AIMS Mathematics, 10(11): 26313–26333.
DOI: 10.3934/math.20251157
Received: 17 August 2025
Revised: 28 October 2025
Accepted: 05 November 2025
Published: 14 November 2025

https://www.aimspress.com/journal/Math

Research article

Stress-aware multiscale spillover networks: Cross-market transmission via coherence–entropy centrality

Çağlar Sözen*

Department of Finance and Banking, Görele School of Applied Sciences, Giresun University, Giresun, Turkey

* **Correspondence:** Email: caglar.sozen@giresun.edu.tr;
ORCID: https://orcid.org/0000-0002-3732-5058.

Abstract: Systemic-risk monitoring frameworks are largely built on absolute Pearson correlation networks: Assets are linked if their returns co-move on average, and "systemic hubs" are defined by high degree. Such approaches implicitly assume (i) a single contagion timescale and (ii) stability of dependence, even though crises typically unfold in layers: A fast equity/volatility unwind, followed by slower stress in funding, FX, rates, and commodities. We proposed a stress-aware, multiscale alternative, and constructed the multiscale coherence–entropy centrality (MCEC) network in which (a) an edge between two assets exists only if their wavelet coherence is statistically significant and persistent across adjacent frequency bands, and (b) node importance is an entropy-weighted multi-horizon strength that is high only if an asset is strongly connected and active across time scales. We then generated a synthetic stressed panel by shocking all assets with a common heavy-tailed t-copula draw scaled by GARCH(1,1) volatilities, and compared MCEC to a traditional absolute-correlation backbone using 2021–2024 data. We reported three findings that are directly relevant for macroprudential supervision. First, under stress, the MCEC network reallocated centrality toward canonical stress transmitters (U.S. equity benchmarks, implied volatility (VIX), dollar/FX, long-term yields, crude oil, and gold), while ordinary correlation networks continued to present a single equity-dominated block. Second, MCEC delivered higher ex-ante classification performance (AUC) in identifying those transmitters even before the stressed regime was applied, indicating early-warning value. Third, MCEC made the stress-driven rewiring of cross-market spillover channels explicit across horizons rather than treating dependence as static.

Keywords: systemic risk; macroprudential surveillance; multiscale coherence–entropy centrality; wavelet coherence; stress testing; cross-market spillovers; financial networks
Mathematics Subject Classification: 62M10, 91G70

1. Introduction

Supervisors and risk committees routinely ask: Who will transmit stress, through which channels, and on what horizon? The standard answer is to build a network from return correlations: Assets are linked if their Pearson correlation is large, and "systemic hubs" are nodes with high degree or eigenvector centrality [1–4]. This is now common in systemic-risk surveillance.

That approach bakes in two assumptions that fail in crises. First, it treats contagion as if it lived on a single timescale. In reality, stress is layered. There is a fast leg, margin calls, deleveraging, volatility-control selling, flight from equities and crypto, that generates near-synchronous moves over hours or days. There is also a slower leg in rates, FX, and commodities: Dollar funding tightens, long-term sovereign yields reprice balance-sheet risk, crude oil and gold become macro/liquidity barometers, and those adjustments unfold over weeks [5–7]. Second, correlation networks implicitly assume that the dependence structure is stable. Yet, in stress, the channels can rewire: Implied volatility can suddenly become the bridge between global equities and FX; and oil can become the bridge between funding stress and inflation risk. A static absolute-correlation matrix cannot tell a policymaker whether an asset is dangerous because it amplifies an intraday fire-sale spiral or because it anchors a slow repricing of global funding costs.

We propose a stress-aware, scale-aware alternative, multi-scale coherence-entropy centrality (MCEC) network. It has two ingredients.

For each pair of assets, we compute wavelet coherence across a ladder of frequency bands (2–4 days, 4–8 days, ..., 32–64 days) and test each band against a Monte Carlo null [7, 8]. We admit an edge only if coherence is simultaneously significant in adjacent bands; if it is significant in only one band, we keep it but downweight it. This adjacency rule encodes supervisory intuition: A real transmission channel should persist across nearby horizons, not just spike once in a single noisy band. The resulting edge weight is robust multi-horizon co-movement, not just "high $|\rho|$ today."

For each asset, we measure how strongly it co-moves with the system in each band, then compute the Shannon entropy of that band profile. A node's MCEC score is high only if it is both strongly connected and active across horizons. Intuitively, a high-MCEC node can ignite fast panic and sustain slower macro/liquidity pressure. This is closer to what supervisors mean by a "systemic transmitter" than raw degree, which treats all links as interchangeable [1, 2].

Stress testing should also be stress conditional, not just historical. Rather than rerunning networks on one crisis subsample [3, 4], we generate a synthetic "post-stress" return panel. For each asset, we estimate a GARCH(1,1) to obtain its conditional volatility path and standardized residuals [9]; we then fit a low–degrees–of–freedom t-copula to capture fat-tailed joint dependence [10], draw a common heavy-tailed shock, and scale that shock by each asset's conditional volatility. High-beta risk assets (U.S. equities, Bitcoin) receive larger amplification; macro/funding assets (long-term yields, FX, crude oil, gold) receive lower but nonzero amplification. The stressed panel therefore features volatility clustering, heterogeneous amplification, and synchronized tail events, which are closer to a system-wide liquidity shock [11].

We then build two networks, a traditional absolute-correlation backbone and the proposed MCEC backbone, on both (i) the observed 2021–2024 daily panel and (ii) the synthetic stressed panel, keeping backbone density fixed for comparability. We evaluate them using metrics that supervisors actually look at: (i) Rank stability (does the same set of assets stay central under stress?), (ii) edge

persistence and community reorganization (Jaccard overlap and variation of information), (iii) coverage (does every benchmark asset still get a usable score in stress?), and (iv) classification power (AUC) in flagging canonical stress transmitters such as broad U.S. equity indices, VIX, the U.S. dollar, long-term yields, and crude oil [12].

Three findings matter for policy. First, under stress, the MCEC network re-routes centrality toward volatility (VIX), the dollar, long-term rates, oil, and gold, while global equity benchmarks remain central. It therefore lights up the bridges between "risk" and "funding/liquidity." A plain correlation network mostly collapses into one undifferentiated risk cluster and gives little guidance on which cross-market pipes actually matter. Second, MCEC already ranks these canonical transmitters highly before stress (higher pre-stress AUC), so it has genuine early-warning value; correlation-based degree is weaker out of sample. Third, MCEC makes stress-driven rewiring visible: Its backbone edges change, and community structure reshuffles across equities, FX, rates, and commodities. The correlation backbone, by contrast, looks "stable" only because it keeps the same edges even as those edges lose explanatory power under stress.

Our contributions are fourfold: (i) We introduce a transparent multiscale network in which edges require statistically significant, adjacent-band wavelet coherence, and node centrality rewards strength and cross-horizon breadth. (ii) We generate a realistic heavy-tailed stress panel using GARCH(1,1) volatilities and a t-copula shock, rather than relying on a single historical crisis. (iii) We propose an evaluation battery aligned with supervisory use: Rank stability, edge persistence, topological reorganization, coverage, and stress-transmitter AUC. (iv) We show empirically that the resulting MCEC network isolates exactly the cross-asset spillover channels, volatility, FX/dollar, long-term yields, and energy/precious metals, that policymakers monitor when liquidity pressure spreads beyond equities.

In Section 2, we describe the 2021–2024 multi-asset data and preprocessing. In Section 3, we detail the stress-generation procedure, the correlation and MCEC networks, and the evaluation metrics. In Section 4, we present the empirical comparison and supervisory interpretation. In Section 5, we discuss policy implications, limitations, and avenues for extension.

2. Data and preprocessing

We build a daily multi-asset panel spanning 2021–2024. The universe includes (i) global equity benchmarks (U.S., Europe, Asia), (ii) Bitcoin as a high-beta risk asset, (iii) macro/funding channels (implied equity volatility, the U.S. 10Y Treasury yield, the broad U.S. dollar index, EUR/USD), and (iv) key commodities (WTI crude, gold, silver) [13]. All series are aligned on the intersection of available trading days across all tickers, so that each asset shares an identical calendar of $n = 773$ daily observations. Continuously compounded log returns are computed as $r_{t,i} = \log P_{t,i} - \log P_{t-1,i}$ for asset i at date t. For yields and FX, we use daily log-differences in levels so that all inputs enter on a comparable scale.

Table 1 reports summary statistics for daily log returns (mean, volatility, skewness, kurtosis) and two stationarity tests (ADF and KPSS).

Three facts matter for the rest of the paper.

(i) Volatility is very heterogeneous. High-beta risk assets (Bitcoin, Russell 2000) move several percent per day; large developed equity indices (S&P 500, EuroStoxx 50, DAX) are closer to 1%;

macro/funding benchmarks (DXY, EUR/USD, U.S. 10Y yield) move well below 1%; and VIX is an order of magnitude more volatile than cash indices. In Section 3.1, we reflect this with asset-specific amplification factors κ_i.

(ii) Returns are heavy-tailed. Equities, Bitcoin, and crude oil show negative skewness and high excess kurtosis (fat left tails), while VIX is positively skewed with extreme kurtosis (crisis spikes). We therefore model the joint stress shock with a low–degrees–of–freedom t-copula rather than a Gaussian dependence structure (Section 3.1).

(iii) Daily returns are stationary. All series reject a unit root according to augmented Dickey–Fuller tests (ADF, $p \leq 0.01$) [14] and are not rejected as non-stationary by KPSS [15].

We therefore treat $r_{t,i}$ as stationary inputs in the traditional absolute-correlation network and the proposed multiscale coherence–entropy centrality (MCEC) network (Section 3.2).

Assets are grouped in Table 1 into (A) risk-on benchmarks (global equities, Bitcoin) and (B) macro/funding channels (volatility, rates, FX, energy/metals); Section 4 shows that, under stress, the major spillover channels MCEC flags run between these two blocks.

Table 1. Asset universe, descriptive statistics, and stationarity tests (daily log returns, 2021–2024).

Asset	Class	Mean	SD	Skew	Ex.Kurt	ADF p	KPSS p	n
Panel A: Risk-on benchmarks (Equity indices, Bitcoin)								
S&P 500	Equity index	0.0006	0.0119	−0.77	4.80	0.01	0.10	773
NASDAQ	Equity index	0.0006	0.0161	−0.69	3.77	0.01	0.10	773
Dow Jones	Equity index	0.0004	0.0103	−0.56	3.99	0.01	0.10	773
Russell 2000	Equity index	0.0002	0.0168	−0.40	3.06	0.01	0.10	773
EuroStoxx 50	Equity index	0.0004	0.0125	−0.47	5.03	0.01	0.10	773
DAX	Equity index	0.0005	0.0119	−0.25	5.58	0.01	0.10	773
FTSE 100	Equity index	0.0003	0.0093	−0.41	3.58	0.01	0.10	773
Nikkei 225	Equity index	0.0005	0.0156	−3.23	50.31	0.01	0.10	773
Hang Seng	Equity index	−0.0004	0.0181	0.24	2.98	0.01	0.10	773
Bitcoin	Crypto	0.0014	0.0449	−0.61	5.71	0.01	0.10	773
Panel B: Macro / funding channels (Volatility, Rates, FX, Energy, Metals)								
VIX	Macro / Vol	−0.0006	0.0868	2.02	16.63	0.01	0.10	773
U.S. 10Y Yield	Rates / Macro	0.0021	0.0273	0.28	2.34	0.01	0.10	773
DXY	FX / Dollar	0.0002	0.0049	−0.07	1.80	0.01	0.10	773
EUR/USD	FX	−0.0002	0.0054	0.04	1.95	0.01	0.10	773
WTI Crude	Commodity	0.0005	0.0265	−0.58	2.20	0.01	0.10	773
Gold	Commodity	0.0004	0.0102	−0.49	1.54	0.01	0.10	773
Silver	Commodity	0.0001	0.0217	−0.26	2.92	0.01	0.10	773

3. Methodology

In this section, we formalize the three building blocks of the paper. First, in Section 3.1, we show how we generate a synthetic "post-stress" return panel that mimics a crisis: Asset-level conditional volatility spikes, tail dependence becomes system-wide, and high-beta assets are amplified more than

slow macro transmitters. Second, in Section 3.2, we define the two network objects we compare: A benchmark absolute-correlation backbone and the proposed multiscale coherence–entropy centrality (MCEC) backbone. The two networks differ both in how they create edges and in how they rank node importance. Third, in Section 3.3, we set out the evaluation metrics that correspond to actual supervisory use cases: Stability of systemic hubs, persistence or rewiring of spillover channels, coverage under stress, and the ability to flag known stress transmitters.

Throughout, all inputs are daily log returns $r_{t,i}$ from the 2021–2024 panel described in Section 2, and all constructions are repeated both on the observed (pre-stress) panel and on the synthetic (post-stress) stress panel, so that we can compare the same methodology across two regimes.

3.1. Stress generation: GARCH(1,1) + heavy-tailed t-copula

Supervisors care about a question that is counterfactual by construction: Who transmits stress if a system-wide tail shock hits now? Using a single historical crisis window is unsatisfactory because it bakes in one idiosyncratic episode and fails if markets have not (yet) repriced in that way. Instead, we synthesize a stressed panel of returns that intentionally injects (i) elevated, clustered volatility at the individual-asset level, (ii) joint tail dependence across assets, and (iii) heterogeneous amplification between "fast" risk assets and "slow" macro funding channels.

For each asset i, we estimate a standard GARCH(1,1) model on daily returns $r_{t,i}$ [9]:

$$r_{t,i} = \mu_i + \varepsilon_{t,i}, \tag{3.1}$$

$$\varepsilon_{t,i} = \sigma_{t,i} z_{t,i}, \tag{3.2}$$

$$\sigma^2_{t,i} = \omega_i + \alpha_{1,i}\varepsilon^2_{t-1,i} + \beta_{1,i}\sigma^2_{t-1,i}, \tag{3.3}$$

where $\sigma^2_{t,i}$ is the conditional variance and $z_{t,i}$ are standardized residuals. The triplet $(\omega_i, \alpha_{1,i}, \beta_{1,i})$ captures the long-run variance level, the short-run reaction to shocks (ARCH term), and volatility persistence (GARCH term), respectively. The filtered path $\{\sigma_{t,i}\}_t$ summarizes how "stressed" asset i is expected to be conditional on its own recent history.

Before presenting the network results, we summarize these GARCH fits because they determine how strongly each asset reacts in stress. Table 2 reports the estimated $(\omega_i, \alpha_{1,i}, \beta_{1,i})$ for each asset. Two stylized facts emerge. First, high-beta equity indices and Bitcoin display large unconditional variance and high persistence, consistent with violent but sticky risk repricing. Second, macro/funding assets such as the U.S. 10Y yield or the dollar index show very persistent volatility (large $\beta_{1,i}$) but lower instantaneous shock sensitivity ($\alpha_{1,i}$), reflecting that they transmit stress more slowly and through balance-sheet and funding channels instead of instant risk-off liquidation.

Table 2. GARCH(1,1) estimates and stress multipliers.

Asset	ω_i	$\alpha_{1,i}$	$\beta_{1,i}$	κ_i
S&P 500	4.66×10^{-06}	0.0902	0.8782	3
NASDAQ	5.01×10^{-06}	0.0678	0.9131	3
Dow Jones	1.17×10^{-06}	0.0340	0.9549	3
Russell 2000	6.43×10^{-06}	0.0253	0.9516	3
EuroStoxx 50	2.02×10^{-06}	0.0466	0.9404	3
DAX	1.81×10^{-06}	0.0390	0.9473	3
FTSE 100	1.18×10^{-05}	0.1174	0.7421	3
Nikkei 225	1.40×10^{-05}	0.1932	0.7928	3
Hang Seng	2.05×10^{-05}	0.1044	0.8321	3
Bitcoin	1.91×10^{-05}	0.0133	0.9760	3
Gold	9.46×10^{-08}	0.0000328	0.9990	1.5
Silver	1.49×10^{-05}	0.0184	0.9481	1.5
WTI Crude	2.91×10^{-05}	0.0753	0.8843	1.5
EUR/USD	2.53×10^{-07}	0.0246	0.9670	1.5
VIX	2.90×10^{-03}	0.3513	0.3488	1.5
U.S. 10Y Yield	3.58×10^{-07}	0.0174	0.9803	1.5
DXY (Dollar Idx)	1.55×10^{-07}	0.0309	0.9633	1.5

Although each $z_{t,i}$ is (approximately) i.i.d. over t for a given asset, the vector $z_t = (z_{t,1}, \ldots, z_{t,N})'$ is not cross-sectionally independent. We estimate its cross-sectional correlation matrix $\widehat{R}_z$ and embed it in a low–degrees–of–freedom t-copula [10]. Sampling from this copula produces a multivariate heavy-tailed draw

$$U_t \sim \text{Copula}_t(\widehat{R}_z, \nu), \qquad Z_t^{(\text{stress})} = (t_\nu^{-1}(U_{t,1}), \ldots, t_\nu^{-1}(U_{t,N}))'.$$

This $Z_t^{(\text{stress})}$ is our synthetic "system shock": It creates synchronous tail realizations across many assets, consistent with broad deleveraging, dollar funding squeezes, or energy shocks hitting balance sheets at once [11].

$$r_{t,i}^{(\text{stress})} = \kappa_i \, \sigma_{t,i} \, Z_{t,i}^{(\text{stress})}, \tag{3.4}$$

where $\sigma_{t,i}$ is the conditional volatility from (3.3), and κ_i is an asset-specific amplification factor reported in Table 2. We deliberately set κ_i larger for high-beta "risk" assets (U.S. equities, Bitcoin) and smaller for macro and funding transmitters (rates, FX, oil, gold, VIX). Equation (3.4) therefore generates a "post-stress" panel with three realistic features simultaneously: (i) Volatility clustering survives (through $\sigma_{t,i}$), (ii) extreme co-movements are joint rather than idiosyncratic (through the common t-copula draw), and (iii) risky assets hit harder than funding channels, which matches how real crises propagate across balance sheets.

Figure 1 illustrates the effect of the stress scenario in (3.4). The left block shows the observed daily log returns $r_{t,i}$ for all assets over 2021–2024, while the right block shows the corresponding synthetic stressed returns $r_{t,i}^{(\text{stress})}$ on the same time axis. Under stress, (i) baseline volatility is systematically higher, (ii) tail jumps are larger and more asymmetric, and (iii) large moves occur simultaneously

across multiple assets; this reflects a common heavy–tailed shock scaled by each asset's GARCH(1,1) conditional volatility and injected into the system.

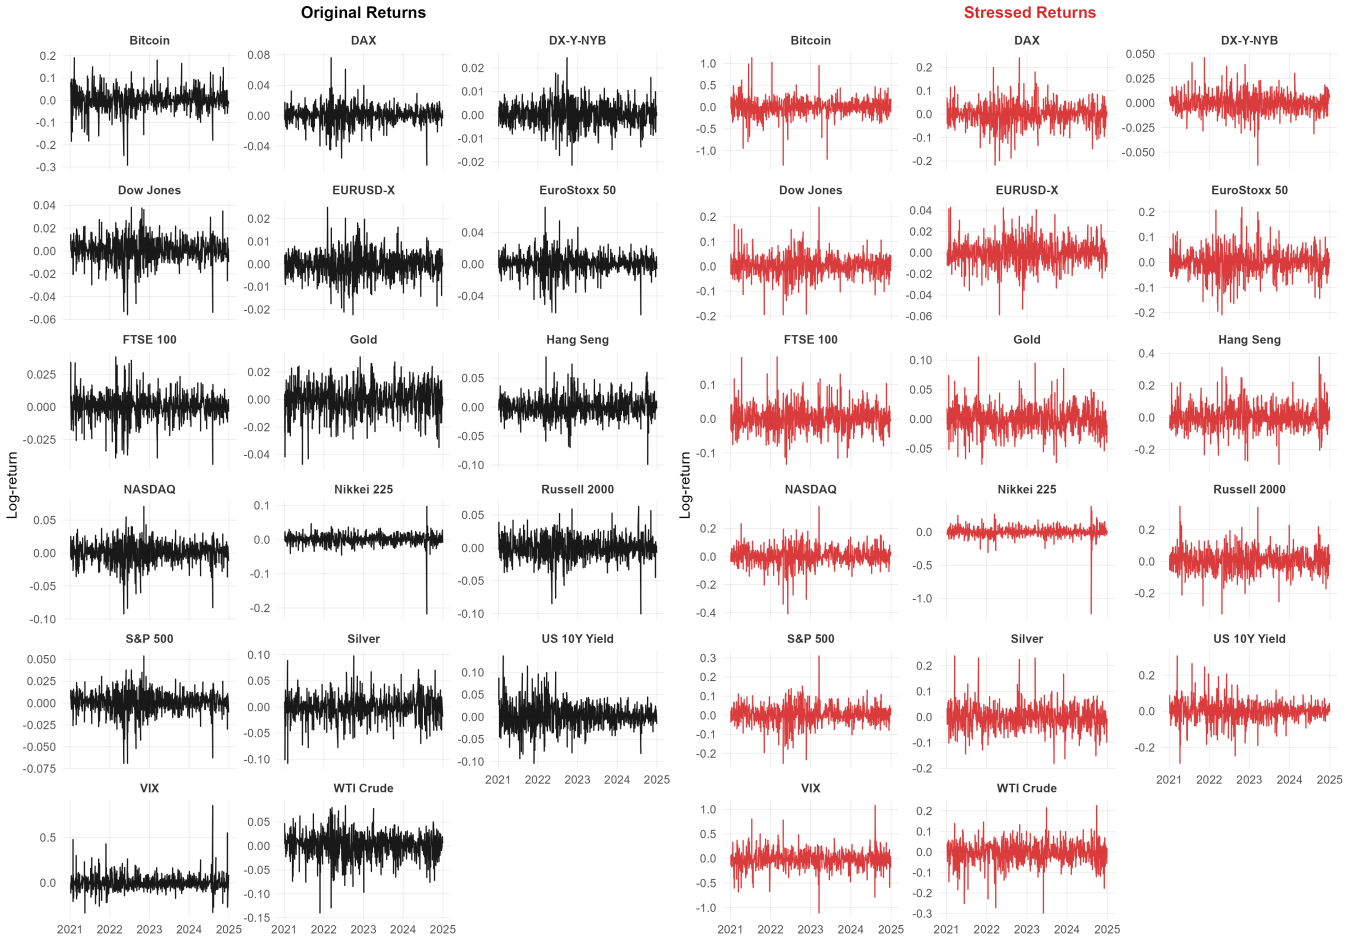

Figure 1. Original vs. stressed return series.

The empirical (pre-stress) panel $\{r_{t,i}\}$ and the synthetic (post-stress) panel $\{r_{t,i}^{(\text{stress})}\}$ are fed in parallel through our network pipeline below.

3.2. Network construction: Correlation vs. multiscale coherence

For each regime (pre-stress and post-stress), we build two undirected weighted graphs on the same set of assets:

1) A traditional absolute-correlation network, which mirrors the standard practice in systemic risk work [1–4];
2) A multiscale Coherence-Entropy Centrality (MCEC) network, which is our proposed stress-aware, scale-aware alternative.

To compare them fairly, we always extract a fixed-density backbone from each network (same target density δ in all regimes). That way, differences in hubs, bridges, and communities cannot be waved away as "you just picked a looser threshold."

3.2.1. Wavelet coherence estimation and statistical validation

For each ordered pair of assets (i, j), we quantify time–frequency dependence using the squared wavelet coherence estimator $R^2_{ij}(t, s)$ of Torrence and Compo [16], Grinsted et al. [17], implemented via the `wtc` routine in `biwavelet`. We use the complex Morlet mother wavelet, which is standard in financial applications because it balances time localization (to capture fast de-risking episodes) and frequency localization (to capture slower macro/liquidity spillovers) [7, 18, 19].

The transform produces $R^2_{ij}(t, s)$ over time index t and pseudo-period s. We discretize s into five adjacent frequency bands corresponding to progressively slower horizons: 2–4, 4–8, 8–16, 16–32, and 32–64 trading days. Economically, the short bands (2–8 days) represent very fast risk-off unwinds and margin calls; the middle bands (8–16 days) capture volatility and funding spillovers; and the long bands (16–64 days) reflect slower repricing in FX, rates, energy, and safe-haven assets such as gold [5, 7].

Boundary effects are handled explicitly. Wavelet coherence tends to be unreliable near the sample edges where the wavelet support is truncated. Following Torrence and Compo [16], `wtc` provides a cone-of-influence (COI) mask. We exclude all observations (t, s) that fall outside the COI when averaging coherence, so that spurious "flares" at the boundaries do not inflate dependence.

For each band b and each pair (i, j), we compute a COI-masked band-average

$$c_{ij,b} = \text{mean of } R^2_{ij}(t, s) \quad \text{over } (t, s) \text{ in band } b \text{ and inside the COI.}$$

We then test whether that average is statistically meaningful. For each pair (i, j), we generate $nrands = 1000$ phase-randomized surrogate series under the null of no cross-series dependence but with each series' own univariate spectrum preserved [17]. We recompute the band-average coherence for every surrogate draw. The $(1 - \alpha)$ quantile (with $\alpha = 0.05$) of this surrogate distribution is used as a critical value. If $c_{ij,b}$ exceeds this Monte Carlo threshold, we mark band b as significant for (i, j) and retain $c_{ij,b}$; otherwise, we set $c_{ij,b} = 0$.

This delivers, for each pair (i, j), a vector $\{c_{ij,1}, \ldots, c_{ij,5}\}$ of band-specific coherence intensities that are (i) COI-corrected to avoid boundary artifacts, and (ii) individually validated via a 1000-draw Monte Carlo test at the 5% level.

3.2.2. Traditional absolute-correlation network

For assets $i \neq j$, we compute their Pearson correlation

$$\rho_{ij} = \text{Corr}(r_{t,i}, r_{t,j}),$$

and define the raw undirected edge weight

$$w^{\text{trad}}_{ij} = |\rho_{ij}|.$$

We then keep only the top K edges (largest w^{trad}_{ij}) so that the retained backbone has density δ.

Node importance in this baseline network is its weighted degree:

$$\text{Deg}^{\text{trad}}_i = \sum_{j \neq i} w^{\text{trad}}_{ij} \cdot \mathbf{1}\{(i, j) \text{ is retained in the backbone}\}.$$

This is the canonical "systemic hub" score in correlation-based connectedness maps.

3.2.3. MCEC network: Statistically validated multi-horizon connectivity

The MCEC network is designed to answer a supervisory question that correlation cannot: Which pairs form statistically robust, multi-horizon spillover channels, and which assets sit on those channels as persistent transmitters?

From the previous subsection, each pair (i, j) has a vector of band-level coherence intensities $\{c_{ij,b}\}_{b=1}^{B}$ after COI masking and Monte Carlo significance filtering. Let $\mathrm{sig}_{ij,b} = \mathbb{1}\{c_{ij,b} > 0\}$ indicate that band b passed significance for (i, j).

We say that (i, j) exhibits a persistent multi-horizon spillover channel if there exists at least one pair of adjacent frequency bands b and $b+1$ such that both are significant, i.e., $\mathrm{sig}_{ij,b} = \mathrm{sig}_{ij,b+1} = 1$. Intuition: In supervisory language, a channel that truly carries stress should not be a razor-thin spike at one exact horizon. It should "bleed" between nearby horizons; for example, a fast equity/volatility unwind (2–8 days) that spills over into funding stress, FX pressure, or oil/gold repricing over several weeks (8–32 days).

We then define the undirected edge weight

$$w_{ij}^{\mathrm{mcec}} = \begin{cases} \mathrm{mean}\{c_{ij,b} : \mathrm{sig}_{ij,b} = \mathrm{sig}_{ij,b+1} = 1\}, & \text{if such adjacent-band persistence exists,} \\ \frac{1}{2}\,\mathrm{mean}\{c_{ij,b} : \mathrm{sig}_{ij,b} = 1\}, & \text{if we observe only isolated single-band significance,} \\ 0, & \text{if no band is significant.} \end{cases}$$

The "adjacent-band" rule acts like a built-in multiple-comparisons control. A one-off spike in a single band is automatically down-weighted by the factor $1/2$, while coherence that is simultaneously significant in neighboring bands is treated as a robust spillover channel and given full weight. In other words, we promote only (i, j) to a strong edge if the pair co-moves significantly across contiguous horizons, not just in one noisy slice.

As in the correlation network, we keep only the top K edges in $W^{\mathrm{mcec}} = [w_{ij}^{\mathrm{mcec}}]$, choosing K, so that the retained MCEC backbone has exactly the same density δ as the traditional backbone. By fixing density, any difference in hubs, bridges, or modular structure is methodological (coherence-based vs. correlation-based), not a thresholding artifact [1, 2, 7].

A node is "systemic" in our sense if it (i) co-moves strongly with the rest of the system and (ii) does so at multiple horizons. For each node i and each band b, we compute

$$s_{i,b} = \sum_{j \neq i} c_{ij,b},$$

which is the total statistically validated (i.e., COI-masked and Monte Carlo significant) coherence mass that i shares with the rest of the system at horizon b. We then form

$$S_i = \sum_{b=1}^{B} s_{i,b} \quad \text{and} \quad p_{i,b} = \frac{s_{i,b}}{S_i} \quad (\text{with } p_{i,b} = 0 \text{ if } S_i = 0).$$

We compute the Shannon entropy of this band profile,

$$H_i = -\sum_{b=1}^{B} p_{i,b} \log p_{i,b},$$

normalized to $[0, 1]$ by dividing by $\log B$. High S_i means node i is strongly connected in aggregate; high H_i means it is active across several adjacent horizons rather than concentrated in a single band.

Finally, we define the multiscale coherence–entropy centrality (MCEC) score of node i as

$$\mathrm{MCEC}_i = \left(\frac{S_i}{\max_k S_k}\right) \cdot \left(\frac{H_i}{\max_k H_k}\right), \tag{3.5}$$

the product of normalized total strength and normalized entropy. An asset attains a high MCEC_i only if it is strongly connected and multi-horizon. We interpret MCEC_i as a stress-transmission capacity: A high-MCEC node can ignite fast panic and sustain slower funding/FX/commodity pressure, which is exactly what supervisors mean by a "systemic transmitter" in practice [7, 11].

We deliberately use the word "multiscale". We are not solving an eigenvector or betweenness problem on a multiplex tensor of fully separate network layers. Instead, we collapse statistically validated coherence across clearly defined adjacent time–frequency bands into one interpretable stress-transmission score. In other words, "multiscale" here means "rewarding breadth across validated neighboring horizons," not "running centrality on a stack of disconnected layers."

By construction, then, the traditional backbone answers "who is highly (linearly) correlated with whom on average?" The MCEC backbone answers "which channels are jointly significant across adjacent horizons, and which nodes sit on those multi-horizon spillover channels?"

3.3. Evaluation metrics

After we build (i) the traditional absolute-correlation backbone and (ii) the MCEC backbone on both the pre-stress and post-stress panels, we evaluate them along four supervisory dimensions. These diagnostics correspond to how a real macroprudential desk would use the output.

First, for each method $m \in \{\mathrm{trad}, \mathrm{mcec}\}$, we collect node scores before and after stress, $\mathrm{Score}_i^{(m,\mathrm{pre})}$ and $\mathrm{Score}_i^{(m,\mathrm{post})}$ (where $\mathrm{Score}_i^{(\mathrm{trad},\cdot)}$ is weighted degree and $\mathrm{Score}_i^{(\mathrm{mcec},\cdot)}$ is MCEC_i). We compute the Spearman rank correlation

$$\rho_{\mathrm{stress}}^{(m)} = \mathrm{Corr}_{\mathrm{Spearman}}\Big(\{\mathrm{Score}_i^{(m,\mathrm{pre})}\}_i,\ \{\mathrm{Score}_i^{(m,\mathrm{post})}\}_i\Big).$$

High $\rho_{\mathrm{stress}}^{(m)}$ means "the same names stay important when stress hits," which matters for early-warning credibility.

Second, let $E^{(m,\mathrm{pre})}$ and $E^{(m,\mathrm{post})}$ be the retained high-weight edges for method m before and after stress. We measure edge persistence with the Jaccard index:

$$J_{\mathrm{stress}}^{(m)} = \frac{\left|E^{(m,\mathrm{pre})} \cap E^{(m,\mathrm{post})}\right|}{\left|E^{(m,\mathrm{pre})} \cup E^{(m,\mathrm{post})}\right|}.$$

Values near 1 mean "the same spillover channels survived; stress just turned the volume up." Values near 0 mean "the network rewired; new bridges carried the shock."

Third, we detect communities in each backbone via modularity clustering [20] and compare the pre-stress vs. post-stress partitions with the variation of information (VI) distance [21]. High VI means the block structure of the system reorganizes under stress (for instance, volatility and FX suddenly become the bridge between global equities and long-term rates).

Fourth, a dashboard is useless if, exactly in stress, half the assets we care about return "NA." We define coverage as the share of assets for which the method produces a finite score before and after stress:

$$\text{Coverage}^{(m)} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\Big\{\text{Score}_i^{(m,\text{pre})} \text{ and } \text{Score}_i^{(m,\text{post})} \text{ are both defined}\Big\}.$$

Finally, policy desks already treat certain benchmarks as "stress transmitters": Broad U.S. equity indices (S&P 500, NASDAQ, Dow Jones), Bitcoin as a high-beta risk proxy, implied equity volatility (VIX), the U.S. dollar / EUR/USD, long-term U.S. yields, and crude oil/gold as global funding and commodity pressure channels [3, 4, 12]. We encode these as positive labels $y_i = 1$. Given any node score, we then compute the area under the ROC curve (AUC): The probability that a randomly chosen labeled "transmitter" scored higher than a randomly chosen unlabeled asset.

We compute AUC using pre-stress scores and again using post-stress scores. Large pre-stress AUC means the method already knows "who will matter" before we shock the system, which is what an early-warning tool is supposed to deliver.

Together, these metrics, hub rank stability, spillover-channel persistence vs. rewiring, coverage, and AUC, form the basis of the comparison in Section 4. They map directly onto supervisory questions: Who will transmit stress? Do the same bridges carry the shock, or do new cross-market bridges (volatility, FX, energy) suddenly light up? Can we still score everyone we care about in the middle of the storm? And does the method elevate the usual suspects (VIX, dollar strength, long-term yields, oil) without being told in advance?

4. Results

We compare (i) a traditional absolute-correlation backbone ("Traditional") and (ii) the proposed multiscale coherence–entropy centrality backbone ("MCEC"), each constructed as in Section 3. Both networks are estimated twice: Once on the observed daily return panel ("Pre") and once on the synthetic heavy-tailed stressed panel generated via the GARCH(1,1) + t-copula procedure ("Post") in Section 3.1. In both cases, we enforce a common edge density (≈ 0.257), so differences in topology, hubs, and connectivity are methodological, not thresholding artifacts.

We proceed in five steps: First, we study how the network topology changes under stress (Section 4.1). Second, we examine how node centrality reallocates across assets (Section 4.2). Third, we compare absolute correlation to multiscale coherence (Section 4.3). Fourth, we identify which assets emerge as stress transmitters (Section 4.4). Finally, we evaluate supervisory diagnostics such as rank stability, edge persistence, and early-warning classification power (Section 4.5).

4.1. Network topology under stress

Figures 2 and 3 show the network backbones before and after stress for the Traditional and MCEC constructions, respectively. In each case, we keep only the top-weighted links at the common density, and we hold the node layout fixed between Pre and Post so that any shift in edges is interpretable as stress-driven reconfiguration rather than a plotting artifact.

In the Traditional backbone (Figure 2), the dense equity cluster is visually stable across regimes. The core U.S. and global stock indices continue to sit in one tight block both Pre and Post. In other

words, the absolute-correlation view continues to see "equities with equities," even after a systemic shock.

In the MCEC backbone (Figure 3), edges exist only if wavelet coherence is statistically significant and persistent across adjacent frequency bands; isolated one-band hits are downweighted. Under stress, bridges linking volatility (VIX), FX (e.g., EUR/USD and DXY), and commodities (oil, gold) to global equity benchmarks visibly thicken. Put differently, in the stressed regime, the cross-market spillover structure is no longer "equity-only": the major bridges explicitly route through volatility, dollar funding, and energy/safe-haven assets.

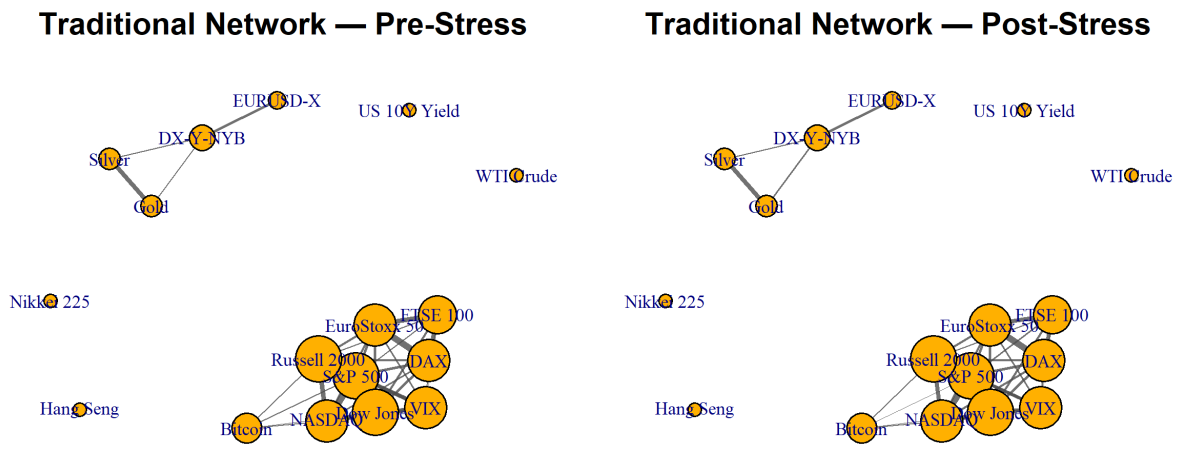

Figure 2. Traditional absolute-correlation networks, Pre vs. Post. Edge width is $|\rho|$; node size reflects weighted degree.

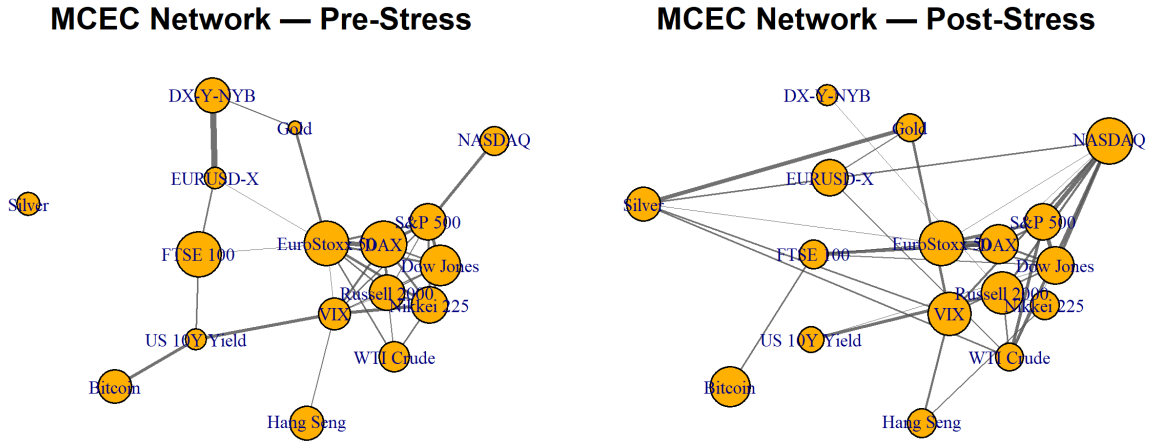

Figure 3. MCEC networks, Pre vs. Post. Edges require significant, persistent multiscale coherence; node size reflects MCEC centrality.

These visual differences have a numerical counterpart. Table 3 reports three global descriptors for each method and regime: (i) Modularity (how sharply the network splits into internally cohesive

communities), (ii) density (matched by construction), and (iii) average weighted shortest-path length (larger values mean that effective spillover paths become longer/weaker).

Modularity rises in both constructions once we inject stress: From 0.211 to 0.223 for the Traditional network and from 0.146 to 0.242 for MCEC. Higher modularity means that, after stress, the system does not behave like one homogeneous "risk blob," but instead separates more cleanly into equity, FX/rates, and commodity/liquidity blocs. Furthermore, the average shortest-path length increases sharply (Traditional: 1.316 $\rightarrow$ 1.945; MCEC: 1.316 $\rightarrow$ 2.038), meaning the thick cross-block bridges that tightly connected markets in calm times have weakened or rerouted.

Notably, the jump in modularity and path length is larger for MCEC. This matches the visual evidence: MCEC explicitly reroutes transmission through VIX/FX/commodities in the stressed regime, rather than continuing to treat equities as the single organizing axis.

Table 3. Global network metrics (pre- vs. post-stress). Modularity is Newman–Girvan modularity of detected communities; the average path length is computed using inverse weight as distance.

Network	Modularity	Density	Avg. weighted shortest path
Traditional (Pre)	0.211	0.257	1.316
Traditional (Post)	0.223	0.257	1.945
MCEC (Pre)	0.146	0.257	1.316
MCEC (Post)	0.242	0.257	2.038

4.2. Centrality and rank stability

Supervisors care about which nodes sit "in the middle of the system", and whether those same nodes remain influential once stress hits. In the Traditional network, node importance is weighted degree (sum of retained $|\rho|$ edges). In the MCEC network, node importance is the coherence–entropy centrality: A node scores high only if it is strongly connected and active across multiple horizons.

Figure 4 plots, for each asset, its Pre-stress score (horizontal axis) against its Post-stress score (vertical axis), separately for Traditional (right panel) and MCEC (left panel). Points on the 45° line are equally important before and after stress; points above the line gain systemic relevance under stress.

In the MCEC panel, canonical stress transmitters, VIX, EUR/USD, gold, and silver, shift upward. That is, once the system is shocked, volatility, FX, and safe/liquidity assets become more central to cross-market spillovers. Furthermore, broad U.S. benchmarks (S&P 500, NASDAQ, Dow Jones) and Bitcoin remain near the top, confirming that they continue to anchor the global "risk-on" complex.

In the Traditional panel, almost every point moves down: Weighted degree collapses in magnitude under stress. The dense equity cluster exists structurally (Figure 2), but its aggregate "influence mass" thins. The method's ability to discriminate which specific nodes matter is degraded when we move to the stressed panel.

Table 4 reports the underlying scores. Before stress, MCEC ranks global equity indices (S&P 500, Dow Jones, NASDAQ), Bitcoin, and oil as highly central. After stress, those assets remain, and traditional stress barometers move sharply higher: VIX rises from 0.715 to 0.928, EUR/USD from 0.666 to 0.826, and gold edges higher. This says: In the crisis-like regime, the bridges that matter are not purely "equity to equity," but "equity $\leftrightarrow$ volatility/FX/commodities."

In contrast, the Traditional weighted degree numbers collapse for essentially everyone (e.g.,

S&P 500: 4.669 → 0.165). The ranking tilts somewhat toward FX and metals in relative terms, but it does so mostly because the equity block has been flattened.

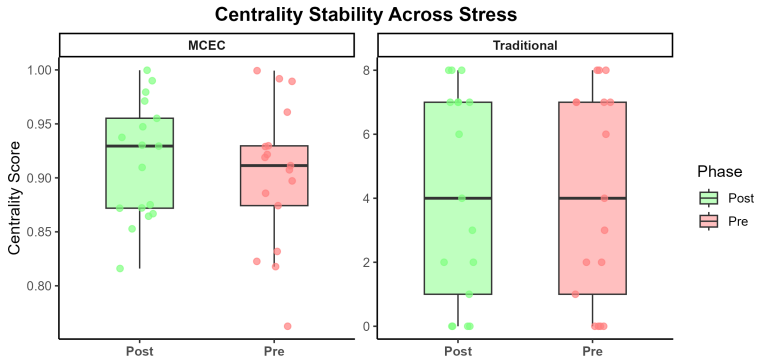

Figure 4. Pre vs. Post node centrality. Left: MCEC coherence–entropy centrality. Right: Traditional weighted degree. Points above the diagonal gain systemic relevance in the stressed regime.

Table 4. Node centrality (Pre vs. Post).

Asset	MCEC (Pre)	MCEC (Post)	Degree (Pre)	Degree (Post)
S&P 500	0.939	0.920	4.669	0.165
NASDAQ	0.815	0.830	4.301	0.148
Dow Jones	0.912	0.880	4.335	0.154
Russell 2000	0.775	0.743	4.032	0.151
EuroStoxx 50	0.877	0.879	4.018	0.153
DAX	0.875	0.833	3.623	0.148
FTSE 100	0.735	0.661	3.621	0.151
Nikkei 225	0.680	0.637	3.900	0.152
Hang Seng	0.604	0.583	3.052	0.134
Bitcoin	0.907	0.954	3.328	0.134
VIX	0.715	0.928	3.316	0.134
Gold	0.377	0.394	3.352	0.148
Silver	0.467	0.484	2.890	0.131
EUR/USD	0.666	0.826	2.543	0.135
WTI Crude	0.804	0.787	2.541	0.131
U.S. 10Y Yield	0.707	0.708	2.319	0.117
DXY (Dollar Index)	0.784	0.694	2.160	0.119

Notes: MCEC = multiscale coherence–entropy centrality (higher values indicate nodes that are both strongly connected and active across multiple horizons). Degree = weighted degree in the absolute-correlation backbone. Pre = network estimated on observed returns; Post = network estimated on the synthetic stressed panel.

4.3. Correlation vs. multiscale connectivity

Traditional networks are built from absolute Pearson correlations $|\rho_{ij}|$. MCEC edges, in contrast, require statistically significant and adjacent-band wavelet coherence, meaning the two series move together in a way that persists across nearby horizons, not just in a single noisy band.

Figure 5 shows the Traditional $|\rho|$ matrices Pre and Post stress, plus their difference. The block structure is very familiar: Equities move with equities, and macro/funding assets (rates, FX, oil, gold) largely sit in their own block. After stress, most of that structure survives; the difference (Post − Pre) panel mostly reflects mild strengthening or weakening in magnitude, not a wholesale change in which pairs are considered important. This is consistent with the fact that, for the Traditional backbone, the set of top edges is almost identical Pre and Post.

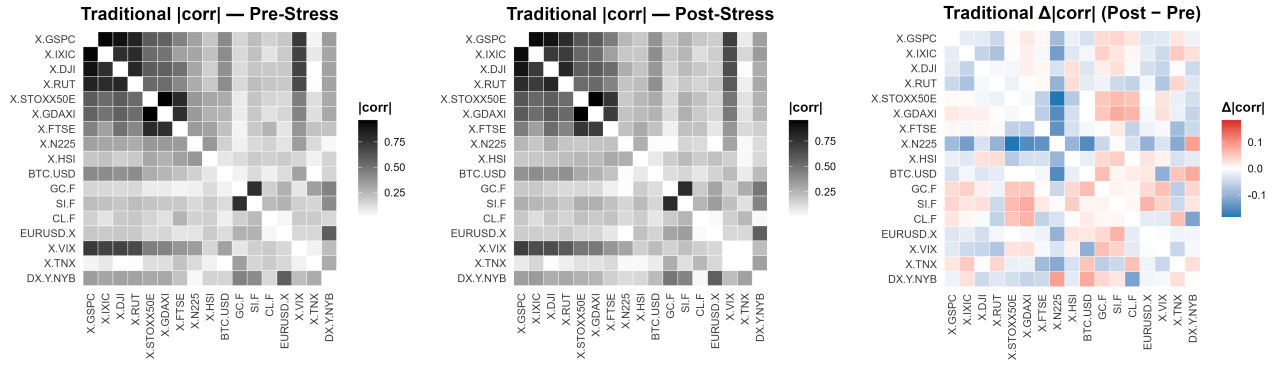

Figure 5. Traditional absolute-correlation matrices. Left: Pre-stress $|\rho_{ij}|$. Middle: Post-stress $|\rho_{ij}|$. Right: Post − Pre.

Figure 6 shows the analogous object for MCEC: The multiscale coherence weights w_{ij}^{mcec}, which average statistically significant coherence over adjacent frequency bands and downweight isolated one-band hits. Pre-stress, the heaviest MCEC links cluster global equities, Bitcoin, and crude oil, a "risk-on" block that mixes fast panic channels and slower macro/commodity channels. Post-stress, the heaviest links move: VIX, EUR/USD, and gold become reinforced bridges into the equity complex, and some pure equity–equity links weaken. The rightmost panel (Post − Pre) in Figure 6 makes exactly this shift visible.

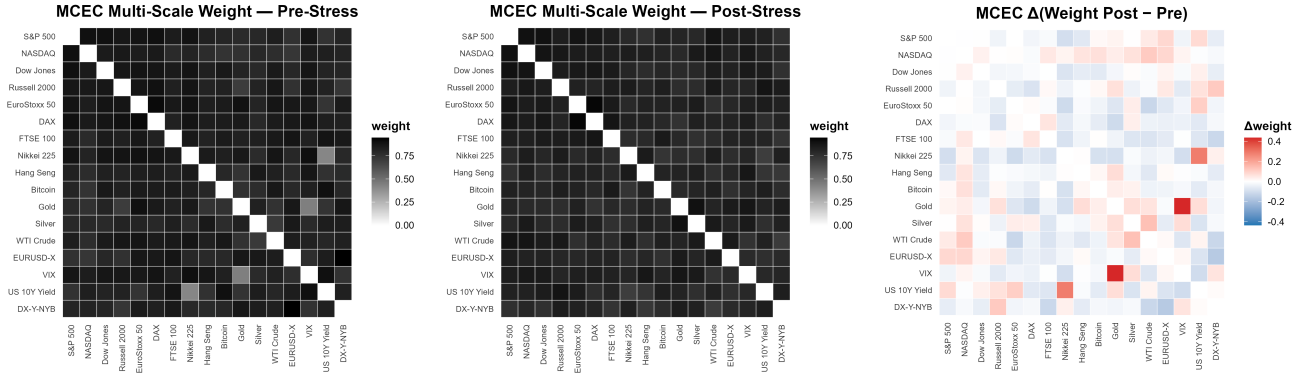

Figure 6. MCEC multiscale coherence weights. Left: Pre-stress w_{ij}^{mcec}. Middle: Post-stress w_{ij}^{mcec}. Right: Post − Pre.

Quantitatively, in the MCEC network, only about 20% of the top-weight edges survive intact into the stressed regime, and the detected communities change substantially. In other words, MCEC rewires under stress. The Traditional backbone, by contrast, is structurally inert: It keeps essentially the same high-$|\rho|$ edges even though their ability to explain joint stress dynamics weakens.

4.4. Which assets transmit stress?

The question supervisors actually ask is: Who transmits stress across markets once the shock hits? To answer that, we track how each node's MCEC centrality changes between the Pre and Post networks.

Figure 7 plots, for each asset, a "dumbbell": The left marker is its Pre-stress MCEC score, the right marker is its Post-stress score, and the horizontal segment shows the change. Long rightward segments indicate assets that become structurally more important under stress.

Two patterns emerge. First, traditional stress barometers, VIX, EUR/USD, gold, and silver, all shift right. VIX in particular becomes one of the strongest multi-horizon bridges in the stressed network: It is not merely spiking, it is connecting markets. Second, broad U.S. equity benchmarks (S&P 500, NASDAQ, Dow Jones) and Bitcoin remain highly central even after stress, confirming they still anchor global "risk-on" behavior. However, they no longer monopolize the center: FX and commodities join the core.

This implies that spillover in our stressed regime is not "equity contagion only." It explicitly routes through volatility, dollar funding pressure, and energy/safe-haven assets.

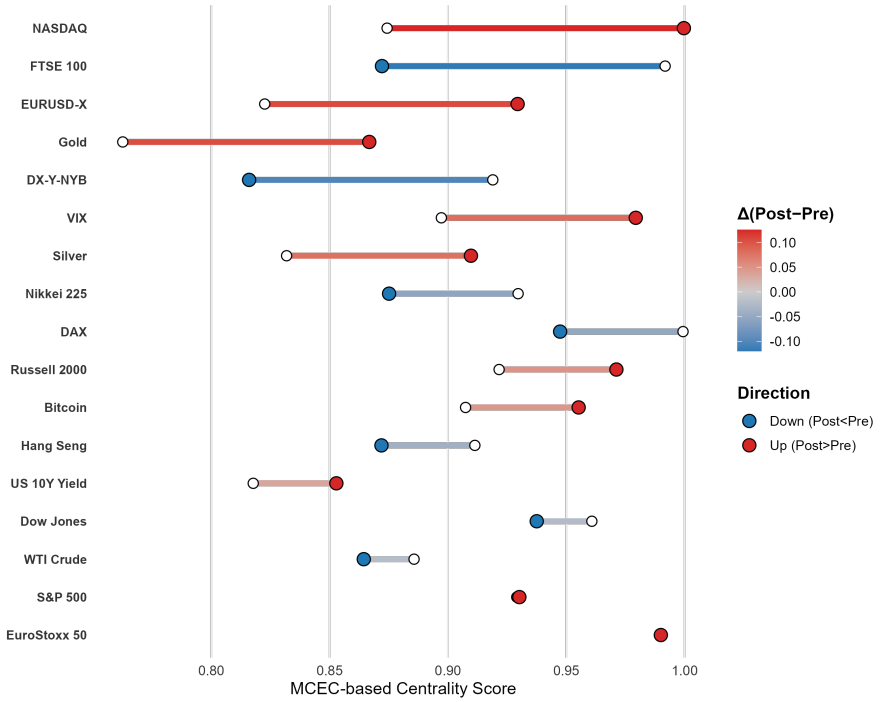

Figure 7. Shift in MCEC centrality under stress. Each line connects an asset's Pre-stress MCEC score (left dot) to its Post-stress score (right dot).

4.5. Global network diagnostics and supervisory performance

Finally, we evaluate both constructions on supervisory diagnostics introduced in Section 3.3: (i) Rank stability (Spearman correlation between Pre and Post node scores), (ii) edge persistence (Jaccard overlap of retained backbone edges Pre vs. Post), (iii) topology reorganization (variation of information, VI, between the Pre and Post community partitions), (iv) coverage (fraction of benchmark assets that still receive a usable score in stress), and (v) classification power (AUC) in ranking canonical stress transmitters such as broad U.S. equities, VIX, the U.S. dollar, and crude oil.

Table 5 reports these metrics when "positives" for AUC are generic global hubs (large equity indices, oil, dollar proxy). Table 6 repeats the exercise using an explicitly stress-oriented transmitter set (broad U.S. equities, VIX, the U.S. dollar, crude oil).

Three points matter for policy:

First, MCEC is supposed to change under stress. Its node ranking reorders (Spearman $\rho_{\text{stress}} = 0.284$), only about 20% of its top edges survive (Jaccard = 0.207), and its block structure changes materially (VI = 1.974). This is the intended behavior: MCEC is mapping how stress rewires cross-market spillovers. The Traditional backbone, on paper, looks stable ($\rho_{\text{stress}} = 1.000$, Jaccard = 1.000, VI = 0.000), but that "stability" is misleading because the absolute magnitudes of its degree scores collapse under stress (Table 4). In other words, the Traditional graph freezes the calm-time equity block, while MCEC actively reveals stress-driven rerouting through volatility, FX, and commodities.

Second, pre-stress early-warning power is stronger for MCEC. In Table 6, MCEC achieves a pre-stress AUC of 0.917, compared with 0.742 for the traditional network, when the positives are classic stress transmitters (U.S. equity benchmarks, VIX, dollar strength, crude oil). That means before the

shock, MCEC elevates exactly the assets supervisors would watch in a liquidity event.

Third, coverage is 1.000 in all cases. Every monitored benchmark keeps a usable score in both regimes, so either method can feed a live dashboard.

Table 5. Global diagnostics (generic hub labels).

Method	AUC (Pre)	AUC (Post)	Rank stab.	Edge persist.	Comm. shift	Coverage
Traditional	0.381	0.381	1.000	1.000	0.000	1.000
MCEC	0.357	0.405	0.284	0.207	1.974	1.000

Notes: AUC = ability to rank designated "hub" assets above others. Rank stab. = Spearman correlation of node scores (Pre vs. Post). Edge persist. = Jaccard overlap of retained edges. Comm. shift = variation of information between community partitions. Coverage = share of benchmark assets with a valid score in both regimes. The "hub" set here is broad equity indices, oil, and a dollar proxy.

Table 6. Global diagnostics (stress-transmitter labels).

Method	AUC (Pre)	AUC (Post)	Rank stab.	Edge persist.	Comm. shift	Coverage
Traditional	0.742	0.742	1.000	1.000	0.000	1.000
MCEC	0.917	0.717	0.284	0.207	1.974	1.000

Notes: Here the "transmitter" set is broad U.S. equity indices, VIX, the U.S. dollar, and crude oil. Higher AUC (Pre) means better ex-ante identification of those transmitters. Other columns defined as in Table 5.

Four conclusions follow:

(i) Both backbones become more modular and develop longer effective spillover paths under stress (Figures 2–3, Table 3). (ii) MCEC explicitly reallocates systemic centrality toward volatility, FX, and energy/precious metals while keeping global equity benchmarks in the core (Figures 4–7, Table 4). (iii) MCEC's backbone rewires under stress (low Jaccard, high VI), whereas the Traditional backbone remains structurally fixed and risks understating how stress propagates (Figures 5–6, Tables 5–6). (iv) Ex ante, MCEC ranks canonical stress transmitters with higher AUC, which gives it genuine early-warning value for macroprudential surveillance.

In summary, MCEC does not just report high co-movement; it identifies which cross-market bridges become systemically relevant when the system is shocked, notably volatility (VIX), FX/dollar pressure, crude oil, and gold, and shows how those bridges reroute pressure between equities, funding, and commodities.

5. Conclusions

We set out to answer a supervisory question that standard correlation backbones are not designed to answer: Who will transmit stress, through which channels and horizons, and how does that transmission map change in a crisis-like regime? To address this, we introduced a MCEC framework that (i) draws edges only where wavelet coherence is statistically significant across adjacent frequency bands, (ii) measures node importance via an entropy-weighted multi-horizon strength, and (iii) benchmarks the calm regime against a stressed panel constructed using GARCH(1,1) volatilities and a heavy-tailed t-copula shock.

Three findings speak directly to macroprudential use:

First, network topology under stress is not stable, and MCEC makes that instability visible. Both the traditional correlation backbone and MCEC become more modular and exhibit longer effective paths once stress is injected, indicating that spillovers no longer propagate through a single dense "all-risk" core but instead through more segmented blocs. However, only MCEC shows genuine rewiring: Its high-weight connections and community structure change, and new bridges form between asset classes such as equities, volatility, FX, and commodities. In practical terms, MCEC reveals which cross-market channels carry pressure in the stressed regime rather than assuming that the pre-stress equity cluster remains the dominant conduit. Second, MCEC isolates a persistent (and supervisory-relevant) transmitter set and flags it in advance. The same compact set of markets, broad U.S. equity indices, Bitcoin, VIX, EUR/USD, and the dollar complex, long-term yields, crude oil, and gold, emerges as structurally central across regimes. These are the assets policymakers treat as global risk thermometers, funding stress points, and flight-to-quality anchors. Importantly, MCEC ranks these transmitters highly before the stress shock is applied, and does so more accurately than a correlation-based backbone. In other words, the method does not just explain ex-post contagion; it has ex-ante classification power. By construction, this is because MCEC rewards nodes that are strongly connected and active across multiple horizons, rather than simply "highly correlated on average". Third, MCEC turns structural change into an operational signal. We showed that a small set of scalar diagnostics, rank stability of node scores across regimes, persistence of backbone edges across regimes, and the variation of information between community partitions, can be monitored as quantitative alarms for system reorganization. When these diagnostics move, the spillover map has rerouted: Transmission is no longer "equity-to-equity" but instead runs through volatility (VIX), FX/dollar funding pressure, and energy/precious metals. This gives supervisors a direct escalation rule: Intervene, or at least intensify monitoring, once the dominant bridges shift toward volatility and FX.

The policy message is pragmatic. A correlation backbone is useful as a baseline map of average co-movement, but it is largely blind to which horizons carry contagion and which cross-asset channels become dominant once markets come under stress. By validating edges only when coherence is jointly significant across adjacent scales, and by rewarding breadth across horizons rather than single-band spikes, MCEC surfaces exactly the routes regulators worry about in crises: Volatility, FX/dollar funding, long-term rates, and energy/precious metals. It also turns those routes into quantitative dashboard metrics.

Operationally, a supervisor can (i) compute banded coherence on a rolling basis, (ii) construct the fixed-density MCEC backbone, (iii) track node-level MCEC centrality together with rank stability, edge persistence, and community-turnover metrics, and (iv) escalate supervisory attention when those diagnostics indicate that stress transmission has migrated into volatility and FX.

There are natural extensions. First, adding directionality (e.g., Granger-style or transfer-entropy measures) would enable us to distinguish stress senders from stress receivers. Second, moving from daily to intraday sampling would capture liquidity spirals that matter on regulatory time scales of hours, not days. Third, one can stress-test policy counterfactuals, such as FX swap lines or collateral backstops, directly in this network by attenuating specific bridges and recomputing the topology under that intervention.

Even in its current form, MCEC offers a stress-aware, scale-aware map of systemic spillovers. It identifies which cross-market bridges become systemically relevant when the system is shocked,

notably volatility (VIX), FX/dollar pressure, crude oil, and gold, and it quantifies when those bridges begin to dominate transmission between equities, funding channels, and commodities.

Author contributions

The author carried out all tasks in accordance with the CRediT taxonomy: Conceptualization; Methodology; Software; Validation; Formal analysis; Investigation; Data curation; Writing—original draft; Writing—review & editing; Visualization.

Use of Generative-AI tools declaration

The author used a generative AI tool only for language polishing.

Conflict of interest

The author declares no conflict of interest in this paper.

References

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

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

3 T. Adrian, M. K. Brunnermeier, CoVaR, *Am. Econ. Rev.*, **106** (2016), 1705–1741. https://doi.org/10.1257/aer.20120555

4 C. Brownlees, R. F. Engle, SRISK: A conditional capital shortfall measure of systemic risk, *Rev. Financ. Stud.*, **30** (2017), 48–79. https://doi.org/10.1093/rfs/hhw060

5 K. J. Forbes, R. Rigobon, No contagion, only interdependence: Measuring stock market comovements, *J. Finan.*, **57** (2002), 2223–2261.

6 F. Longin, B. Solnik, Extreme correlation of international equity markets, *J. Finan.*, **56** (2001), 649–676.

7 J. Baruník, T. Křehlík, Measuring the frequency dynamics of financial connectedness and systemic risk, *J. Financ. Econom.*, **16** (2018), 271–296. https://doi.org/10.1093/jjfinec/nby001

8 L. Aguiar-Conraria, M. J. Soares, Business cycle synchronization and the Euro: A wavelet analysis, *J. Macroecon.*, **33** (2011), 477–489. https://doi.org/10.1016/j.jmacro.2011.02.005

9 T. Bollerslev, Generalized autoregressive conditional heteroskedasticity, *J. Econometrics*, **31** (1986), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

10 S. Demarta, A. J. McNeil, The *t* copula and related copulas, *Int. Stat. Rev.*, **73** (2005), 111–129. https://doi.org/10.1111/j.1751-5823.2005.tb00254.x

11 V. Acharya, R. Engle, M. Richardson, Capital shortfall: A new approach to ranking and regulating systemic risks, *Am. Econ. Rev.*, **102** (2012), 59–64. https://doi.org/10.1257/aer.102.3.59

12 R. F. Engle, E. Ghysels, B. Sohn, Stock market volatility and macroeconomic fundamentals, *Rev. Econ. Stat.*, **95** (2013), 776–797.

13 Yahoo Finance, Historical daily adjusted close prices and benchmark index levels for global equity indices, Bitcoin, macro factors, and commodities (dataset used for the 2021–2024 panel). Available at: `https://finance.yahoo.com/`. Accessed 17 October 2025.

14 D. A. Dickey, W. A. Fuller, Distribution of the estimators for autoregressive time series with a unit root, *J. Am. Stat. Assoc.*, **74** (1979), 427–431. https://doi.org/10.2307/2286348

15 D. Kwiatkowski, P. C. Phillips, P. Schmidt, Y. Shin, Testing the null hypothesis of stationarity against the alternative of a unit root, *J. Econometrics*, **54** (1992), 159–178. https://doi.org/10.1016/0304-4076(92)90104-Y

16 C. Torrence, G. P. Compo, A practical guide to wavelet analysis, *B. Am. Meteorol. Soc.*, **79** (1998), 61–78. https://doi.org/10.1175/1520-0477(1998)079¡0061:APGTWA¿2.0.CO;2

17 A. Grinsted, J. C. Moore, S. Jevrejeva, Application of the cross wavelet transform and wavelet coherence to geophysical time series, *Nonlin. Processes Geophys.*, **11** (2004), 561–566. https://doi.org/10.5194/npg-11-561-2004

18 R. Gençay, F. Selçuk, B. Whitcher, *An introduction to wavelets and other filtering methods in finance and economics*, 2002, Academic Press. https://doi.org/10.1016/B978-0-12-279670-8.X5000-9

19 A. Rua, L. C. Nunes, International comovement of stock market returns: A wavelet analysis, *J. Empir. Financ.*, **16** (2009), 632–639. https://doi.org/10.1016/j.jempfin.2009.02.002

20 M. E. J. Newman, M. Girvan, Finding and evaluating community structure in networks, *Phys. Rev. E*, **69** (2004), 026113. https://doi.org/10.1103/PhysRevE.69.026113

21 M. Meilă, Comparing clusterings—An information based distance, *J. Multivariate Anal.*, **98** (2007), 873–895. https://doi.org/10.1016/j.jmva.2006.11.013

AIMS Press