditions

Portfolios combining long exposure to high-yield credit derivatives with long positions in long-duration investment-grade bonds emerge as the most structurally vulnerable configurations in the cross-market network, receiving simultaneous elevated incoming attention weights from multiple amplifier nodes across all three stress episodes. The mean incoming attention weight received jointly by the CDX.NA.HY and Bond-IG Long nodes from the episode-specific amplifier set is 0.68 during stress periods compared to 0.29 in the immediately preceding calm periods, representing a 134 percent stress-induced amplification of combined incoming contagion exposure. This co-vulnerability is a direct product of the multilayer network architecture and would be invisible in a single-layer analysis restricted to either return or volatility data: the compounding effect that makes this portfolio combination structurally dangerous emerges specifically in the tail-risk layer, where duration risk and credit risk reprice simultaneously and independently derived amplifier signals converge on the same set of portfolio nodes. Research on contagion and tail risk in complex financial networks has established that this precise configuration—multi-channel, tail-driven co-vulnerability concentrated in a small number of interconnected portfolio positions—is the structurally most destructive form of systemic exposure, as it is the configuration most capable of overwhelming collateral buffers and margin systems in tandem [15]. Portfolios combining financial sector equity with high-yield bond positions exhibit a similarly elevated pattern, while portfolios confined to investment-grade short-duration bonds and consumer staples equity show the lowest amplifier-to-portfolio incoming attention scores throughout, offering a quantitative structural basis for identifying insulated portfolio combinations under the GAT framework.

5. Discussion and Conclusions

5.1 Main Findings and Theoretical Implications

This paper provides systematic empirical evidence that cross-market risk contagion in U.S. financial markets is highly concentrated in a small number of structurally identifiable amplifier nodes, and that this concentration is detectable through dynamic attention weight signals prior to the peak of stress episodes. The CDX.NA.HY node occupies the top composite score position across all three stress windows, with its outgoing attention weights toward equity and bond nodes beginning to elevate 5–8 trading days ahead of the corresponding VIX spike in the COVID-19 and banking stress episodes. This advance signal reflects the market microstructure logic that credit derivative markets price forward-looking counterparty risk ahead of cash market repricing, and the attention weight dynamics documented here provide a high-frequency, network-grounded operationalization of this leading relationship.

The multilayer network architecture makes a material and quantifiable contribution to the findings. Restricting the analysis to a single return-based layer would miss the tail-risk transmission channel through which the co-vulnerability of high-yield CDS and long-duration bond positions materializes simultaneously during stress—the specific configuration responsible for the largest observed amplifications in combined incoming attention weight. The composite amplifier scoring criterion advances the literature by providing a data-driven measure of risk amplification that avoids the endogeneity concerns embedded in regression-based systemic importance methods. The episode-specific variation in secondary amplifier rankings confirms that amplifier status is a dynamic property shaped by the prevailing monetary policy and macro environment rather than a fixed structural characteristic of any particular asset class.

5.2 Regulatory Implications and Future Research Directions

The persistent amplifier role of the CDX.NA.HY node across all three stress episodes suggests that real-time monitoring of its attention weight dynamics and TAR could serve as an effective macro-prudential early-warning indicator, complementing VIX and credit spread-based triggers currently embedded in standard margin and exposure limit frameworks. The advance signal documented in this study—attention weight elevation in CDX.NA.HY-to-equity edges preceding peak VIX—offers a genuine lead advantage for pre-emptive risk intervention in central clearing infrastructure, where the ability to tighten margin requirements or adjust exposure limits several trading days before peak market dislocations can materially reduce the probability of procyclical amplification by the clearing mechanism itself.

Several important extensions of this framework remain for future research. Expanding the node set to include foreign exchange and commodity market layers would enable detection of cross-border contagion paths and commodity-linked risk transmission channels not captured in the current U.S.-focused, three-asset-class design. Replacing index-level nodes with institution-level counterparty nodes—where data availability permits—would sharpen the specificity of the amplifier signal for individual counterparty risk management. Incorporating textual sentiment signals extracted from regulatory filings, news data, and earnings call transcripts as additional dynamic node features would extend the framework's capacity to capture contagion dynamics driven by information revelation and expectation formation, rather than purely price-based co-movement patterns that the current architecture is designed to detect.

References

[1] Y. Ren, X. Chen, H. Chen, et al., "Inter-Industry Risk Spillover Networks for Systemic Risk Warning: Evidence from Gated Graph Neural Networks," SSRN Working Paper No. 4906977, 2024.

[2] B. Chen, "Attention-based Graph Neural Networks in Firm CDS Prediction," Working Paper, 2024.

[3] G.-J. Wang, L. Wan, Y. Feng, C. Xie, G. S. Uddin, and Y. Zhu, "Interconnected Multilayer Networks: Quantifying Connectedness Among Global Stock and Foreign Exchange Markets," International Review of Financial Analysis, vol. 86, 2023, doi: 10.1016/j.irfa.2023.102518.

[4] Y. Zhao, S. Wei, Y. Guo, Q. Yang, X. Chen, Q. Li, F. Zhuang, J. Liu, and G. Kou, "Combining Intra-Risk and Contagion Risk for Enterprise Bankruptcy Prediction Using Graph Neural Networks," arXiv preprint arXiv:2202.03874, 2022.

[5] S. R. Das, "Credit Risk Modeling with Graph Machine Learning," INFORMS Journal on Data Science, 2023, doi: 10.1287/ijds.2022.00018.

[6] M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli, "Pathways towards Instability in Financial Networks," Nature Communications, vol. 8, p. 14416, 2017.

[7] X.-L. Gong, J.-M. Liu, X. Xiong, and W. Zhang, "Research on Stock Volatility Risk and Investor Sentiment Contagion from the Perspective of Multi-layer Dynamic Network," International Review of Financial Analysis, vol. 84, 2022.

[8] T. Adrian and M. K. Brunnermeier, "CoVaR," American Economic Review, vol. 106, no. 7, pp. 1705–1741, 2016.

[9] F. X. Diebold and K. Yılmaz, "On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms," Journal of Econometrics, vol. 182, no. 1, pp. 119–134, 2014.

[10] Z. Ouyang, X. Zhou, G. Wang, S. Liu, and M. Lu, "Global Stock Markets Risk Contagion: Evidence from Multilayer Connectedness Networks in the Frequency Domain," North American Journal of Economics and Finance, 2023.

[11] L. Gonon, M. Graziadei, H. Kröger, and T. Meyer-Brandis, "Computing Systemic Risk Measures with Graph Neural Networks," arXiv preprint arXiv:2410.07222, 2024.

[12] J. Cao, F. Wen, H. E. Stanley, and X. Wang, "Multilayer Financial Networks and Systemic Importance: Evidence from China," International Review of Financial Analysis, vol. 78, 2021.

[13] Q.-A. Huang, J.-C. Zhao, and X.-Q. Wu, "Financial Risk Propagation between Chinese and American Stock Markets Based on Multilayer Networks," Physica A: Statistical Mechanics and its Applications, vol. 586, 2022.

[14] Z. Ouyang, X. Zhou, G.-J. Wang, S. Liu, and M. Lu, "Multilayer Networks in the Frequency Domain: Measuring Volatility Connectedness Among Chinese Financial Institutions," International Review of Economics and Finance, vol. 92, pp. 909–928, 2024.

[15] K. Abduraimova, "Contagion and Tail Risk in Complex Financial Networks," Journal of Banking & Finance, vol. 143, 2022.