Chapter 1

Mapping Cross-Asset Information Flow and Network Vulnerability: A Directed Graph Approach Across Equities, Cryptocurrencies, Commodities, and Macro Indicators

Cemal Öztürk[1]

Mehmet Polat[2]

Abstract

This study maps cross-asset information flow and network vulnerability across major equities, cryptocurrencies, commodities, and macro/FX indicators by translating Granger-causality relations into a directed network. Using daily data for 23 instruments (seven equity indices, five cryptocurrencies, six commodities, five macro/FX series) from April 13, 2020 to December 30, 2025 (1,485 observations), the sample is aligned to weekdays to avoid weekend-induced bias between 24/7 crypto and business-day markets before computing log returns. Stationarity is verified via a triangulation strategy (ADF, KPSS, and Zivot–Andrews). Pairwise linear Granger causality tests are then applied to all 506 ordered pairs (maximum lag of 5), and statistically significant links are mapped to a directed graph. Network roles are quantified using in-/out-degree, betweenness, eigenvector centrality, and density, while time variation is assessed through rolling-window estimation (250-day window, 20-day step) and sub-period regime networks. The full-sample network contains 181 significant directed links (density = 0.358), indicating moderately high cross-asset connectivity. EURUSD and HSI are the most interconnected nodes (total degrees 25 and 24), SP500 is the strongest transmitter (out-degree 14), and HSI acts as the key bridge (betweenness 0.138). Rolling density ranges from 0.144 to 0.374 (mean 0.241; σ 0.055), implying regime-dependent spillover intensity and higher vulnerability when

1 Res. Asst. Dr., Department of Economics, Iğdır University, cemal.ozturk@igdir.edu.tr, ORCID ID: 0000-0003-3850-7416

2 Assoc. Prof. Dr., Department of Economics, Iğdır University, mehmetpolat@igdir.edu.tr, ORCID ID: 0000-0002-6930-1499

https://doi.org/10.58830/ozgur.pub1259.c5125

density exceeds 0.296. Sub-period density peaks in 2022 (0.283) and remains elevated in 2023–2025 (0.269). The results imply that diversification benefits weaken during dense regimes; monitoring high out-degree and bridge nodes can provide early warning of contagion pathways.

1. Introduction

The global financial system has undergone a structural transformation in recent years, driven by the proliferation of digital assets, the unprecedented monetary policy interventions during COVID-19, and the cascading geopolitical shocks that followed. Understanding how information flows across distinct asset classes—equities, cryptocurrencies, commodities, and macroeconomic indicators—is a practical necessity for portfolio managers, central bankers, and regulators seeking to anticipate contagion pathways before they materialize.

The traditional approach to measuring cross-asset dependencies relies on correlation analysis, which is inherently symmetric and captures only contemporaneous co-movement. Granger (1969) causality testing offers a directional alternative: it identifies whether the past values of one asset contain predictive information about another, establishing a lead–lag structure that correlation alone cannot reveal. Crucially, when these pairwise directional relationships are aggregated into a directed graph—where each node represents an asset and each edge represents a statistically significant causal link—the resulting topology provides a map of information highways across the financial system. This Granger-to-graph mapping is the methodological core of this chapter.

This chapter makes three contributions. First, we construct a comprehensive cross-asset Granger causality network encompassing 23 financial instruments spanning four distinct classes using daily data from 2020-04-13 through 2025-12-30. Second, we characterize the network's architecture through graph-theoretic metrics—degree centrality (Freeman, 1978), betweenness centrality (Freeman, 1977), eigenvector centrality (Bonacich, 1987), and network density (Wasserman & Faust, 1994). Third, we deploy a rolling-window framework to track how the network's topology evolves across market regimes.

The remainder of this chapter proceeds as follows. Section 2 reviews the relevant literature. Section 3 describes the data and methodological framework—including why stationarity testing is a necessary precondition for Granger causality, how the Granger test itself operates, and how its results are

transformed into a directed graph. Section 4 presents the empirical findings. Section 5 discusses implications. Section 6 concludes.

2. Theoretical Background and Literature Review

2.1. Financial Network Analysis

The application of network science to financial markets gained momentum after the 2008–2009 crisis. Billio et al. (2012) demonstrated that Granger causality networks among financial institutions could serve as an early-warning indicator of systemic risk. Diebold and Yılmaz (2014) proposed an alternative framework based on forecast error variance decompositions. While the Diebold–Yılmaz approach offers a continuous connectedness measure, the Granger causality network approach provides a cleaner binary identification of significant causal links, which is more amenable to graph-theoretic analysis.

2.2. Cross-Asset Information Transmission

Early studies focused on equity–equity spillovers, documenting the dominant role of the U.S. market (Eun & Shim, 1989; Hamao et al., 1990). Corbet et al. (2018) found limited initial crypto integration, but subsequent research documented increasing connectedness. Tang and Xiong (2012) studied commodity financialization, while Bouri et al. (2021) showed that extreme conditions amplify cross-asset spillovers.

2.3. Granger Causality in Financial Applications

Granger (1969) formalized the notion that variable X Granger-causes Y if past values of X improve the prediction of Y beyond Y's own history. Extensions include nonlinear Granger causality (Diks & Panchenko, 2006), frequency-domain variants (Breitung & Candelon, 2006), and quantile approaches (Troster, 2018). This chapter employs the standard linear framework for transparency and interpretability.

3. Data and Methodology

3.1. Data Description

We collect daily closing prices for 23 financial instruments spanning four asset classes from Yahoo Finance. The sample period extends from 2020-04-13 to 2025-12-30, yielding 1,485 trading-day observations. Table 1 details the asset universe. A critical data alignment issue arises because cryptocurrency markets operate continuously (24/7), while equity, commodity, and FX markets trade only on business days. If weekend observations are retained, forward-

filling equity prices creates artificial zero returns on Saturdays and Sundays while cryptocurrencies exhibit genuine price changes—introducing systematic bias into the Granger causality tests. To address this, we restrict the sample to business days (Monday through Friday) before computing returns. All prices are then converted to continuously compounded (log) returns:

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

where $P_{i,t}$ is the closing value of asset i on day t. Small gaps from public holidays are forward-filled up to a maximum of three consecutive days; remaining missing observations are dropped.

The selection of 23 instruments across four asset classes is designed to capture the major channels through which information propagates in the modern financial system. The equity indices represent the world's largest capital markets—spanning the Americas (S&P 500, NASDAQ), Europe (FTSE 100, DAX), and Asia-Pacific (Nikkei 225, Hang Seng, KOSPI)—and collectively account for a substantial share of global equity market capitalization. The cryptocurrency cluster includes the five largest digital assets by market capitalization, representing a market segment that has grown from a niche speculative vehicle to a trillion-dollar asset class over the sample period. The commodity selection covers energy (WTI crude oil, natural gas), precious metals (gold, silver), industrial metals (copper), and agriculture (wheat), each serving as a benchmark for distinct segments of the real economy. The macro indicators—VIX (investor fear gauge), DXY (dollar strength), EUR/USD, USD/JPY, and the U.S. 10-year Treasury yield—function as transmission channels between monetary policy decisions and asset prices. Table 1 provides the complete asset universe with ticker symbols and classifications.

Table 1. Asset Universe and Classification (23 Assets)

Asset	Ticker	Class	Description
DAX	^GDAXI	Equity	German blue-chip index
FTSE100	^FTSE	Equity	U.K. blue-chip index
HSI	^HSI	Equity	Hong Kong index
KOSPI	^KS11	Equity	South Korean index
NASDAQ	^IXIC	Equity	U.S. tech-heavy index
Nikkei225	^N225	Equity	Japanese blue-chip index
SP500	^GSPC	Equity	U.S. large-cap index
BNB	BNB-USD	Crypto	Binance native token
Bitcoin	BTC-USD	Crypto	Largest cryptocurrency
Ethereum	ETH-USD	Crypto	2nd largest cryptocurrency
Solana	SOL-USD	Crypto	High-throughput blockchain
XRP	XRP-USD	Crypto	Payment-focused crypto
Copper	HG=F	Commodity	Copper futures
Gold	GC=F	Commodity	Gold futures (COMEX)
NatGas	NG=F	Commodity	NatGas futures
Oil_WTI	CL=F	Commodity	Crude oil futures
Silver	SI=F	Commodity	Silver futures
Wheat	ZW=F	Commodity	Wheat futures (CBOT)
DXY	DX-Y.NYB	Macro	U.S. Dollar Index
EURUSD	EURUSD=X	Macro	Euro/Dollar rate
US10Y	^TNX	Macro	10-Year Treasury yield
USDJPY	JPY=X	Macro	Dollar/Yen rate
VIX	^VIX	Macro	CBOE Volatility Index

Before proceeding to Granger causality analysis, we first examine the contemporaneous correlation structure of log returns. While correlation does not imply causation and cannot capture lead–lag dynamics, it provides a useful baseline for understanding the co-movement landscape. Figure 1 presents the Pearson correlation matrix. Several patterns merit attention. Within the equity cluster, correlations are uniformly positive and often exceed 0.50, reflecting the well-documented phenomenon of global equity market synchronization—driven by common exposures to U.S. monetary policy, global risk appetite, and institutional cross-listing. The cryptocurrency cluster displays even stronger intra-class correlations, often above 0.70, consistent with the hypothesis that crypto assets are primarily driven by aggregate speculative sentiment rather than idiosyncratic fundamentals. Cross-class correlations are notably weaker, particularly between crypto and traditional assets, suggesting that diversification benefits exist—at least in the unconditional sense. Gold shows near-zero or slightly negative correlations with equities, confirming its traditional role as a safe-haven asset. The VIX exhibits negative correlations

with most equity indices, consistent with its construction as an inverse measure of market confidence. Crucially, correlation is symmetric—it tells us that the S&P 500 and DAX move together, but not whether one leads the other. This limitation motivates the directional Granger causality analysis that follows.

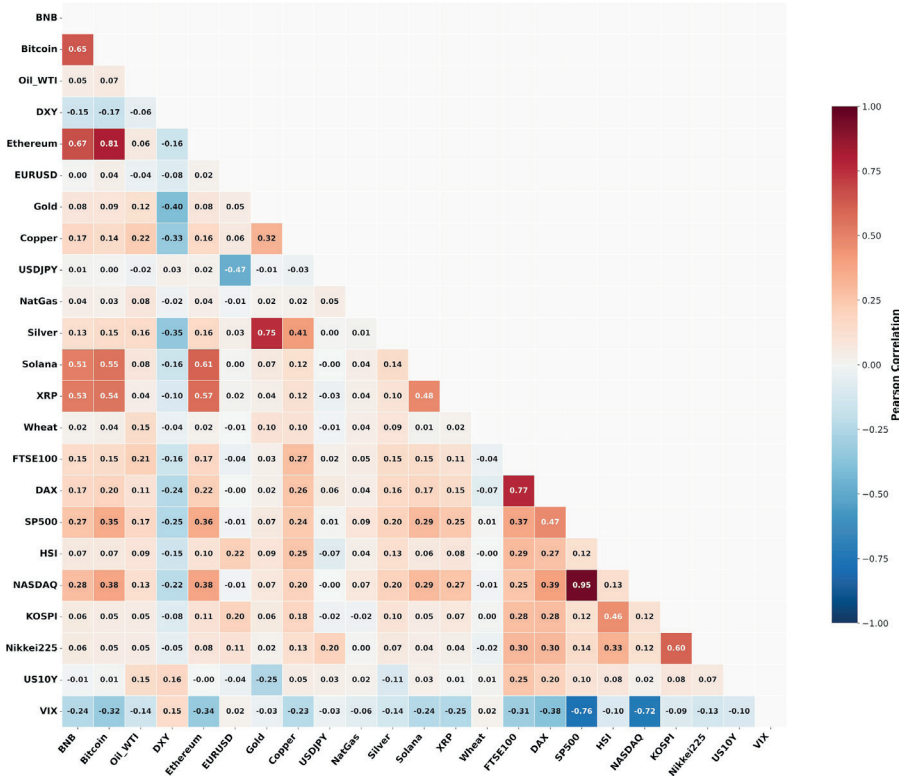

Figure 1. Pearson correlation matrix of daily log returns (23 assets, 2020–2025).

3.2. Stationarity Testing: A Prerequisite for Granger Causality

Before running Granger causality tests, we verify that all return series are stationary. This step matters because Granger inference is built on OLS-based autoregressive regressions; if the variables contain a unit root, the usual F-test theory can fail and the results may become spurious, inflating test statistics and creating false rejections (Granger & Newbold, 1974; Phillips, 1986). In a network setting, that would translate into artificial edges and a distorted topology.

To make the stationarity assessment robust, we use three complementary tests. We estimate the ADF (Dickey & Fuller, 1979) regression

$$\Delta y_t = \mu + \phi y_{t-1} + \sum_{i=1}^{p} \gamma_i \Delta y_{t-i} + \varepsilon_t,$$

and test the unit-root null $H_0 : \phi = 0$ against $H_1 : \phi < 0$. The lag order p is selected using the Akaike Information Criterion (AIC). Rejection of H_0 indicates stationarity.

Following Kwiatkowski et al. (1992), we test the null of level stationarity in the decomposition

$$y_t = \mu + \eta_t + \varepsilon_t, \eta_t = \eta_{t-1} + u_t,$$

where η_t is a random-walk component. The KPSS statistic is based on the partial sum process of residuals from regressing y_t on an intercept,

$$KPSS = \frac{1}{T^2 \hat{\sigma}^2} \sum_{t=1}^{T} S_t^2, S_t = \sum_{j=1}^{t} \hat{\varepsilon}_j,$$

with $\hat{\sigma}^2$ a long-run variance estimator. Failure to reject the null hypothesis supports stationarity and complements the ADF test, particularly in near-unit root settings.

To account for a single endogenous structural break, we apply the Zivot–Andrews unit-root test

(Zivot & Andrews, 1992), which augments the ADF framework by allowing a break date τ to be determined from the data. In the break-in-intercept specification, the regression takes the form

$$\Delta y_t = \mu + \theta DU_t(\tau) + \phi y_{t-1} + \sum_{i=1}^{p} \gamma_i \Delta y_{t-i} + \varepsilon_t,$$

where $DU_t(\tau) = 1(t > \tau)$. The test selects τ to yield the most negative t-statistic for ϕ . This is especially relevant in our sample given major regime shifts (e.g., COVID-19, tightening cycle, and the FTX episode), where breaks can otherwise be mistaken for unit-root behavior.

Table 2 reports the results. All 23 return series reject the unit root null at the 1% level under both the ADF and ZA tests, confirming stationarity. The KPSS test confirms stationarity for 21 of 23 series. 2 assets (Oil, Gold) show KPSS rejection, a well-documented artifact of the KPSS test's sensitivity to heteroskedasticity in commodity return series. Since both ADF and break-accommodating ZA test unanimously confirm stationarity, we conclude that all series are suitable for Granger causality testing.

Table 2. Stationarity Test Results: ADF, KPSS, and Zivot–Andrews

Asset	ADF	p	KPSS	p	ZA	p	Break Date
BNB	-11.48	0.0000	0.442	0.059	-12.17	<0.001	2021-05-03
Bitcoin	-13.66	0.0000	0.248	>0.10	-14.14	<0.001	2021-02-19
Oil_WTI	-10.58	0.0000	1.002†	0.010	-10.99	<0.001	2022-03-08
DXY	-28.76	0.0000	0.142	>0.10	-29.01	<0.001	2022-09-27
Ethereum	-14.14	0.0000	0.433	0.063	-14.70	<0.001	2021-05-11
EURUSD	-38.52	0.0000	0.144	>0.10	-38.72	<0.001	2022-09-28
Gold	-40.77	0.0000	0.639†	0.019	-40.98	<0.001	2022-03-08
Copper	-16.75	0.0000	0.154	>0.10	-17.02	<0.001	2021-05-11
USDJPY	-23.87	0.0000	0.088	>0.10	-24.09	<0.001	2022-10-21
NatGas	-16.80	0.0000	0.105	>0.10	-17.06	<0.001	2022-08-22
Silver	-26.98	0.0000	0.350	0.099	-27.20	<0.001	2021-02-22
Solana	-38.79	0.0000	0.360	0.094	-39.14	<0.001	2021-11-04
XRP	-37.93	0.0000	0.058	>0.10	-38.07	<0.001	2021-04-14
Wheat	-18.85	0.0000	0.109	>0.10	-19.20	<0.001	2022-03-07
FTSE100	-17.08	0.0000	0.044	>0.10	-17.14	<0.001	2022-02-10
DAX	-16.53	0.0000	0.084	>0.10	-16.76	<0.001	2022-09-29
SP500	-12.93	0.0000	0.147	>0.10	-13.17	<0.001	2023-10-27
HSI	-19.09	0.0000	0.156	>0.10	-19.26	<0.001	2022-10-31
NASDAQ	-11.57	0.0000	0.173	>0.10	-11.91	<0.001	2022-12-28
KOSPI	-39.55	0.0000	0.318	>0.10	-39.88	<0.001	2021-04-20
Nikkei225	-23.45	0.0000	0.081	>0.10	-23.61	<0.001	2021-02-22
US10Y	-20.74	0.0000	0.407	0.074	-21.00	<0.001	2022-10-24
VIX	-15.08	0.0000	0.055	>0.10	-15.12	<0.001	2022-10-12

Note: † denotes KPSS rejection at 5% (null = stationarity). All ADF and ZA tests reject H_0 at 1%.

3.3. Granger Causality Testing

Having verified stationarity, we test directional predictability using the standard linear Granger causality framework (Granger, 1969). For each ordered pair (X, Y) with $X \neq Y$, we estimate the unrestricted model

$$Y_t = \alpha_0 + \sum_{i=1}^{L} \alpha_i Y_{t-i} + \sum_{j=1}^{L} \beta_j X_{t-j} + \varepsilon_t ,$$

and compare it with the restricted model that excludes the lagged X . The null hypothesis $H_0 : \beta_1 = \cdots = \beta_L = 0$ is evaluated with an F-test. We set

$L = 5$ in the full-sample analysis and $L = 3$ in the rolling-window analysis. For each pair, we evaluate lags $\ell = 1, \ldots, L$ and select the lag that yields the minimum p-value,

$$p_{\min} = \min_{1 \le \ell \le L} p(\ell).$$

We record the corresponding F-statistic and optimal lag and classify the pair as significant if $p_{\min} < 0.05$. Finally, we stress that "Granger-causality" here refers to predictive precedence, not structural or philosophical causation: X Granger-causes Y only in the sense that lagged X improves forecasts of Y in the linear model (Granger, 1969).

3.4. From Granger Causality Results to a Directed Graph

To move from pairwise predictability results to a system-level view, we translate statistically significant Granger-causality links into a directed network. In this graph representation, each asset is a node, and each significant lead–lag relation is a directed edge, allowing us to study cross-asset information flow using standard network metrics.

Step 1 — Pairwise testing. With $n = 23$ assets, we run $n(n-1) = 23 \times 22 = 506$

Granger-causality F-tests over all ordered pairs. Because the relation is directional, $X \to Y$ does not imply $Y \to X$.

Step 2 — Significance filtering (within-pair lag selection). For each ordered pair, we evaluate lags $L = 1, \ldots, 5$ and obtain lag-specific p-values. We select the lag that yields the minimum p-value, $p_{\min} = \min_{1 \le \ell \le L} p(\ell)$, and record its corresponding F-statistic and selected lag. A directed link is retained if $p_{\min} < 0.05$.

Step 3 — Graph construction (nodes, edges, weights). We build the directed graph $G = (V, E)$, where V is the set of assets. We include a directed edge $(X \to Y) \in E$ if and only if the selected lag for the pair yields $p_{\min} < 0.05$. For visualization, edge strength is encoded as $w(X, Y) = -\log_{10}(p_{\min})$.

Step 4 — Graph analysis. Each edge remains directly traceable to its underlying Granger output (F-statistic, selected lag, and p-value).

3.5. Graph-Theoretic Metrics

We summarize the network topology of the directed Granger graph $G = (V, E)$ using standard graph-theoretic centrality measures. Each metric has a clear interpretation in our setting: who receives information,

who transmits it, which assets act as bridges in the cross-asset system, and how strongly connected the system is overall. In vulnerability terms, these measures also indicate (i) which assets are more exposed to incoming spillovers and (ii) which assets are capable of propagating shocks across the network; hence, higher centrality can be interpreted as greater vulnerability or systemic importance depending on direction.

- **In-degree centrality** (Freeman, 1978):

$$C_{in}(v) = |\{u \in V : (u, v) \in E\}|$$

Counts how many significant links point into node v. In our context, a high C_{in} indicates an asset that is more often predicted by others (information receiver) and, therefore, more vulnerable to upstream shocks transmitted from multiple sources.

- **Out-degree centrality** (Freeman, 1978):

$$C_{out}(v) = |\{u \in V : (v, u) \in E\}|$$

Counts how many significant links originate from node v. A high C_{out} marks an asset that tends to lead others (information transmitter) and thus represents a potential source of systemic

vulnerability, since shocks originating in v can spread to many downstream assets.

- **Betweenness centrality** (Freeman, 1977):

$$C_B(v) = \sum_{s \neq v \neq t} \frac{\sigma_{st}(v)}{\sigma_{st}}$$

Where σ_{st} is the number of shortest paths from s to t, and $\sigma_{st}(v)$ is the number of those

paths that pass through v. Assets with a high betweenness function as bridges that connect otherwise distant parts of the network and are "critical points of vulnerability," because stress can be routed through them to cross market segments.

- **Eigenvector centrality** (Bonacich, 1987):

$$\mathbf{Ax} = \lambda \mathbf{x} \quad (\text{equivalently } x_i = \frac{1}{\lambda} \sum_j a_{ij} x_j)$$

assigns higher scores to nodes that are connected to other highly connected nodes, capturing "influence" beyond simple degree counts and highlighting

vulnerability-amplifying hubs whose connections link them to other influential nodes.

- **Network density** (Wasserman & Faust, 1994):

$$D = \frac{|E|}{n(n-1)}$$

where $n = |V|$. Density ranges from 0 to 1 and summarizes overall connectedness: higher density implies more widespread spillovers and, correspondingly, a more vulnerable system in which contagion can travel through many active pathways.

3.6. Rolling-Window Framework

We re-estimate the full Granger causality network over a rolling window of 250 trading days, advanced in steps of 20 days. At each step, we reconstruct the directed graph and compute density and in-degree centrality. At each window, for every ordered pair, we again select $p_{\min} = \min_{1 \le \ell \le 3} p(\ell)$ and count an edge if $p_{\min} < 0.05$. The resulting density series summarizes time variation in cross-asset predictive linkages.

3.7. Sub-Period Analysis

We partition the sample into four regimes: (i) COVID-19 crash (April 2020 to June 2020),

(ii) recovery and boom (Jul 2020–Dec 2021), (iii) monetary tightening (2022), (iv) post-shock stabilization (2023–2025). Note that our data begins on 2020-04-13, so the COVID-19 crash sub-period contains substantially fewer observations than the other regimes, which has implications for statistical power discussed in Section 4.4.

4. Empirical Results

4.1. Granger Causality Test Results

The full-sample pairwise Granger causality analysis tests 506 directed pairs (23×22). Of these, 181 reject the null hypothesis at the 5% significance level, yielding a network density of 0.358. This density implies that approximately 35.8% of all possible directional information channels are statistically active—a non-trivial level of interconnectedness that underscores the integrated nature of modern financial markets. Of the 181 significant links, 133 (73%) are cross-class links (connecting assets from different classes), while 48 (27%) are intra-class links. The predominance of cross-class transmission channels suggests that

information does not remain confined within asset class boundaries—a finding with direct implications for portfolio construction and risk management.

Table 3 reports the 25 strongest causal links ranked by F-statistic. The highest F-statistic belongs to DXY → EURUSD (F = 5362.47, lag = 1). However, this link reflects a near-mechanical relationship rather than genuine information transmission: the DXY index is a trade-weighted basket in which the euro carries approximately 57.6% weight, so DXY and EUR/USD are linked by construction rather than by market-driven information flow. The same applies to DXY → USD/JPY, since the yen constitutes roughly 13.6% of the DXY basket. We retain these links in the table for completeness but exclude them from substantive interpretation.

Table 3. Granger Causality Test Results: Top 25 Strongest Significant Links

Cause	Effect	Lag	F-Statistic	p-value	Decision
DXY	EURUSD	1	5362.465	<0.001	Reject H_0 ‡
DXY	USDJPY	2	365.252	<0.001	Reject H_0 ‡
SP500	Nikkei225	2	228.361	<0.001	Reject H_0
NASDAQ	Nikkei225	2	215.642	<0.001	Reject H_0
Gold	USDJPY	1	208.277	<0.001	Reject H_0
Gold	EURUSD	1	186.876	<0.001	Reject H_0
VIX	Nikkei225	2	170.284	<0.001	Reject H_0
SP500	KOSPI	2	160.168	<0.001	Reject H_0
Silver	EURUSD	1	148.530	<0.001	Reject H_0
NASDAQ	KOSPI	2	137.902	<0.001	Reject H_0
Copper	EURUSD	1	121.871	<0.001	Reject H_0
VIX	KOSPI	2	119.415	<0.001	Reject H_0
VIX	DAX	1	91.292	<0.001	Reject H_0
DAX	Nikkei225	2	87.083	<0.001	Reject H_0
NASDAQ	DAX	1	83.998	<0.001	Reject H_0
Silver	USDJPY	1	74.402	<0.001	Reject H_0
DAX	KOSPI	2	72.631	<0.001	Reject H_0
DXY	HSI	1	68.798	<0.001	Reject H_0
DXY	KOSPI	1	66.649	<0.001	Reject H_0
Copper	KOSPI	1	65.520	<0.001	Reject H_0
SP500	HSI	2	62.873	<0.001	Reject H_0
VIX	HSI	2	55.610	<0.001	Reject H_0
NASDAQ	HSI	2	53.075	<0.001	Reject H_0
FTSE100	KOSPI	2	52.961	<0.001	Reject H_0
SP500	EURUSD	2	51.854	<0.001	Reject H_0

Note: Ranked by F-statistic (descending). All links are significant at $p < 0.05$. Total significant links: 181/506. ‡ denotes near-mechanical links (DXY basket components) excluded from substantive interpretation.

The strongest genuine information flow link is SP500 → Nikkei225 ($F = 228.36$, $p < 0.001$, lag = 2), indicating that past SP500 returns contain significant predictive information for Nikkei225 returns. This is consistent with the market microstructure literature on cross-border information transmission. As the world's most liquid equity benchmark, SP500's closing prices set the information tone that Nikkei225 incorporates at its next opening, with the full effect materializing over 2 trading days. Among the top 25 links, lag-2 relationships are the most frequent (56%), suggesting that the dominant information transmission horizon in our cross-asset universe spans 2 trading days—suggesting a two-step propagation mechanism rather than instantaneous overnight diffusion. This is consistent with a sequential transmission chain—for example, U.S. market movements first affect European markets the next day, which in turn propagate to Asian markets on the second day—or with the time required for market participants to fully digest and reprice cross-asset signals.

4.2. Full-Sample Network Architecture

Figure 2 visualizes the directed graph constructed from the 181 significant Granger causality links. Each node represents an asset, scaled by total degree (the sum of incoming and outgoing edges), and each arrow represents a statistically significant causal relationship—directly traceable to an F-test with a specific F-statistic, p-value, and optimal lag. The resulting network exhibits a density of 0.358 (181 significant links out of 506 possible directed pairs), indicating that roughly 35.8% of all potential directional information channels are statistically active in the full sample. The visual topology immediately reveals several economically meaningful patterns. The network is not uniformly connected: a small number of hub nodes attract and emit a disproportionate share of edges, while peripheral nodes participate in fewer transmission channels. This hub-and-spoke structure mirrors the hierarchical organization of global capital markets, where the most connected assets—EURUSD, HSI, and SP500 (Table 4)—function as information clearinghouses through which shocks propagate to less connected markets. From a vulnerability perspective, these hub nodes can be viewed as potential systemic points of fragility: shocks that originate in (or are transmitted into) highly connected assets have more opportunities to diffuse across the network, increasing the likelihood of broad spillovers.

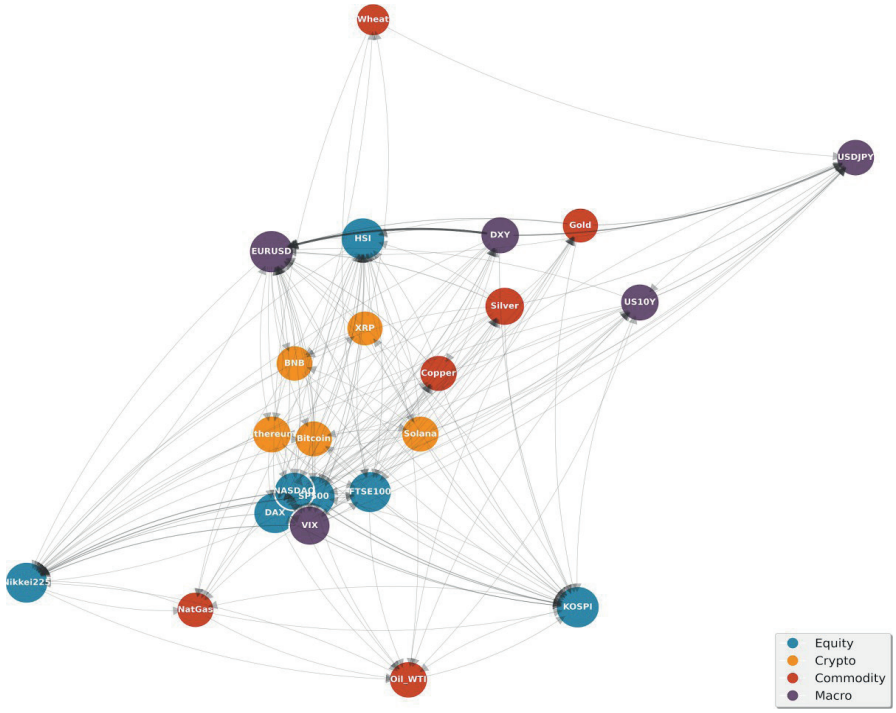

Figure 2. Full-sample Granger causality directed graph (2020–2025, 23 assets, 181 edges). Node size: total degree. Edge width and opacity: $-\log_{10}$(p-value). Arrow direction: cause → effect.

The clustering of same-class nodes in the graph is also informative. Equity indices tend to cluster together, as do cryptocurrencies, reflecting the stronger intra-class information transmission documented in the correlation analysis. However, the cross-class edges—visible as arrows bridging different-colored clusters—are the most economically significant finding, as they reveal the specific channels through which, for example, a monetary policy shock (captured by DXY or US10Y) propagates to commodity prices, or a crypto market crash spills over into equity sentiment. These cross-class connections are also central to vulnerability analysis: bridge-type assets (high betweenness) can transmit stress between otherwise more weakly connected market segments, so disruptions affecting such nodes may disproportionately elevate contagion risk. Moreover, the relatively high density implies that spillovers can travel through multiple alternative pathways rather than relying on a single channel, which further increases system-wide vulnerability during stress episodes.

Table 4 reports the centrality metrics that quantify each asset's role in the information network. From a vulnerability perspective, these metrics

identify (i) exposure to incoming spillovers (high C_{in}), (ii) systemic shock-propagation potential (high C_{out}), and (iii) cross-segment fragility points that relay stress across clusters (high C_B). EURUSD records the highest total degree (25), making it the most interconnected node in the system. A caveat is warranted: EURUSD's high connectivity partly reflects its near-mechanical link with DXY, which inflates both its in-degree and out-degree. Nevertheless, even excluding the DXY link, EURUSD remains among the most connected nodes, consistent with the euro's role as the world's second most traded currency and a key conduit for monetary policy transmission between the U.S. and the Eurozone. HSI follows with a total degree of 24. The distinction between in-degree and out-degree reveals the directionality of information flow. SP500 has the highest out-degree (14), identifying it as the system's primary information transmitter—its past returns Granger-cause the most other assets. In vulnerability terms, SP500 is therefore a key systemic "shock emitter": disturbances originating in SP500 have the highest potential to diffuse widely through downstream links. This is economically intuitive: SP500 is among the most liquid and widely watched instruments, and its price movements are rapidly incorporated by market participants trading related assets. Conversely, KOSPI has the highest in-degree (19), marking it as the most information-receptive asset—a "follower" that reacts to signals from many other markets. This makes KOSPI a high-exposure node: it is relatively more vulnerable to shocks transmitted from elsewhere in the network because many assets contain predictive information about its returns. HSI's betweenness centrality of 0.138 confirms its role as a critical bridge node: removing it from the network would disrupt more shortest paths than removing any other asset, suggesting it serves as a relay between otherwise weakly connected market segments. Bridge nodes like HSI are particularly important for vulnerability assessment because they can transmit stress across otherwise partially separated market communities, turning localized shocks into cross-class contagion.

At the asset-class level, the average out-degree reveals a clear transmission hierarchy: equity (9.6), crypto (8.8), macro (7.2), commodity (5.7). Equity indices lead the ranking, which is consistent with the depth, liquidity, and institutional following of major stock markets—price discovery in these benchmark indices occurs rapidly and diffuses outward to other asset classes. Notably, however, the cryptocurrency class ranks a close second, with an average out-degree nearly matching equities. This is a striking departure from earlier findings: Corbet et al. (2018) documented that crypto markets were largely isolated from traditional finance. Our results suggest that by the 2020–2025 period, digital assets have become deeply embedded in the global

information network, both receiving and broadcasting signals to traditional asset classes. Commodities, by contrast, occupy the lowest transmission tier, consistent with their role as receivers of macro and equity-driven demand signals rather than originators of cross-asset information. Taken together, this hierarchy also implies a vulnerability ordering in shock propagation: equity and crypto markets are the primary sources of system-wide spillover risk, whereas commodities are more often downstream recipients of stress transmitted from other classes.

Table 4. Directed Graph Centrality Metrics

Asset	C_{in}	C_{out}	Total	C_B	C_E	Class
EURUSD	18	7	25	0.1162	0.4013	Macro
HSI	16	8	24	0.1380	0.2916	Equity
SP500	9	14	23	0.1076	0.2707	Equity
KOSPI	19	4	23	0.0407	0.4075	Equity
FTSE100	9	13	22	0.0672	0.1484	Equity
Nikkei225	18	4	22	0.0342	0.3924	Equity
DAX	11	11	22	0.0706	0.1699	Equity
VIX	7	11	18	0.0293	0.2011	Macro
Silver	5	11	16	0.0163	0.1177	Commodity
NASDAQ	3	13	16	0.0179	0.0708	Equity
US10Y	7	7	14	0.0280	0.1544	Macro
Ethereum	3	11	14	0.0179	0.1080	Crypto
DXY	5	9	14	0.0100	0.1177	Macro
USDJPY	11	2	13	0.0141	0.2255	Macro
Oil_WTI	9	4	13	0.0141	0.2654	Commodity
Solana	2	10	12	0.0289	0.0509	Crypto
Copper	6	6	12	0.0011	0.1268	Commodity
Bitcoin	4	8	12	0.0093	0.0969	Crypto
BNB	4	7	11	0.0365	0.0804	Crypto
NatGas	6	5	11	0.0439	0.1787	Commodity
Gold	5	5	10	0.0015	0.1106	Commodity
XRP	2	8	10	0.0106	0.0180	Crypto
Wheat	2	3	5	0.0099	0.0269	Commodity

Note: The table is sorted in descending order by total degree $(C_{in}+C_{out})$.

Figure 3 decomposes the centrality structure into three bar charts. The in-degree panel reveals which assets are most responsive to external information—these are the assets whose returns are best predicted by the lagged returns of others. From a market efficiency perspective, a high in-degree suggests that an asset's price discovery process is partially driven by signals originating

elsewhere, which may reflect slower information incorporation or structural dependence on upstream markets. At the same time, high in-degree implies greater exposure to incoming spillovers because the asset is influenced by a broader set of upstream signals (e.g., KOSPI and Nikkei225 rank among the highest in-degree nodes).

The out-degree panel identifies the system's information leaders: assets whose past movements carry predictive power for many others. For portfolio managers, high out-degree assets serve as early-warning indicators—monitoring their price action can provide advance notice of movements in downstream assets. High out-degree also signals greater systemic impact: disturbances originating in these nodes have more downstream pathways through which they can spread (notably SP500 and NASDAQ in Figure 3).

The betweenness centrality panel identifies bridge assets that connect otherwise disparate market segments. These bridge nodes are critical from a systemic risk perspective: a shock to a high-betweenness asset can rapidly propagate across multiple asset classes by traversing these bridging pathways. Accordingly, high-betweenness nodes act as relay points between clusters; disruptions affecting them can transmit stress across market segments rather than remaining confined within a single class (e.g., HSI and EURUSD).

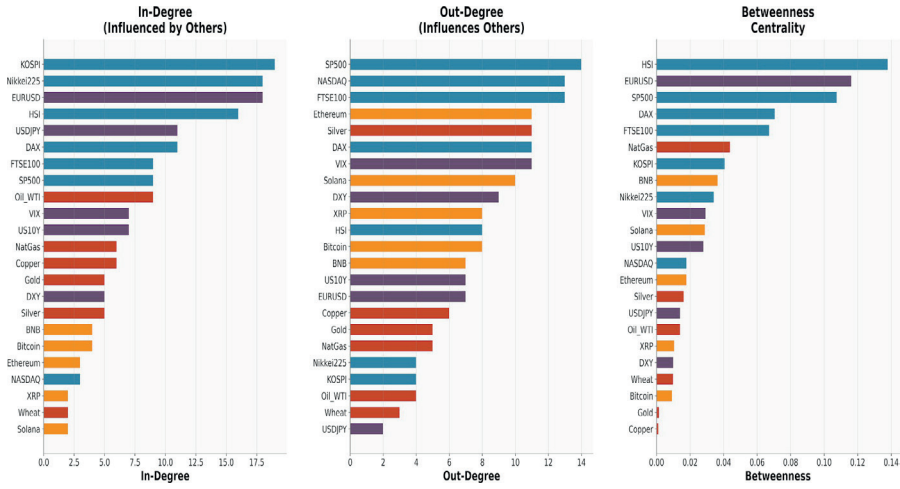

Figure 3. Centrality bar charts: in-degree (information receivers), out-degree (information transmitters), and betweenness centrality (bridge nodes).

Figure 4 presents the Granger causality heatmap, where the color intensity of cell (i, j) represents the statistical strength of the causal link from asset j (column) to asset i (row), measured as $-log_1(p-value)$. This visualization

provides a comprehensive view of the full 23×23 causality matrix (506 ordered pairs), with non-significant links appearing near zero intensity.

The bright hot spots along certain rows and columns correspond to the hub assets identified in the centrality analysis—their rows (as receivers) and columns (as transmitters) show consistently high values. The within-class clustering visible for equities and cryptocurrencies is consistent with stronger intra-class predictability, while the off-diagonal bright cells reveal the specific cross-class channels. The asymmetric nature of the heatmap—where cell (i, j) can be bright while cell (j, i) is dark—is the clearest visual demonstration that information flow is directional, justifying the use of Granger causality over symmetric correlation analysis.

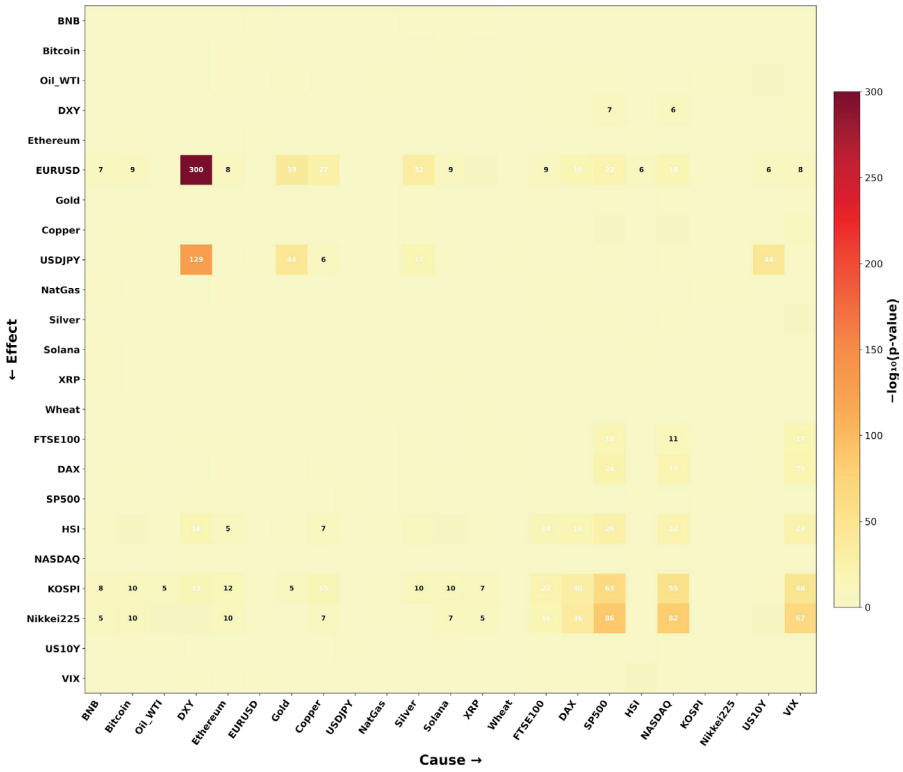

Figure 4. Granger causality significance heatmap. Cell (i, j) = $-log_{10}$(p-value) for the test of whether asset j (column) Granger-causes asset i (row). Brighter = stronger evidence.

4.3. Rolling-Window Dynamics

The full-sample network provides a time-averaged picture, but financial markets are inherently non-stationary in their interconnectedness. The 2008 crisis, the COVID-19 pandemic, and the 2022 monetary tightening cycle all demonstrated that cross-asset correlations and spillovers surge during stress episodes—a phenomenon known as contagion. To capture this time-varying structure, we re-estimate the Granger causality network over a rolling window of 250 trading days, generating a time series of network density that reveals how the intensity of cross-asset information flow evolves across market regimes.

Figure 5 plots this rolling network density alongside the number of significant links. The density exhibits pronounced regime dependence, ranging from 0.144 during tranquil periods to 0.374 during stress episodes ($mean = 0.241, \sigma = 0.055$). The economic interpretation is clear: during calm markets, assets are driven primarily by their own idiosyncratic factors, and cross-asset predictability is limited. When uncertainty rises—whether from a pandemic, a monetary policy shock, or a major institutional failure—market participants engage in broad-based risk repricing, creating synchronized movements that manifest as new Granger causality links. The increases visible around the COVID-19 aftermath, the Fed rate hike cycle, and the FTX episode are consistent with this mechanism; however, the most pronounced spike occurs toward the end of the sample, when both density and the number of significant links jump sharply, indicating an abrupt tightening of cross-asset linkages.

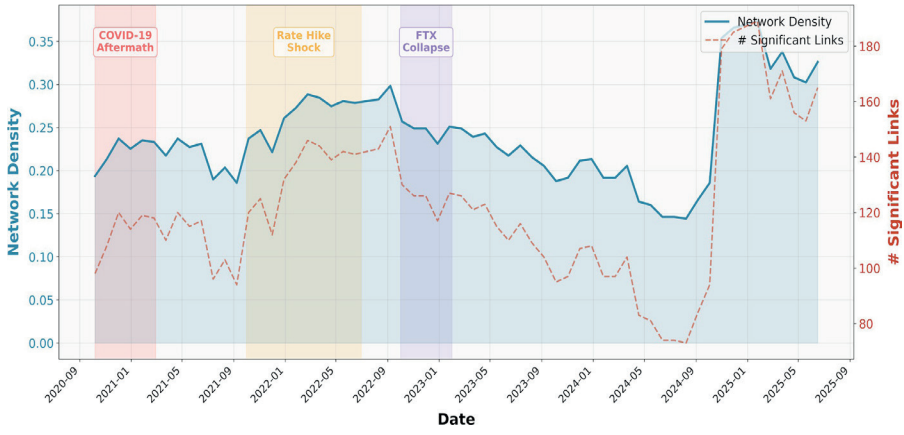

Figure 5. Rolling network density (250-day window, step = 20 days). Shaded regions: crisis periods.

Left axis: density. Right axis: number of significant Granger links.

While the aggregate density captures the system-wide intensity of spillovers, it does not reveal which specific assets are gaining or losing centrality over time. Figure 6 tracks the rolling in-degree centrality of eight key assets across the four asset classes. In-degree centrality at each window measures how many other assets Granger-cause a given asset—a high in-degree at a particular time indicates that the asset is temporarily receiving information from many sources, suggesting it is being swept up in broad market dynamics rather than moving on its own fundamentals. In practical terms, periods of elevated in-degree mark heightened exposure to incoming spillovers, since the asset is being influenced by a wider set of upstream signals. Several patterns are noteworthy. Equity in-degrees tend to spike together during stress episodes, reflecting the well-documented contagion channel among global stock markets. Bitcoin's in-degree shows pronounced time variation—rising sharply during certain turbulent windows when the crypto market becomes more tightly linked to the broader financial system. Gold's in-degree also exhibits episodic increases, consistent with its tendency to respond more strongly to system-wide repricing episodes.

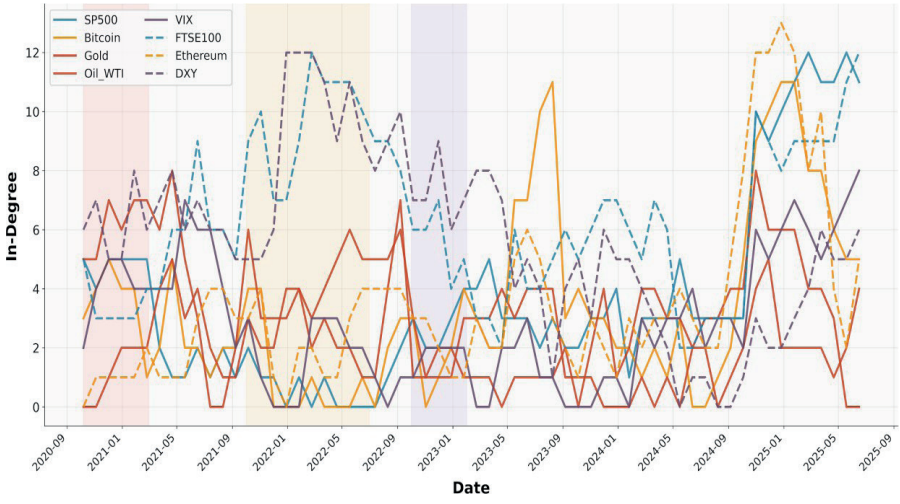

Figure 6. Rolling in-degree centrality (250-day window) for selected assets. Higher values indicate the asset is receiving more predictive signals from the rest of the network.

4.4. Sub-Period Network Comparison

To complement the rolling-window analysis with a more structured comparison, we partition the sample into four economically distinct regimes and estimate a separate Granger causality network for each. This approach sacrifices the continuity of the rolling framework but gains the ability to

compare network topologies side-by-side, revealing how the architecture of information flow—not just its intensity—changes across market environments.

Figure 7 displays the four sub-period networks. The visual contrast is striking. The Rate Hike Era (2022) network is the densest (density = 0.283), with arrows connecting many pairs of asset classes—consistent with a regime in which cross-asset linkages intensify and information spillovers become more widespread across markets. During such episodes, the usual fundamental drivers of individual asset classes can be overshadowed by common risk factors—aggregate risk aversion and global repricing—creating temporary but statistically significant Granger causality links even between economically unrelated assets. A denser sub-period network also implies higher system-wide vulnerability, since shocks have more active pathways through which they can propagate across asset classes.

The COVID-19 crash (April 2020 to June 2020) network is the sparsest (density = 0.134). An important methodological caveat applies here: this sub-period is the shortest window in our partition, containing substantially fewer observations than the other regimes. Since the Granger F-test requires sufficient data to detect lead–lag relationships, the lower density in this sub-period partly reflects reduced statistical power rather than genuinely weaker cross-asset linkages. Indeed, the rolling-window analysis in Section 4.3—which uses a consistent 250-day estimation window throughout—shows elevated density during the COVID aftermath, suggesting that the pandemic did intensify cross-asset spillovers as one would expect from a systemic shock. The sub-period results for the COVID window should therefore be interpreted with caution.

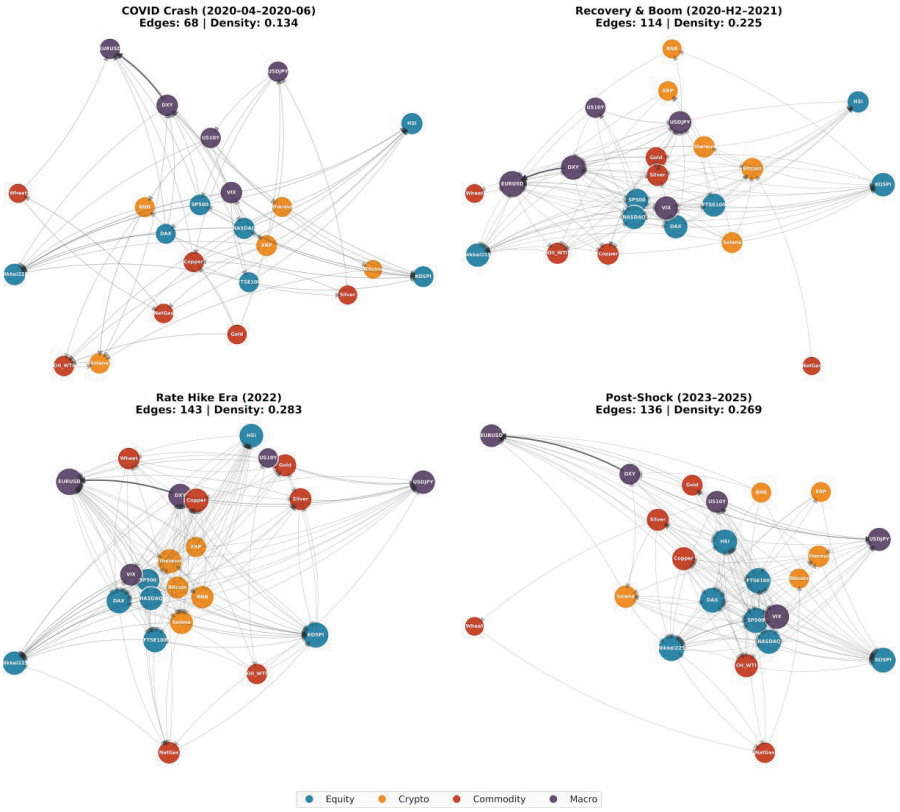

Figure 7. Sub-period Granger causality directed graphs (23 assets). Each panel shows the network for a distinct market regime. Node size: total degree. Arrow direction: cause → effect.

Figure 8 provides a quantitative comparison across regimes. The left panel confirms the density pattern: the network is sparsest during the COVID-19 crash sub-period (0.134) and becomes substantially denser in the subsequent regimes, peaking in the Rate Hike Era (2022) at 0.283 and remaining elevated in the Post-Shock period (0.269). The right panel decomposes information transmission by asset class, revealing which classes drive these regime differences. The increase in connectedness is not uniform across classes. In the Rate Hike Era, cryptocurrencies exhibit the highest average out-degree, indicating that crypto markets play an unusually strong transmitter role in that regime, while equities remain influential across regimes and become the dominant transmitter class in the Post-Shock sample. Macro variables and commodities contribute more moderately and show less consistent dominance across regimes. Overall, the shifting class-level out-degree profile suggests time-varying cross-market integration, with implications for diversification:

when transmitter roles concentrate in equities or crypto and the network is denser, spillover channels multiply and system-wide vulnerability to contagion increases, raising the likelihood that shocks propagate across asset classes.

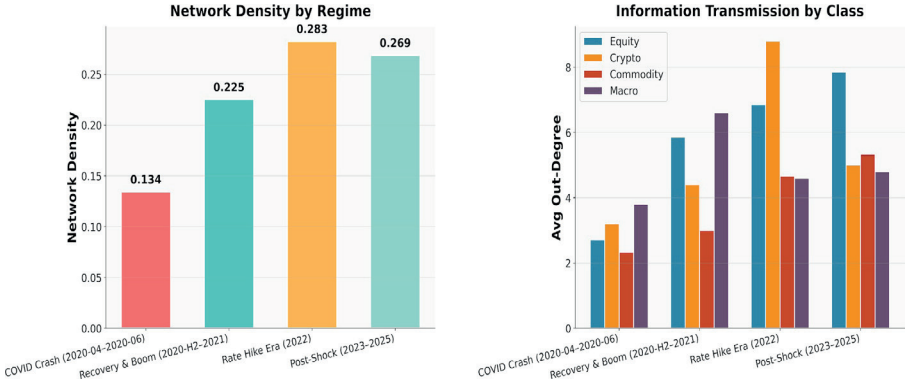

Figure 8. Left: Network density by regime. Right: Average out-degree by asset class across regimes. A higher out-degree indicates greater information transmission from the asset class.

5. Discussion and Implications

5.1. Implications for Portfolio Management

The regime-dependent nature of the network has direct and actionable implications for dynamic portfolio construction. During tranquil periods, the sparser network (density as low as 0.144) indicates that cross-asset predictability is limited, and assets are driven primarily by their own asset-specific (idiosyncratic) dynamics. In this regime, cross-class diversification is effective: a portfolio combining equities, crypto, commodities, and macro instruments benefits from low inter-class correlations and minimal spillover risk. However, during stress episodes, the sharp increase in density to 0.374 implies that diversification benefits erode precisely when they are most needed—a phenomenon known as "diversification meltdown" in the portfolio theory literature. The practical implication is that static asset allocation strategies that assume constant cross-asset relationships will systematically underestimate tail risk.

The centrality metrics suggest a monitoring strategy. Among the highest-degree nodes, SP500 (highest out-degree, 14) and HSI (highest betweenness, 0.138) are the most informative for anticipating broad market shifts, as their connectivity reflects genuine information transmission rather than

compositional artifacts (recall that EURUSD's high total degree partly reflects the DXY mechanical link, as discussed in Section 4.2). A portfolio manager who observes an unusual movement in a high-out-degree asset can anticipate—with a lag of 1 to 5 trading days—that downstream assets will respond. Similarly, the rolling density can be used as a dynamic allocation signal: when density exceeds its historical average, the manager can reduce cross-asset positions or increase hedging, anticipating that contagion risk is elevated.

5.2. Implications for Systemic Risk Monitoring

For regulators and central bankers, the rolling network density series can serve as a timely, operational barometer of systemic interconnectedness. The mean density of 0.241 with a standard deviation of 0.055 provides a natural threshold: when density exceeds 0.296 (one standard deviation above the mean), it signals elevated connectivity and contagion risk. Unlike market-based systemic risk measures such as CoVaR or SRISK, which require balance sheet data or high-frequency market microstructure data, the Granger causality density can be computed from publicly available daily price data alone, making it accessible to a broader range of regulators and market participants.

The identification of bridge nodes (high betweenness centrality) also has regulatory implications. These assets serve as conduits through which shocks propagate between otherwise weakly connected market segments. Monitoring bridge assets—and particularly tracking changes in their betweenness centrality over time—can help regulators identify emerging contagion pathways before they fully materialize.

5.3. Limitations and Future Research

Several limitations warrant acknowledgment and suggest directions for future research. First, the linear Granger causality test may miss nonlinear dependencies that are particularly relevant during market stress (Diks & Panchenko, 2006). Extending the framework to quantile Granger causality (Troster, 2018) would allow the detection of asymmetric tail-driven spillovers. Second, the rolling-window approach involves a fundamental trade-off between window length and regime detection precision: shorter windows improve temporal resolution but increase estimation noise. Regime-switching models could provide a more principled alternative. Third, intraday data would enable within-day transmission analysis, capturing the rapid information diffusion that occurs during market hours. Fourth, the inclusion of additional asset classes—such as real estate investment trusts (REITs), sovereign credit default swaps, or carbon emission allowances—could reveal additional transmission

channels. Fifth, the current framework treats all significant links equally at the binary level; future work could employ weighted network analysis to incorporate the strength of causality into the centrality computations. Sixth, the significance tests are conducted at a nominal 5% level without family-wise or false-discovery-rate correction (Benjamini & Hochberg, 1995). With 506 simultaneous tests, approximately 25 links may represent false discoveries; applying FDR control would yield a more conservative but potentially more robust network.

6. Conclusion

This chapter has demonstrated that graph-based Granger causality network analysis provides a powerful and transparent lens for understanding cross-asset information flow. The methodology proceeds through a rigorous pipeline: (i) stationarity verification through a triangulation strategy (ADF

+ KPSS + Zivot–Andrews) to ensure the validity of the Granger F-test, (ii) pairwise Granger causality testing across all 506 directed asset pairs, (iii) transformation of significant test results into a directed graph where each edge is directly traceable to a specific F-statistic, p-value, and optimal lag, and (iv) graph-theoretic characterization through degree, betweenness, and eigenvector centrality metrics.

Three principal findings emerge from the empirical analysis. First, EURUSD and HSI occupy the positions of highest total centrality (degrees of 25 and 24, respectively), confirming their roles as the primary hubs through which information flows across the 23-asset universe. This finding has direct implications for portfolio monitoring: monitoring highly central assets—especially strong transmitters (high out-degree) and bridge nodes (high betweenness)—provides early signals for movements in downstream assets. Second, network density is strongly regime-dependent, ranging from 0.144 during tranquil periods to 0.374 during market stress, with the rolling density series offering a real-time systemic risk indicator. The economic implication is that diversification benefits are time-varying and tend to erode precisely during the stress episodes when they are most needed. Third, the expanded 23-asset universe spanning four classes reveals cross-class transmission pathways—from monetary policy variables to commodities, from crypto to equities, and from equity benchmarks to emerging market indices—that smaller, single-class studies cannot detect in a time-varying manner across regimes.

The methodological contribution is equally important. By formalizing the four-step Granger-to-graph pipeline, this chapter provides a reproducible framework that can be applied to any set of time series—financial or

otherwise—where directional predictive relationships are of interest. The code accompanying this chapter is fully automated: a single execution produces the complete analysis from data download through Word document generation, ensuring that every number in this chapter is directly traceable to the underlying computations.

References

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological), 57*(1), 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

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

Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology, 92*(5), 1170–1182. https://doi.org/10.1086/228631

Bouri, E., Cepni, O., Gabauer, D., & Gupta, R. (2021). Return connectedness across asset classes around COVID-19. *International Review of Financial Analysis, 73*, 101646. https://doi.org/10.1016/j.irfa.2020.101646

Breitung, J., & Candelon, B. (2006). Testing for short- and long-run causality. *Journal of Econometrics, 132*(2), 363–378. https://doi.org/10.1016/j.jeconom.2005.02.004

Corbet, S., Meegan, A., Larkin, C., Lucey, B., & Yarovaya, L. (2018). Exploring dynamic relationships between cryptocurrencies and other financial assets. *Economics Letters, 165*, 28–34. https://doi.org/10.1016/j.econlet.2018.01.004

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association, 74*(366a), 427–431. https://doi.org/10.1080/01621459.1979.10482531

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

Diks, C., & Panchenko, V. (2006). A new statistic for nonparametric Granger causality testing. *Journal of Economic Dynamics and Control, 30*(9–10), 1647–1669. https://doi.org/10.1016/j.jedc.2005.08.008

Eun, C. S., & Shim, S. (1989). International transmission of stock market movements. *Journal of Financial and Quantitative Analysis, 24*(2), 241–256. https://doi.org/10.2307/2330774

Freeman, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry, 40*(1), 35–41. https://doi.org/10.2307/3033543

Freeman, L. C. (1978). Centrality in social networks: Conceptual clarification. *Social Networks, 1*(3), 215–239. https://doi.org/10.1016/0378-8733(78)90021-7

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica, 37*(3), 424–438. https://doi.org/10.2307/1912791

Granger, C. W. J., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics, 2*(2), 111–120. https://doi.org/10.1016/0304-4076(74)90034-7

Hamao, Y., Masulis, R. W., & Ng, V. (1990). Correlations in price changes and volatility across international stock markets. *Review of Financial Studies, 3*(2), 281–307. https://doi.org/10.1093/rfs/3.2.281

Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity. *Journal of Econometrics, 54*(1–3), 159–178. https://doi.org/10.1016/0304-4076(92)90104-Y

Phillips, P. C. B. (1986). Understanding spurious regressions in econometrics. *Journal of Econometrics, 33*(3), 311–340. https://doi.org/10.1016/0304-4076(86)90001-1

Tang, K., & Xiong, W. (2012). Index investment and the financialization of commodities. *Financial Analysts Journal, 68*(6), 54–74. https://doi.org/10.2469/faj.v68.n6.5

Troster, V. (2018). Testing for Granger-causality in quantiles. *Econometric Reviews, 37*(8), 850–866. https://doi.org/10.1080/07474938.2016.1172400

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge University Press. https://doi.org/10.1017/CBO9780511815478

Zivot, E., & Andrews, D. W. K. (1992). Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *Journal of Business & Economic Statistics, 10*(3), 251–270. https://doi.org/10.1080/07350015.1992.10509904