Crisis-Origin Dependence in Sector Volatility Spillovers: Evidence from U.S. Equity Markets, 2015–2025

March 13, 2026

Abstract

This paper examines whether the origin of a financial crisis determines how volatility shocks transmit through equity sectors. We apply a quantile vector autoregression (QVAR) framework to daily changes in option-implied volatility for nine U.S. sector ETFs from 2015 to 2025. This timeframe covers six crises with fundamentally different origins: the 2015–2016 oil/China shock, the 2018 Volmageddon tightening, the COVID-19 pandemic, the 2022–2023 high-inflation monetary tightening, the 2023 SVB banking crisis, and the 2025 tariff crisis. Panel regressions reject equal transmission patterns across all nine quantiles, and all six crises yield unique top-3 transmitter compositions. Monetary tightening elevates industrials and cyclical sectors; COVID-19 features the distinctive emergence of energy-sector transmission; the tariff crisis produces the unprecedented emergence of financials as the leading transmitter; and the oil/China shock features consumer discretionary and technology leadership—with no sector appearing consistently in the top-3 across all crises. We also document a consistent cyclical-defensive divide, where cyclical sectors transmit—and defensive sectors absorb—volatility. The divide holds across all regimes. Utilities qualify as the only pure safe haven, serving as a net receiver in every regime. Our findings demonstrate that crisis-origin shapes the volatility transmission network, with implications for scenario-based stress testing, portfolio construction, and crisis-specific risk management.

Keywords: Quantile VAR, Volatility Spillovers, Sector ETFs, Crisis Transmission, Financial Contagion, Implied Volatility

JEL Classification: C32, C58, G01, G11

1 Introduction

Financial crises are typically characterized by heightened uncertainty, increased cross-asset interdependence, and loss of diversification benefits. This description is helpful in broad terms, but it may hide a more basic empirical problem. For example, a shock to banking undermines balance sheets and funding conditions, a pandemic changes production, mobility, and demand, an oil-price collapse propagates through commodity exposure and sectoral cash flow sensitivity, monetary tightening reprices discount rates and financing costs, and trade policy shocks redistribute expected profitability

across industries that are globally exposed. With fundamental differences in crisis origins, there is little basis for presuming that the volatility network of transmission is stable across crises.

This paper examines whether the initial cause of a crisis is a systematic factor that reshapes the sectoral architecture of volatility spillovers in the U.S. equity markets. More precisely, we examine whether the identities of net volatility transmitters and receivers are dependent on the initiating event. This issue has implications for the theory of asset pricing and the practice of financial stability. For investors, it provides information on hedging, sector rotation, and the dependability of diversification strategies under stress. For regulators and macroprudential authorities, it means that scenario design may depend on the nature of the initiating shock, rather than the nature of a crisis template. If spillover leadership varies by crisis types, models calibrated by a short historical episode might not work well in another. A framework based on financial-sector dominance may be helpful in the case of a banking event, but much less helpful in the case of a pandemic shock or a trade policy conflict when the impulse emanates elsewhere and spreads through different channels.

Existing literature on connectedness has documented that spillovers are time-varying and tend to increase during turbulent times. Network-based measures of connectedness have been used to show that financial linkages are dynamic and state dependent and are frequently amplified during periods of market stress (Diebold and Yilmaz, 2009, 2012, 2014). However, the most relevant question regarding crisis heterogeneity is not fully answered. The variation in connectedness over time alone does not demonstrate the dependence on the origin of the crisis in the transmission architecture. Much of the empirical literature compares a small number of episodes, observes changes in sector rankings, and interprets them as a sign of the specificity of the crisis. While such comparisons are suggestive, they cannot cleanly parse out systematic differences across crisis origins from sampling noise, window-induced smoothing, or contamination across adjacent regimes. This limitation is especially relevant in the case of short-lived crises using long rolling windows, which tend to blur regime boundaries.

We address this limitation by pooling a wider crisis design with formal cross-crisis inference. Our dataset comprises daily 30-day at-the-money implied volatility of nine U.S. sector ETFs from January 2015 through March 2025. The sample covers six major episodes: the oil/China shock of 2015–2016, the Volmageddon tightening episode of 2018, the COVID-19 crisis of 2020, the inflation-driven monetary tightening cycle of 2022–2023, the SVB banking stress of 2023, and the tariff shock of 2025. The value of this design is that it brings together commodity, policy, pandemic, banking, trade policy disturbances, and puts this in one internally consistent framework. Compared to conventional two-crisis studies, this expanded set of episodes offers greater variation in the source of the stress and, therefore, a better basis for discrimination on whether or not spillover structures are true crisis-specific. We focus on implied rather than realized volatility as they provide a forward-looking assessment of risk. In rapidly evolving crises, option markets tend to take the information into account before lower frequency macro variables and realized volatility fully capture the repricing of uncertainty. Therefore, for a study of the transmission of expected risk across sectors, implied volatility is the appropriate state variable.

Methodologically, we utilize the quantile connectedness framework by Ando et al. (2022) implemented in the form of a Quantile Vector Autoregressive (QVAR). The quantile dimension is crucial because patterns measured at the conditional mean or median may hide the presence of asymmetries that are economically important only in the tails, the very place where the dynamics of a crisis may be most pertinent. A sector that seems peripheral during normal states can be systemically important in bad states. Accordingly, we estimate connectedness at nine quintiles (deciles). We use the median as the main reference while leveraging the entire quantile structure to examine how crisis dependence changes with the market situation. Our empirical strategy is to use both rolling window estimation of bidirectional spillovers and formal tests of transmission patterns across crisis types. Specifically, we examine the directional TO, FROM and NET connectedness and the system-wide Total Connectedness Index and test whether there are systematic differences in the cyclical-defensive transmission gap between crises. To ensure that inference is not sensitive to one single measure, we complement the evidence from panel (i.e., the single measure) with rank correlation tests across crises, bootstrap interval comparisons, and analysis on the composition of the top transmitting sectors. Because rolling windows necessarily overlap, we also quantify regime purity to demonstrate how much each window is informed by the observations inside the assigned crisis vs. the observations drawn from the surrounding periods.

The paper has three major contributions. First, it shows formally that the sectoral volatility spillovers in U.S. equity markets depend on the origin of the crisis. Although previous research has shown that connectedness is time-varying and increases during times of stress, that finding does not demonstrate that transmission structures are crisis-specific in any systematic way. Our contribution is to identify and test that distinction directly. Rather than inferring the crisis-specificity from descriptive contrasts from a small number of episodes, we put the claim on an explicit inferential footing and assess whether volatility transmission patterns differ between crises that are caused by different economic shocks. The contribution thus counts as conceptual as much as empirical, as it changes the focus from the general crisis amplification to the origin dependence of the spillover architecture itself.

Second, the paper distinguishes between episodic rotation and persistent structure in sectoral transmission networks. The evidence indicates that spillover leadership is not stable across crises, as the sectors that dominate transmission vary with the source of stress. Nevertheless, the results further show that such rotation takes place within an underlying durable structural ordering, where cyclical sectors always behave as net transmitters and defensive sectors as net receivers. This is an important distinction, since it explains that heterogeneity of crisis does not imply lack of regularity; on the other hand, ongoing regularity does not imply crisis invariance. Consequently, the manuscript provides a more nuanced characterization of connectedness under stress: Volatility transmission is neither solely linked to specific episodes nor to fixed patterns, but is the result of interplay between a persistent sectoral backbone and crisis-contingent reordering.

Third, the manuscript presents regime-purity diagnostics as an explicit layer of inference-quantity in rolling window connectedness analysis. This addition addresses an understudied though central

challenge of crisis-based network research. By analyzing short-lived episodes with large (long and overlapping) estimation windows, the identified spillover pattern attributed to a crisis may reflect, in part, regimes that are adjacent to one another, rather than the crisis. Absent some kind of direct measurement of such a contamination, comparisons between crisis risks can conflate the economic structure with window mechanics. By quantifying the fraction of each rolling estimate that is borne out by the within-regime observations, the manuscript increases the credibility of cross-crisis inference and provides a replicable diagnostic to be used more widely in studies of dynamic spillovers. In this respect, the contribution is of a methodical nature, but it also has material implications, since the reliability of any claim to crisis-specificity depends on the quality of the purity of the regime from which it is deduced.

Our results show that volatility transmission is not crisis invariant. Tests of equality of cross-crisis transmission patterns are rejected at all quantiles, suggesting that the architecture of spillovers is systematically dependent upon stress origin. Sector leadership shows meaningful turnover from episode to episode. Cyclical sectors are the primary drivers during tightening episodes, Energy sector reaches a certain significance during the pandemic shock, and Financials becomes the main transmitter during the tariff crisis. No sector is consistently found in the top-three set of transmitters across all of the crises. At the same time, the analysis shows a stable structural regularity: Among the fourteen regimes for whose coverage we have valid data, cyclical sectors are net transmitters and defensive sectors net receivers, with Utilities being the only sector that shows uniformly negative net spillovers over the period. Thus, for example, the picture which emerges is not one of total instability but of structured variation, with the origin of crisis changing the hierarchy of transmission in a larger and continuing cyclical - defensive division.

The remainder of the paper is organized as follows. Section 2 reviews related literature on connectedness, crisis transmission, and quantile methods. Section 3 describes data construction, sector classification, and regime definitions. Section 4 presents the QVAR and inference framework. Section 5 reports empirical findings and robustness evidence. Section 6 concludes and discusses policy implications and extensions.

2 Literature Review

Volatility spillovers across the financial and commodity markets have attracted substantial research because of their role in understanding risk propagation in the financial system. The literature on sector-level volatility dynamics shows that spillover patterns are not always consistently present or similar across all cases, and that these are dependent on the root causes of economic and financial disarray. Early studies of connectivity found that the cross-markets volatility transmissions become stronger during times of high uncertainty and that the top transmitters change with macroeconomic periods (Diebold and Yilmaz, 2009, 2012, 2014). More recent research focused on crisis-specific events gives additional evidence that disturbances arising from specific areas of the economy - bank failures and commodity disruptions, for instance - create unique channels of transmission through

which risk spreads through the equity market (Jawadi and Sellami, 2022; Akhtaruzzaman et al., 2021; Rehman et al., 2026). Research on the COVID-19 pandemic shows the role of non-traditional transmitters that can take over as the main risk sources in case of systemic disruptions of health and supply chains, but the studies of oil-price and credit shocks show that it is the energy and financial industries that act as the main transmitters of risk (Jawadi and Sellami, 2022; Akhtaruzzaman et al., 2021; Rehman et al., 2026). Collectively, these findings highlight that the dynamics of volatility are time-dependent, crisis-dependent, and distribution-dependent by nature, thus providing motivation for the methodological approach that will be adopted in this investigation.

Studies examining volatility spillovers have shown that market linkages are significantly altered during crises. Following initial studies that measured return and volatility spillovers across global equity markets (Diebold and Yilmaz, 2009), Diebold and Yilmaz (2012) provided a basic framework that utilized generalized forecast error variance decomposition (GFEVD) to measure connectedness that is invariant to the order of the variables. Diebold and Yilmaz (2014) extended this framework to include time-varying parameters utilizing rolling window estimation and introduced measures of network topology to assess the structure of connectivity among financial firms, allowing examination of how connectivity dynamics evolve during crisis and non-crisis periods. Their work showed that connectedness among financial firms evolves over time, with volatility spillovers increasing during crisis periods and return spillovers displaying gentler trends. Antonakakis et al. (2020) enhanced their earlier work by using time-varying parameter vector autoregressions to estimate their measures, thereby removing the need to select arbitrary rolling windows and enhancing the reliability of their estimates.

One of the questions that has emerged in the literature is which sectors or assets lead the transmission of volatility during crisis periods. There is evidence that the type of crisis determines the identity of the volatility transmitters. Laborda and Olmo (2021) estimated volatility spillovers across European equity sectors during the 2008 Global Financial Crisis (GFC) and the 2020 COVID-19 shock. They discovered that banks were the most significant volatility transmitters across Europe during the 2008 GFC, whereas energy and technology sectors dominated the transmission of volatility during the COVID-19 pandemic. Similar to the findings of Laborda and Olmo (2021), Ngene (2021) estimated spillovers across U.S. equity sectors and reported that typically, cyclical (or aggressive) sectors (sectors more sensitive to economic cycles) serve as net transmitters of volatility, whereas defensive sectors serve as net receivers of volatility; however, these roles can be reversed depending on the business cycle phase. These results support the traditional view that during periods of great uncertainty, investors flee to defensive sectors (such as utilities or consumer staples) and make those sectors absorb volatility, whereas cyclical sectors amplify the effects of shocks during expansions or in specific crisis conditions.

More recently, the literature has extended spillover analysis to quantile-based frameworks, recognizing that patterns of connectivity may differ between normal and extreme market conditions. Ando et al. (2022) introduce a new approach to quantile-based connectedness that utilizes quantile vector autoregression (QVAR) to extend the Diebold-Yilmaz framework to various portions of the

conditional distribution, demonstrating that credit risk shocks are propagated much more forcefully in both tails than at the conditional median. Chatziantoniou et al. (2021) applied a similar approach to interest rate swaps, illustrating that large and directional movements in interest rates have a profound impact on connectedness and demonstrating how the transmission of monetary policy functions in different market conditions. Kim et al. (2025) utilize this methodology to analyze the behavior of U.S. sector indexes, demonstrating that during extreme conditions spillover effects are relatively evenly distributed; however, during non-extreme conditions effects are concentrated, with elevated levels of connectedness post-COVID-19. Rehman et al. (2026) also found that there were substantial increases in U.S. sectoral return integration during and subsequent to the COVID-19 pandemic, with larger impacts at short-term horizons and in extreme market conditions.

Despite these important contributions, several limitations remain. First, the existing work on the origins of dependence disparities uses mostly descriptive comparisons of two or three crises, and does not formally statistically test to determine whether the observed differences reflect systematic patterns or simply random variation. Second, the methodological transparency about the quality of data is limited. Particularly, the influence of rolling window contamination on the validity of conclusions for short-duration crises. Third, existing studies are mostly restricted to studies before 2020 or the COVID-19 pandemic, missing crisis-origin dependence research at a wider range of crisis types, such as post-pandemic episodes of banking stress and trade policy shocks. The current study fills these gaps by investigating crisis-origin dependence across six different crisis types over a decade-spanning sample (2015–2025) using rigorous statistical methods: (1) panel regressions with Wald tests, (2) Spearman rank correlations, bootstrap confidence intervals, and rank-based analysis of the top-three transmitters, and (3) an explicit regime purity evaluation (to calibrate the reliability of our inferences with respect to the underlying data).

3 Data

3.1 Data and Sample Period

The daily implied volatility data for this study was collected from FirstRate Data for nine major U.S. sector exchange-traded funds (ETFs) from January 5, 2015 through March 14, 2025, totaling 2,435 trading days. The data are used for nine SPDR Select Sector ETFs that together cover all nine major sectors of the U.S. economy within the S&P 500 Index. They include: Materials (XLB), Energy (XLE), Financials (XLF), Industrials (XLI), Technology (XLK), Consumer Staples (XLP), Utilities (XLU), Healthcare (XLV), and Consumer Discretionary (XLY). These nine ETFs have a combined market capitalization of over 95% of the S&P 500, allowing us to provide comprehensive representation of the sectoral composition of the U.S. equity markets. Two sectors were excluded due to the quality of the data provided by FirstRate Data. Real Estate (XLRE) was excluded due to insufficient liquidity, limiting our ability to collect reliable data. Communication Services (XLC) was also excluded due to data sparseness in the 2015–2020 time frame,

For each of the nine sector ETFs included in the sample, we calculate the 30-day at-the-money

(ATM) implied volatility. The ATM IV represents the current market's forward-looking expectation of future volatility in that specific sector of the U.S. equity market. It provides better information for studying how volatility expectations across sectors change over time, compared to backward-looking measures such as realized volatility. Using a consistent 30-day horizon across all sectors enables comparison of volatility expectations both across sectors and over time. This eliminates confounding effects that may result from differences in the volatility term structure among the sectors in the sample.

3.2 Construction of Implied Volatility Series

To derive a consistent implied volatility metric for each sector ETF, we followed the CBOE's VIX calculation methodology (Chicago Board Options Exchange, 2019). To calculate the 30-day constant-maturity IV, we linearly interpolate between implied variances from options maturing before and after 30 days. Each day for each ETF, we obtained the mid-quote implied volatilities for ATM strikes (options with the closest strike to the current ETF price). We convert these to implied variances and interpolate them to exactly 30 days. If T_1 and T_2 are the two maturities (in days) and σ_1^2 and σ_2^2 are the implied variances for those maturities, then the 30-day variance is:

$$Var_{30} = Var_1 + \frac{(30 - T_1)}{(T_2 - T_1)}(Var_2 - Var_1). \tag{1}$$

The implied volatility for 30 days is then calculated as $\sigma_{30} = \sqrt{Var_{30}/30}$. This is the standard method (see Chicago Board Options Exchange (2019) for the VIX white paper) so that we have a way to compare our sector IVs both over time and between sectors. We used call options alone because put-call parity implies equivalent implied volatilities at the same strike and expiration (Hull, 2018). While ETF options are American-style, the early exercise premium is negligible for at-the-money options on non-dividend-paying or low-dividend ETFs over short horizons, making the European approximation appropriate for our purposes. We defined the at-the-money (ATM) strikes as the ones closest to the ETF spot price on each date, where spot prices were derived from the daily adjusted closing prices (Bollen and Whaley, 2004). We focused on ATM options in order to maximize liquidity and minimize model-dependent biases affecting out-of-the-money options.

We applied several filters to ensure data quality. We required a minimum open interest of 100 contracts and a maximum bid-ask spread of 30% relative to the midpoint implied volatility. We excluded options with no bids or asks and those exhibiting crossed markets (bid exceeding ask). When there were multiple expiries surrounding the 30-day target, we used the pair with the least interpolation gap and rejected all gaps larger than 40 days. To compute implied volatility, we used the midpoint of the bid and ask implied volatilities.[1] The resulting IV series were log-differenced ($\Delta \log \mathrm{IV}_t$) prior to estimation. All series are winsorized at the 1st and 99th percentiles to mitigate interpolation artifacts in the raw options data. A stationarity rationale for the transformation is

[1]Complete technical implementation details, including handling of ETF-specific liquidity constraints and alternative methodologies considered, are available from the authors upon request.

provided in Section 4.1. Table 1 reports descriptive statistics for the winsorized log-differenced series.

Table 1: Descriptive Statistics of Log-Differenced Implied Volatility Series

Sector	Obs.	Mean	SD	Skewness	Kurtosis	Min	Max	AC(1)
XLB	2435	-0.0004	0.0655	0.43	0.92	-0.1640	0.2078	-0.06
XLE	2435	-0.0001	0.0557	0.54	1.34	-0.1395	0.1906	-0.05
XLF	2435	-0.0000	0.1050	0.13	3.01	-0.3636	0.3598	-0.29
XLI	2435	-0.0000	0.0727	0.57	1.19	-0.1767	0.2420	-0.08
XLK	2435	0.0000	0.0726	0.50	1.06	-0.1738	0.2426	-0.10
XLP	2435	0.0002	0.0713	0.60	1.37	-0.1762	0.2499	-0.14
XLU	2435	-0.0000	0.0571	0.47	1.24	-0.1526	0.1925	-0.14
XLV	2435	-0.0002	0.0694	0.45	1.00	-0.1764	0.2165	-0.10
XLY	2435	0.0002	0.0684	0.51	1.32	-0.1744	0.2328	-0.07

Note: Each series represents daily log changes ($\Delta \log \mathrm{IV}_t$) in 30-day at-the-money implied volatility for the corresponding SPDR sector ETF. Sample period: January 2015 to March 2025 (2435 observations after differencing). All series are winsorized at the 1st and 99th percentiles to mitigate interpolation artifacts in the raw options data. Kurtosis is reported in excess of 3 (normal distribution). AC(1) denotes the first-order autocorrelation coefficient.

3.3 Cyclical-Defensive Classification

Following Ngene (2021), we group the nine sectors into three categories based on their sensitivity to the economic cycle. *Cyclical* sectors—Industrials (XLI), Consumer Discretionary (XLY), and Technology (XLK)—have a high beta value and a significant relationship between their earnings and GDP growth, making them sensitive to economic cycles. *Defensive* sectors—Utilities (XLU), Consumer Staples (XLP), and Healthcare (XLV)—provide essential products or services, and therefore have a steady stream of income regardless of economic conditions. The remaining three—Energy (XLE), Materials (XLB), and Financials (XLF)—are influenced by both commodity prices and monetary policy, as well as sector-specific factors, and therefore cannot be cleanly categorized under the cyclical-defensive paradigm; they are treated as a separate category.

3.4 Economic Regime Definitions

To identify crisis vs. non-crisis behavior, we segmented our dataset into distinct economic regimes. Each regime represents a distinct period characterized by specific macroeconomic or market conditions, classified as a crisis, calm, or transitional period. We identified fifteen distinct regimes from 2015 to 2025, based on significant events, Federal Reserve policy phases, and market volatility conditions. A summary of these regimes is provided in Table 2, along with their respective start/end dates, classification (Crisis, Calm, or Transitional), and a description of the environment during each regime.

There are six distinct crisis regimes with fundamentally different origins: the Oil/China Crisis

Table 2: Economic Regime Definitions: 2015–2025

Regime	**Period**	**Type**	**Description**
Pre Oil Calm	Jan 02, 2015 – Jun 30, 2015	Calm	Post-QE3 environment; moderate growth; stable oil prices; low volatility
Oil/China Crisis	Jul 01, 2015 – Feb 29, 2016	Crisis	China stock market crash (Aug 2015); oil price collapse below $30; yuan devaluation
Post Oil Recovery	Mar 01, 2016 – Dec 30, 2016	Transition	Oil price recovery; post-Brexit adjustment; US election; gradual Fed tightening begins
Low Vol Expansion	Jan 03, 2017 – Jan 31, 2018	Calm	Record low volatility; synchronized global growth; tax reform optimism; gradual rate hikes
Volmageddon Tightening	Feb 01, 2018 – Dec 31, 2018	Crisis	VIX spike Feb 2018 (Volmageddon); Q4 selloff from aggressive Fed tightening; trade war
Fed Pivot Recovery	Jan 02, 2019 – Feb 19, 2020	Transition	Fed reverses course with rate cuts; trade war de-escalation; pre-pandemic calm
COVID Crisis	Feb 20, 2020 – May 31, 2020	Crisis	Pandemic shock; fastest market crash in history; global lockdowns; emergency Fed intervention
COVID Recovery	Jun 01, 2020 – Dec 31, 2020	Transition	V-shaped recovery; massive fiscal/monetary stimulus; vaccine development
Post COVID Normalization	Jan 01, 2021 – Dec 31, 2021	Calm	Economic reopening; labor market recovery; supply chain disruptions; inflation begins rising
High Inflation Tightening	Jan 04, 2022 – Mar 07, 2023	Crisis	Fed rapid rate hikes (0% to 4.75%); CPI peak 9.1%; Russia-Ukraine war; energy price spike
SVB Banking Crisis	Mar 08, 2023 – Apr 30, 2023	Crisis	Silicon Valley Bank collapse; First Republic failure; Signature Bank; FDIC interventions
Post Banking Stress	May 01, 2023 – Jul 31, 2023	Transition	Banking sector stabilization; continued rate hikes to 5.25%; regional bank stress subsides
Rate Peak Hold	Aug 01, 2023 – Aug 31, 2024	Calm	Fed holds rates at 5.25–5.50%; inflation moderates; longest hold period; market stability
Easing Pivot	Sep 01, 2024 – Jan 19, 2025	Transition	Markets price rate cuts; Fed signals pivot; soft landing narrative; first rate cut Sep 2024
Tariff Crisis	Jan 20, 2025 – Mar 14, 2025	Crisis	Trump inauguration; sweeping tariff announcements; trade war escalation

Note: Sample spans January 5, 2015 to March 14, 2025 (2,435 trading days). Fifteen regimes: six crisis, four calm, and five transition periods. Regime definitions based on macroeconomic conditions, Federal Reserve policy, and financial market events.

(2015–2016, commodity and emerging-market shock), the Volmageddon Tightening (2018, volatility spike and aggressive Fed tightening), COVID-19 (2020, pandemic demand shock), High Inflation Tightening (2022–2023, policy-induced tightening), SVB Banking Crisis (2023, financial sector-specific), and the Tariff Crisis (2025, trade-policy shock). The remaining regimes comprise four calm periods (Pre Oil Calm, Low Vol Expansion, Post COVID Normalization, Rate Peak Hold) and five transitional periods. The inclusion of six crises spanning a decade strengthens generalizability relative to specifications limited to two or three episodes. Because rolling window estimation requires 200 trading days, the earliest regime (Pre Oil Calm) has no windows ending within its dates and is excluded, leaving fourteen regimes with available data.

4 Methods

This section presents the Quantile Vector Autoregression (QVAR) framework for measuring volatility spillovers, the hypothesis tests for crisis-origin dependence, and our approach to statistical inference under the serial and cross-sectional dependence induced by rolling window estimation.

4.1 Quantile Vector Autoregression (QVAR)

We apply the quantile connectedness methodology of Ando et al. (2022), which extends the Diebold-Yilmaz framework to measure spillovers at each quantile of the conditional distribution. Unlike a standard VAR that estimates conditional means, QVAR estimates coefficients at each quantile τ, providing a more robust measure of central tendency at the median ($\tau = 0.5$) and enabling distributional analysis across $\tau \in \{0.1, 0.2, \ldots, 0.9\}$.

Following Ando et al. (2022), let $\mathbf{Y}_t = (Y_{1t}, Y_{2t}, \ldots, Y_{Kt})'$ denote the K-dimensional vector of log-differenced implied volatility series at time t, where $Y_{kt} = \log(\mathrm{IV}_{kt}) - \log(\mathrm{IV}_{k,t-1})$ represents the daily change in log implied volatility for sector k. The τ-th quantile VAR model of order p is specified as:

$$Q_{Y_t}(\tau \mid \mathcal{F}_{t-1}) = \mathbf{c}(\tau) + \sum_{j=1}^{p} \mathbf{A}_j(\tau)\mathbf{Y}_{t-j} \tag{2}$$

where $Q_{Y_t}(\tau \mid \mathcal{F}_{t-1})$ denotes the conditional quantile (vector) of Y_t given past information $\mathcal{F}_{t-1}$. Here $\mathbf{c}(\tau)$ is a $K \times 1$ vector of intercepts in quantile τ, and each $\mathbf{A}_j(\tau)$ is a $K \times K$ matrix of autoregressive coefficients in quantile τ. We estimate the model using quantile regression, which minimizes the check function of Koenker and Bassett Jr (1978) applied to the VAR system:

$$\min_{\mathbf{c}(\tau), \mathbf{A}_1(\tau), \ldots, \mathbf{A}_p(\tau)} \sum_{t=1}^{T} \rho_\tau \left(\mathbf{Y}_t - \mathbf{c}(\tau) - \sum_{j=1}^{p} \mathbf{A}_j(\tau)\mathbf{Y}_{t-j} \right) \tag{3}$$

where $\rho_\tau(u) = u(\tau - \mathbb{I}(u < 0))$ is the asymmetric loss function. Following Ando et al. (2022), we obtain the QVAR companion form at each quantile and derive the corresponding moving average representation, from which the generalized forecast error variance decomposition is computed. The

innovation covariance matrix required for the GFEVD is estimated from the quantile regression residuals at each τ, as detailed in Ando et al. (2022).

We use log-differenced implied volatility ($\Delta \log \mathrm{IV}_t$) rather than log levels. The log-level series produce conflicting stationarity results: ADF tests reject the unit root null ($p < 0.05$) while KPSS tests reject the stationarity null ($p < 0.01$) for all nine sectors—a pattern typical of highly persistent but mean-reverting processes (Hamilton, 1994). Log-differencing resolves this ambiguity: all differenced series are clearly stationary under both tests. While differencing precludes identification of long-run cointegrating relationships, our research question concerns the short-run propagation of volatility shocks across sectors during crises rather than long-term comovements. The log-differenced form has the intuitive interpretation of daily percentage changes in volatility expectations.

To capture time-varying dynamics, we estimate the QVAR using rolling windows of 200 trading days—large enough for reliable parameter estimation with $K = 9$ variables, yet small enough to track changing market conditions. With 2,435 observations, this yields 2,236 rolling windows estimated at nine quantiles ($\tau \in \{0.1, 0.2, \ldots, 0.9\}$). Each 200-day window inevitably spans multiple regime periods. We assign each window to the regime prevailing on its end date, ensuring a unique and reproducible regime classification.

We calculate the regime purity for each window as the proportion of the 200 days contained within the designated regime. Rolling windows with high regime purity offer clearer identification of the dynamic effects of crises on markets, whereas rolling windows with low regime purity combine elements of multiple regimes. The regime assignment statistics are provided in Table 6 to enable transparency regarding the quality of data that underlies our comparisons of regimes.[2]

We determine the optimal lag order for $K = 9$ sectors using the Bayesian information criterion (BIC) on the entire sample, which selects $p = 2$ as the best choice. This lag order is kept constant for all rolling windows and quantiles to ensure that spillover estimates can be compared across regimes.

4.2 Spillover Measures

Following Diebold and Yilmaz (2012), we compute spillover measures from the generalized forecast error variance decomposition (GFEVD), which is invariant to variable ordering. We set the forecast horizon to $H = 10$ trading days (two weeks), which is the standard choice in the spillover literature (Diebold and Yilmaz, 2012, 2014). The GFEVD decomposes the H-step-ahead forecast error variance of variable i into contributions from shocks to all variables $j = 1, \ldots, K$.

Let $\tilde{d}^g_{ij}(H)$ denote the contribution of variable j to the forecast error variance of variable i at horizon H. Normalizing by row sums yields:

$$\tilde{\phi}^g_{ij}(H) = \frac{\tilde{d}^g_{ij}(H)}{\sum_{j=1}^{K} \tilde{d}^g_{ij}(H)} \tag{4}$$

These normalized shares satisfy $\sum_{j=1}^{K} \tilde{\phi}^g_{ij}(H) = 1$ and $\sum_{i=1}^{K} \sum_{j=1}^{K} \tilde{\phi}^g_{ij}(H) = K$. From the

[2]Alternative assignment rules (majority rule, 75% purity threshold) yield qualitatively identical results. These are available from the corresponding author upon request.

GFEVD, we compute the directional spillover measures by sector i from/to all others as:

$$\text{FROM}_i^g(H) = \sum_{j=1, j\neq i}^{K} \tilde{\phi}_{ij}^g(H) \tag{5}$$

$$\text{TO}_i^g(H) = \sum_{j=1, j\neq i}^{K} \tilde{\phi}_{ji}^g(H) \tag{6}$$

The net directional spillover is the difference between "To others" and "From others". A positive *NET* indicates that the sector is a net transmitter of volatility shocks, while a negative *NET* indicates a net receiver.

$$\text{NET}_i^g(H) = \text{TO}_i^g(H) - \text{FROM}_i^g(H) \tag{7}$$

The system-wide total connectedness index (TCI) measures the average contribution of spillovers from shocks to all other variables:

$$\text{TCI}^g(H) = \frac{\sum_{i,j=1, i\neq j}^{K} \tilde{\phi}_{ij}^g(H)}{\sum_{i,j=1}^{K} \tilde{\phi}_{ij}^g(H)} \times 100. \tag{8}$$

The TCI ranges from 0 to 100, with higher values indicating greater interconnectedness throughout the system.

4.3 Testing Crisis-Origin Dependence

We formulate three hypotheses. First, cyclical sectors consistently transmit more volatility than defensive sectors across all market regimes—the cyclical-defensive divide is positive and never reverses sign, even during crises. Second, the sectors leading volatility transmission differ depending on the origin of the crisis: a pandemic, a monetary tightening, and a banking crisis each produce distinct top transmitters. This is our central hypothesis of crisis-origin dependence. Third, Utilities exhibits safe-haven behavior in all regimes, remaining a net receiver of volatility regardless of crisis type.

To test these hypotheses, we employ a fixed-effects panel regression complemented by cross-crisis correlation analysis, bootstrap confidence interval tests, rank-based top-3 transmitter analysis, and regime purity assessment.

4.3.1 Panel Regression Specification

We employ fixed effects panel regression to test whether the cyclical-defensive divide differs systematically across crisis types. Our interaction specification is:

$$\text{NET}_{it}(\tau) = \alpha_i(\tau) + \sum_{c\in C} \beta_c(\tau) \cdot D_t^c + \sum_{c\in C} \gamma_c(\tau) \cdot D_t^c \times \text{Cyclical}_i + \varepsilon_{it}(\tau) \tag{9}$$

where $\text{NET}_{it}(\tau)$ is the NET spillover for sector i in window t at quantile τ, $\alpha_i(\tau)$ are sector fixed effects, $C = \{\text{Oil/China}, \text{Volmageddon}, \text{COVID}, \text{Inflation}, \text{Banking}, \text{Tariff}\}$ denotes crisis types with D_t^c as indicator variables, and Cyclical_i equals one for cyclical sectors (XLI, XLY, XLK) and zero for defensive sectors (XLU, XLP, XLV). The three remaining sectors (XLE, XLB, XLF) are excluded from the panel regression because they do not fit cleanly into either classification; however, these sectors are included in the full QVAR estimation and appear in the complementary rank-based and top-3 transmitter analyses. The interaction coefficients $\gamma_c(\tau)$ measure the cyclical-defensive divide during each crisis type. We test crisis-origin dependence through a Wald test:

$$H_0 : \gamma_{\text{Oil/China}}(\tau) = \gamma_{\text{Volmageddon}}(\tau) = \gamma_{\text{COVID}}(\tau) = \gamma_{\text{Inflation}}(\tau) = \gamma_{\text{Banking}}(\tau) = \gamma_{\text{Tariff}}(\tau) \quad (10)$$

Rejection indicates that the aggregate cyclical-defensive divide varies significantly by crisis origin.[3] We note that this formal test captures group-level asymmetry rather than individual sector leadership; the complementary rank-based and top-3 transmitter analyses reported in Section 5 provide the sector-level evidence for crisis-specific leadership patterns. Standard errors are clustered by rolling window to account for temporal dependence from the overlapping window structure.

4.3.2 Complementary Evidence

The panel regression is validated using four additional methods: (i) cross-crisis Spearman rank correlations of sector NET spillovers; (ii) bootstrap confidence intervals for the cyclical-defensive divide by crisis type; (iii) rank-based analysis of top-3 transmitter compositions; and (iv) regime purity assessment. The baseline cyclical-defensive divide is assessed by comparing mean NET spillovers for cyclical versus defensive sectors across all fourteen regimes. Utilities' safe-haven status is evaluated through sign-stability analysis.

4.4 Statistical Inference Under Dependence

Rolling window estimation with 200-day windows creates substantial overlap (99.5%) between consecutive estimates, inducing serial dependence that violates standard test assumptions. Additionally, joint QVAR estimation across all nine sectors creates cross-sectional dependence. We address these issues through two approaches. First, for the panel regression, we cluster standard errors by rolling window, which accounts for cross-sectional dependence within each window but may understate uncertainty due to the near-complete overlap between consecutive windows. Second, for divide comparisons across regimes, we employ regime-level aggregation, which substantially reduces the effective overlap problem as regime means are computed over non-overlapping groups of windows. We acknowledge that the effective sample size for inference is considerably smaller than the nominal number of rolling windows, and interpret marginal significance levels with appropriate caution.

[3] Because this Wald test is applied at each of nine quantiles, a multiple-testing concern arises. Applying a Bonferroni correction ($\alpha_{\text{corrected}} = 0.05/9 \approx 0.0056$), all nine quantiles survive the stricter threshold, with F-statistics ranging from 66.0 to 303.3 and all p-values effectively zero.

5 Results

We first establish the baseline transmission patterns that hold in all market conditions (Section 5.1), then examine how these patterns shift during each type of crisis (Section 5.2), and finally assess the safe-haven status of utilities (Section 5.3). All results are based on QVAR spillover measures computed at the median quantile ($\tau = 0.5$) unless otherwise noted.

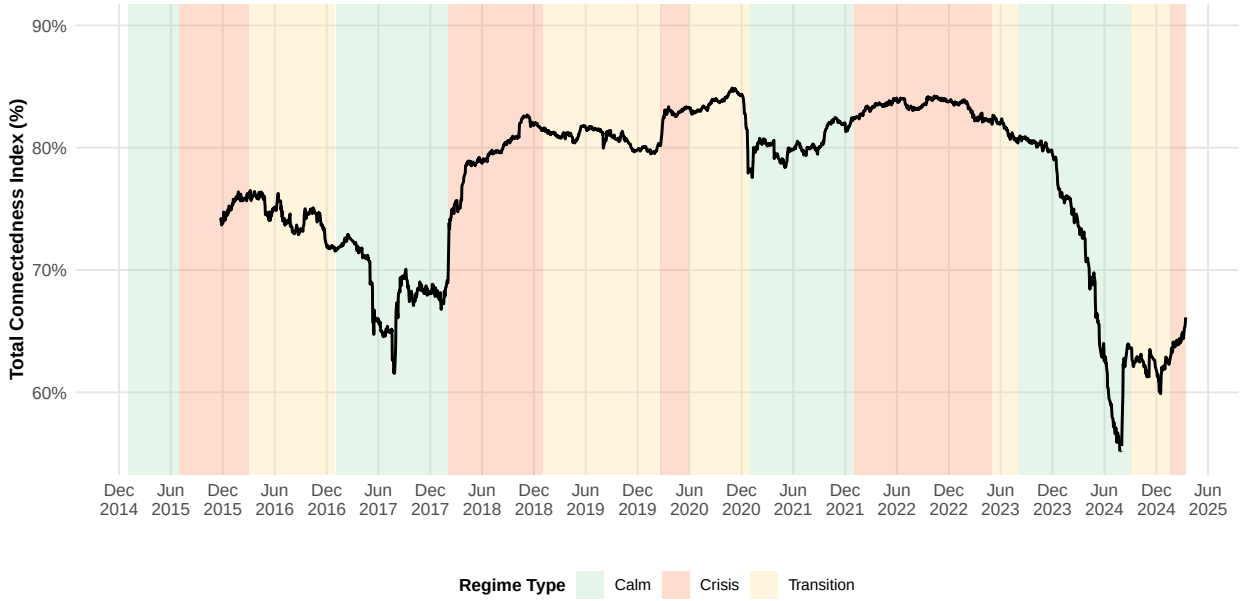

Figure 1: Total Connectedness Index: 2015–2025

Note: Total Connectedness Index at median quantile (q=0.5) with economic regime shading. Red regions indicate crisis periods, green indicates calm periods, and orange indicates transition periods. TCI measures system-wide spillovers in volatility changes across all sector ETFs.

5.1 Baseline Transmission Patterns and the Cyclical-Defensive Divide

Figure 1 shows the total connectedness index (TCI) from 2015–2025. We see that the TCI increases during crisis periods (red shaded areas) and decreases during calm periods (green shaded areas). This suggests that system-wide volatility spillovers respond to macroeconomic conditions. Before determining specific transmission patterns, we establish a baseline: cyclical sectors are consistently net transmitters of volatility while defensive sectors are net receivers.

The cyclical-defensive divide (as reported in Table 3) is positive for all fourteen regimes, ranging from +1.19 (COVID Recovery) to +27.44 (Low Vol Expansion), and an average value of +13.80 percentage points. Although the cyclical-defensive divide narrowed during some crises (SVB Banking: +7.33; High Inflation: +8.85), it remained large during others (Tariff: +19.21; Volmageddon: +19.96). Even during times of extreme stress, the divide persisted. The largest differences were observed during calm and transitional regimes, consistent with the idea that crises reduce the degree of sectoral asymmetry through the absorption of systemic risks (Baruník et al., 2016; Rehman et al.,

2026).

Figure 2 shows that the cyclical mean (blue line) remains higher than the defensive mean (red line) for the whole duration of the sample period. During crises, the gap narrows. Figure 3 confirms this on the individual sector level. Cyclical sectors (XLI, XLY, XLK) are positive in most regimes while defensive sectors (XLU, XLP, XLV) are negative.

Table 3: Cyclical-Defensive Divide Across Economic Regimes

Regime	**Cyclical**	**Defensive**	**Divide**
Oil/China Crisis	11.33	-2.38	13.70
Post Oil Recovery	10.35	-3.89	14.24
Low Vol Expansion	16.57	-10.87	27.44
Volmageddon Tightening	11.07	-8.89	19.96
Fed Pivot Recovery	9.22	-10.57	19.79
COVID Crisis	3.27	-6.48	9.74
COVID Recovery	-0.10	-1.29	1.19
Post COVID Normalization	6.02	-5.70	11.72
High Inflation Tightening	6.44	-2.41	8.85
SVB Banking Crisis	6.34	-1.00	7.33
Post Banking Stress	6.43	-2.92	9.35
Rate Peak Hold	7.44	-9.81	17.26
Easing Pivot	2.63	-10.80	13.43
Tariff Crisis	5.19	-14.02	19.21

Note: Mean NET spillovers (TO minus FROM) at median quantile (q=0.5). Cyclical sectors: XLI, XLY, XLK. Defensive sectors: XLU, XLP, XLV. Divide = Cyclical – Defensive. All values in percentage points. Positive divide indicates cyclical sectors transmit more than defensive sectors receive. Pre Oil Calm regime omitted because no rolling windows fall entirely within its dates.

5.2 Crisis-Origin Dependence in Sector Transmission

Table 4 presents the interaction coefficients (γ_c) from the panel regression model specified in Equation 9. This captures how the cyclical-defensive divide varies across crisis types after controlling for sector fixed effects.

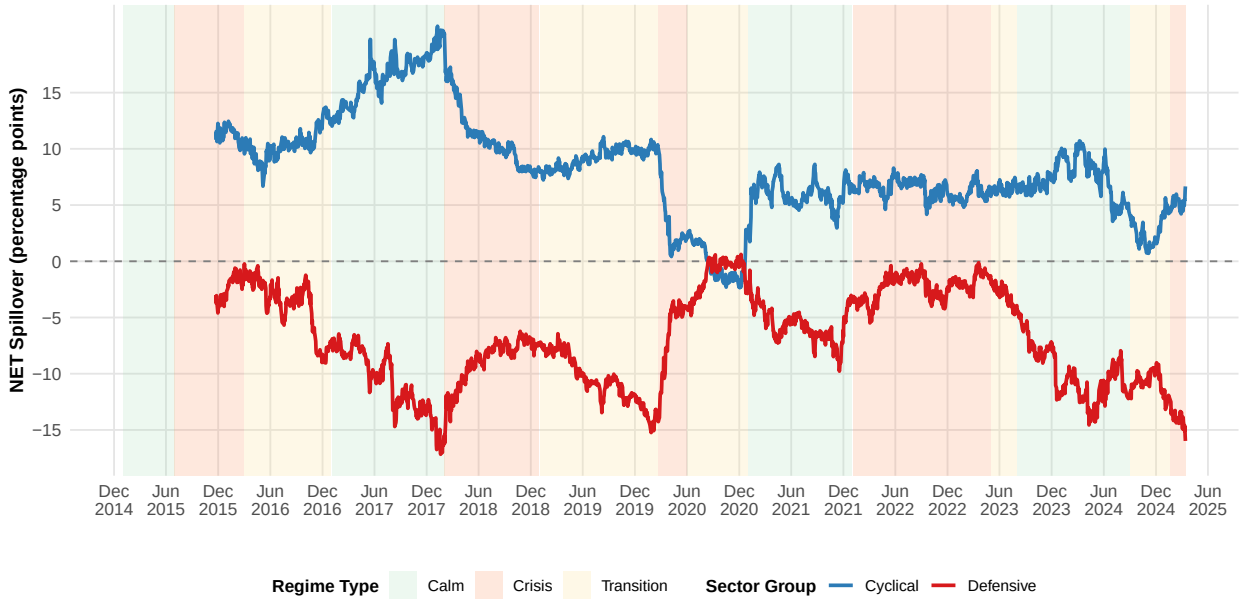

Figure 2: Cyclical vs. Defensive Sector Spillovers: 2015–2025

Note: Mean NET spillovers for cyclical sectors (XLI, XLY, XLK) and defensive sectors (XLU, XLP, XLV) at median quantile (q=0.5). The persistent gap between blue and red lines visualizes the cyclical-defensive divide. The gap narrows during crises and expands during calm periods throughout the sample.

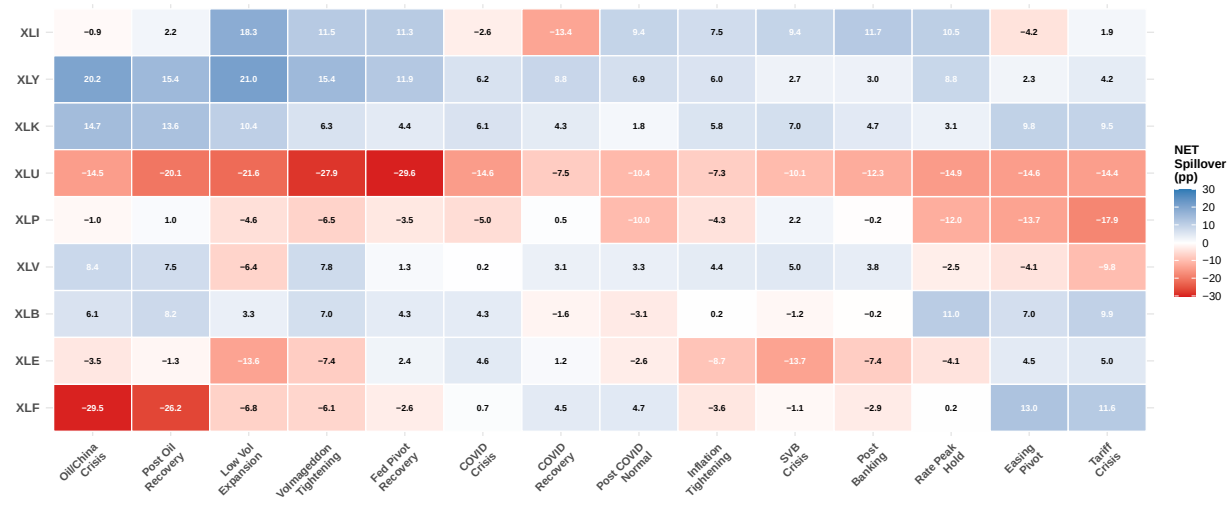

Figure 3: Sector NET Spillovers Across All Economic Regimes: 2015–2025

Note: Heatmap shows mean NET spillovers for all sectors across all fourteen regimes with available data at median quantile (q=0.5). Cyclical sectors generally transmit more than defensive sectors. Utilities exhibits safe-haven characteristics with negative NET across all regimes. Variation in transmission patterns illustrate crisis-origin dependence.

Table 4: Panel Regression Tests: Crisis-Origin Dependence in Cyclical-Defensive Divide

	Interaction Coefficients (γ_c)						**Wald Test**	
Quantile	Oil/China	Volmageddon	COVID	Inflation	SVB Banking	Tariff	F-stat	p-value
$\tau = 0.1$	-4.08^{***}	4.62^{***}	-2.06^{***}	-0.35^{*}	-2.53^{***}	8.32^{***}	179.3	<0.001
	(0.61)	(0.23)	(0.45)	(0.19)	(0.51)	(0.51)		
$\tau = 0.2$	-1.47^{***}	4.53^{***}	-2.35^{***}	-2.25^{***}	-3.85^{***}	13.57^{***}	287.9	<0.001
	(0.35)	(0.26)	(0.45)	(0.20)	(0.36)	(0.54)		
$\tau = 0.3$	-1.48^{***}	6.36^{***}	-1.40^{***}	-1.74^{***}	-2.53^{***}	12.55^{***}	253.8	<0.001
	(0.53)	(0.30)	(0.53)	(0.24)	(0.43)	(0.60)		
$\tau = 0.4$	1.33^{**}	8.58^{***}	0.12	-2.01^{***}	-3.44^{***}	10.26^{***}	241.3	<0.001
	(0.63)	(0.37)	(0.61)	(0.27)	(0.55)	(0.76)		
$\tau = 0.5$	2.58^{***}	8.84^{***}	-1.38^{**}	-2.27^{***}	-3.79^{***}	8.09^{***}	222.3	<0.001
	(0.67)	(0.38)	(0.58)	(0.28)	(0.55)	(0.82)		
$\tau = 0.6$	2.17^{***}	8.43^{***}	-1.65^{***}	-1.85^{***}	-3.47^{***}	7.13^{***}	192.1	<0.001
	(0.57)	(0.38)	(0.54)	(0.28)	(0.52)	(0.76)		
$\tau = 0.7$	0.23	6.31^{***}	-1.53^{***}	-1.71^{***}	-4.13^{***}	7.15^{***}	163.4	<0.001
	(0.51)	(0.34)	(0.52)	(0.25)	(0.51)	(0.60)		
$\tau = 0.8$	0.33	7.26^{***}	2.05^{***}	-0.44^{**}	-3.61^{***}	11.47^{***}	303.3	<0.001
	(0.38)	(0.26)	(0.48)	(0.19)	(0.40)	(0.48)		
$\tau = 0.9$	-0.23	0.83^{*}	1.16^{**}	0.48	-2.33^{***}	7.82^{***}	66.0	<0.001
	(0.48)	(0.43)	(0.51)	(0.55)	(0.61)	(0.50)		
Mean	**−0.03**	**6.20**	**−0.78**	**−1.35**	**−3.30**	**9.59**	**212.2**	**<0.001**

Note: Interaction coefficients $\gamma_c(\tau)$ from Equation 9 capture crisis-specific variation in the cyclical-defensive divide after controlling for sector fixed effects. Robust standard errors (HC1) in parentheses. Wald F-test tests H_0: γ equal across all six crisis types (df=5). N=4,314 sector-window observations per quantile (6 sectors $\times$ 719 crisis windows). All nine quantiles reject H_0 at $p < 0.001$, including under Bonferroni correction ($\alpha/9 \approx 0.0056$). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

The Wald test rejects equality of the cyclical-defensive divide across crisis types at all nine quantiles, with F-statistics ranging from 66.0 to 303.3 and all p-values effectively zero, including under Bonferroni correction ($\alpha/9 \approx 0.0056$). At the median quantile, the divide is largest during Volmageddon (+19.96) and Tariff (+19.21), and smallest during SVB Banking (+7.33) and High Inflation (+8.85) (Wald $F = 222.3$, $p < 0.001$).

Our results are robust to small perturbations of the regime boundaries: all Wald tests were still significant at $p < 0.001$ across all five perturbations (± 5 and ± 10 trading days) and all nine quantiles, with F-statistics ranging from 207 to 241 at the median quantile. Additionally, permutation tests on top-3 transmitter composition demonstrate that crisis-specific sector leadership is statistically significant at all nine quantiles for the intersection test ($p < 0.01$). Table 5 presents the percentage of quantiles for each sector in the top three transmitter positions, separated by crisis type, showing a high degree of sector role switching among the six crises.

Table 5: Sector Role Switching Across Crisis Types: Top-3 Transmitter Frequencies

Sector	**Group**	**Oil/China**	**Volmag.**	**COVID**	**Inflation**	**SVB**	**Tariff**	**Mean**
Cyclical Sectors:								
XLK (Technology)	Cyclical	100%	11%	67%	89%	100%	100%	78%
XLY (Cons. Disc.)	Cyclical	78%	100%	100%	89%	0%	33%	67%
XLI (Industrials)	Cyclical	0%	89%	0%	100%	100%	0%	48%
Defensive Sectors:								
XLV (Healthcare)	Defensive	67%	78%	0%	11%	100%	0%	43%
XLP (Cons. Staples)	Defensive	0%	0%	0%	0%	0%	0%	0%
XLU (Utilities)	Defensive	0%	0%	0%	0%	0%	0%	0%
Other Sectors:								
XLB (Materials)	Other	56%	22%	22%	11%	0%	67%	30%
XLF (Financials)	Other	0%	0%	22%	0%	0%	100%	20%
XLE (Energy)	Other	0%	0%	89%	0%	0%	0%	15%
Group Representation in Top-3 Slots:								
Cyclical (3 sectors)	Exp. 33%	59%	67%	56%	93%	67%	44%	64%
Defensive (3 sectors)	Exp. 33%	22%	26%	0%	4%	33%	0%	14%
Other (3 sectors)	Exp. 33%	19%	7%	44%	4%	0%	56%	22%

Note: Proportion of quantiles (out of 9) in which each sector appears in top-3 NET transmitter positions, by crisis type. No sector appears in the top-3 for $\geq$50% of quantiles across all six crises. Key role switches: Energy (XLE) 89% in COVID, 0% elsewhere; Financials (XLF) 100% in Tariff, 22% in COVID, 0% in all other crises; Industrials (XLI) 100% in Inflation and SVB Banking, 0% in COVID, Oil/China, and Tariff.

The Inflation Crisis provides the strongest test case, with the best data quality among the six crises (65.8% regime purity, 291 windows). Cyclical sectors dominate the top-3 positions: Industrials (XLI) appears in the top-3 at all nine quantiles, followed by Consumer Discretionary (XLY) and Technology (XLK) at 89% each. The defensive sectors represent less than 4% of the top-3 spots. The dominance of cyclical sectors during the Inflation Crisis happens primarily due to their exposure to Fed policy through capital expenditures and earnings cyclicality.

The COVID-19 Crisis represents an example of significant role-switching behavior. Energy

(XLE) did not occur in Inflation top-3; however, it represented 89% of the COVID top-3 spots, consistent with the collapse in oil demand during lockdowns. Industrials was the dominant Inflation transmitter, but was absent from the COVID top-3 transmitters. This complete reversal of roles between two crises occurring within two years replicates Laborda and Olmo (2021)'s finding and cannot be attributed to contamination alone.

The Tariff Crisis (2025) had the most distinct pattern. Financials (XLF), which appeared in the top-3 at only 2 of 45 quantile-crisis observations across the five previous crises, emerged as the leader in transmission (NET = +11.59) at every quantile. Materials (XLB) also rose to prominence (67% of quantiles), consistent with direct trade-policy exposure. This first occurrence of Financials as the leading transmitter clearly indicates that crisis origins determine sectoral leadership.

Finally, the remaining crises reinforce this variety. In the Oil/China Crisis (2015–2016), both Consumer Discretionary and Technology were the leaders, given their sensitivity to emerging-market demand. Volmageddon (2018) was led by Consumer Discretionary and Industrials, with the surprising appearance of Healthcare (XLV) in 78% of quantiles. In the SVB Banking Crisis, Healthcare was present in the top-3 at every quantile—a level of defensive sector prominence that was uniquely found in this crisis.

The Spearman rank correlation coefficients quantify the magnitude of role reversals. The average of the 15 pairwise crisis comparisons is $\rho_s = 0.47$, ranging from $\rho_s = 0.03$ (Tariff vs. SVB Banking) to $\rho_s = 0.90$ (Inflation vs. SVB Banking, temporally adjacent crises). None of the sectors appear in the top-3 in all six crises, and therefore we conclude that there is crisis-specific leadership. The cyclical-defensive divide by crisis type is illustrated in Figure 4, along with bootstrap confidence intervals confirming that the divide varies significantly across crisis origins.

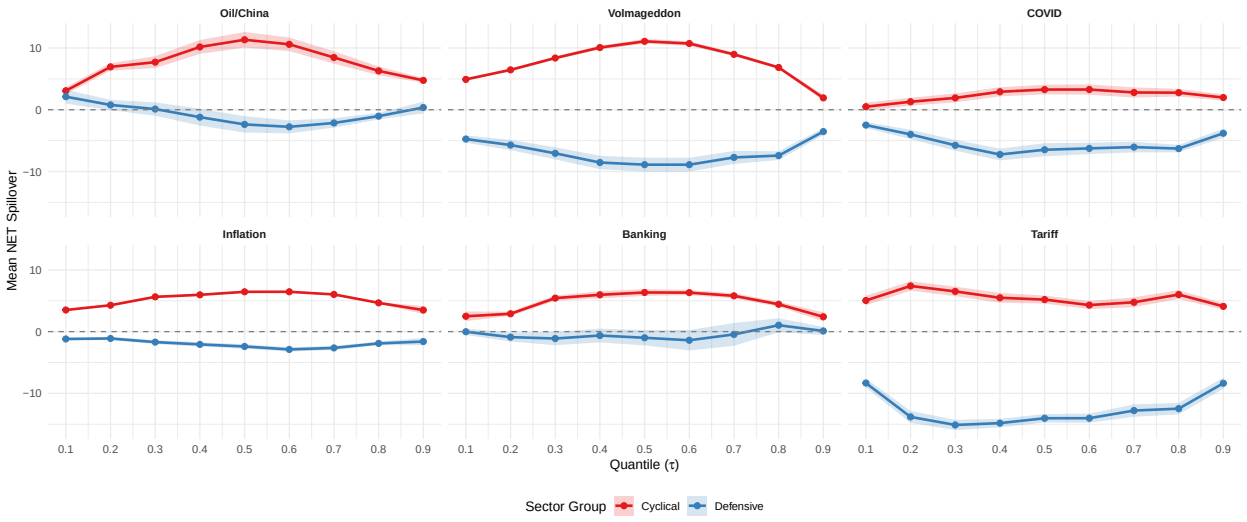

Figure 4: Cyclical-Defensive Divide by Crisis Type: Mean NET Spillovers Across Quantiles
Note: Shaded regions represent bootstrap 95% CIs (1,000 replications). The gap between the lines represents the cyclical-defensive divide.

Figure 5 shows the directed spillover networks for the six crises and illustrates how the network

topology reconfigures with each crisis origin.

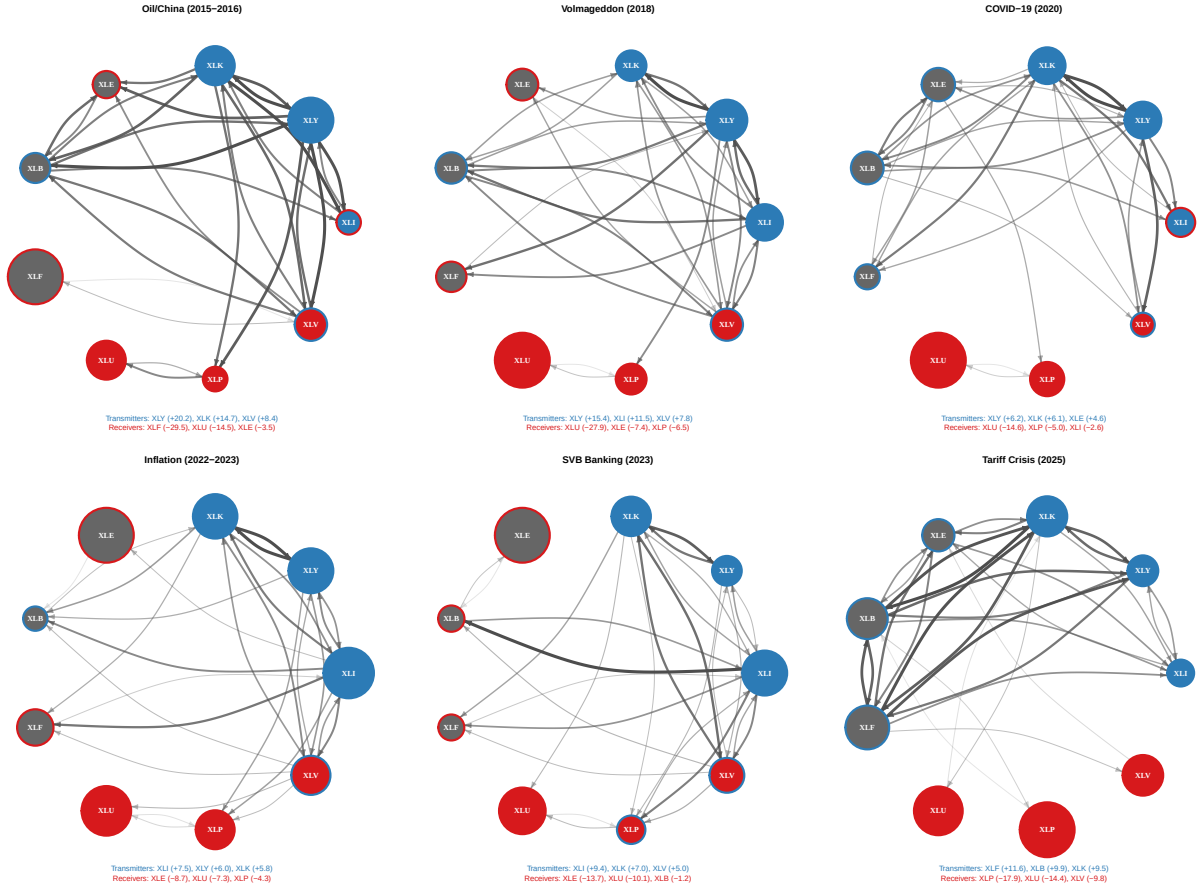

Figure 5: Directed Spillover Networks Across Crisis Types

Note: Directed network graphs of pairwise volatility spillovers at the median quantile ($\tau = 0.5$) for each crisis regime. Node color indicates sector classification: blue = cyclical (XLI, XLY, XLK), red = defensive (XLU, XLP, XLV), grey = other (XLE, XLB, XLF). Node size is proportional to the absolute value of NET spillovers. Node border color indicates direction: blue border = net transmitter, red border = net receiver. Edge thickness and opacity reflect the magnitude of pairwise spillovers; only edges above the 65th percentile are shown, with minimum connectivity ensured for all nodes. Layout is fixed across panels for comparability.

Table 6 reports the regime purity statistics. The Inflation and Volmageddon periods have the best data quality (65.8% and 55.6% mean purity) and provide the strongest basis for inference. The COVID, SVB Banking, and Tariff crises have shorter durations and therefore lower purity levels. Nevertheless, the differences in sectoral leadership between crises cannot be explained by contamination, as the Wald test achieves F-statistics greater than 200 at most quantiles.

The results are robust in multiple dimensions. The cyclical-defensive divide is positive at all nine quantiles ($\tau \in \{0.1, \ldots, 0.9\}$), ranging from 6.15 to 13.80, with the divide largest at the middle quantiles and more compressed in the tails. Detailed robustness analyzes such as quantile-by-quantile results, Bonferroni corrections, regime boundary sensitivity, permutation tests, and pre-COVID subsample comparisons are available from the corresponding author upon request.

Table 6: Regime Purity Statistics: Data Quality Assessment by Crisis Type

Crisis	Duration	Windows	Mean Purity	Median Purity	Quality
High Inflation Tightening	291	291	65.8%	73.0%	High
Oil/China Crisis	156	66	61.8%	61.7%	Moderate
Volmageddon Tightening	224	224	55.6%	56.2%	Moderate
COVID Crisis	67	67	17.0%	17.0%	Low
SVB Banking Crisis	37	37	9.5%	9.5%	Low
Tariff Crisis	34	34	8.8%	8.8%	Low

Note: Regime purity measures the proportion of days within each 200-day rolling window that fall within the assigned crisis regime. Crises ordered by mean purity. High Inflation's long duration yields reliable estimates; Oil/China and Volmageddon provide moderate quality. COVID, SVB Banking, and Tariff have short durations producing severe contamination. Duration in trading days; windows assigned based on end date.

5.3 Utilities Safe-Haven Characteristics

Table 7: Sector-Level Spillover Statistics Across All Regimes

Sector	Group	Mean NET	SD	Positive/Total
XLY	Cyclical	9.47	6.34	14/14
XLK	Cyclical	7.24	3.83	14/14
XLI	Cyclical	5.19	8.37	10/14
XLV	Defensive	1.56	5.49	10/14
XLP	Defensive	-5.37	6.06	3/14
XLU	Defensive	-15.70	6.86	0/14
XLB	Commodity	3.95	4.51	10/14
XLE	Commodity	-3.20	6.37	5/14
XLF	Financial	-3.15	12.02	6/14

Note: Mean NET spillovers across all 14 regimes with available data at median quantile ($\tau = 0.5$). Positive values indicate net transmission, negative values indicate net reception. SD measures variability across regimes. XLY and XLK are positive in all regimes; XLU is negative in all regimes.

Utilities (XLU) is the only sector to exhibit negative NET spillovers in all fourteen regimes. With a mean NET of -15.70 across the full sample—2.9 times larger than the next-most defensive sector (XLP: -5.37) and opposite in sign to Healthcare (XLV: $+1.56$)—Utilities is the dominant net receiver of volatility (Table 7). Even during the least negative regime (High Inflation Tightening: -7.28), XLU receives more than most sectors transmit. During the most negative (Fed Pivot Recovery: -29.57), it receives nearly five times the typical cyclical transmission. The safe-haven characteristic persists across commodity shocks, monetary tightening, pandemic disruptions, banking crises, and trade-policy shocks alike.

Figure 6 compares the three defensive sectors. XLU (dark blue) is negative in all regimes, while

XLP (light blue) in eleven of the fourteen regimes. XLV (orange) is predominantly a net transmitter, with positive values in ten regime periods. Therefore, the "defensive" label alone does not guarantee the safe-haven characteristics. Rather, XLP serves as a conditional safe haven and XLV functions primarily as a transmitter. XLU's safe-haven status reflects its regulated monopoly structure, stable cash flows, and high dividend yields (typically 3–5%). The bond-like income characteristics attract defensive capital flows during uncertainty.

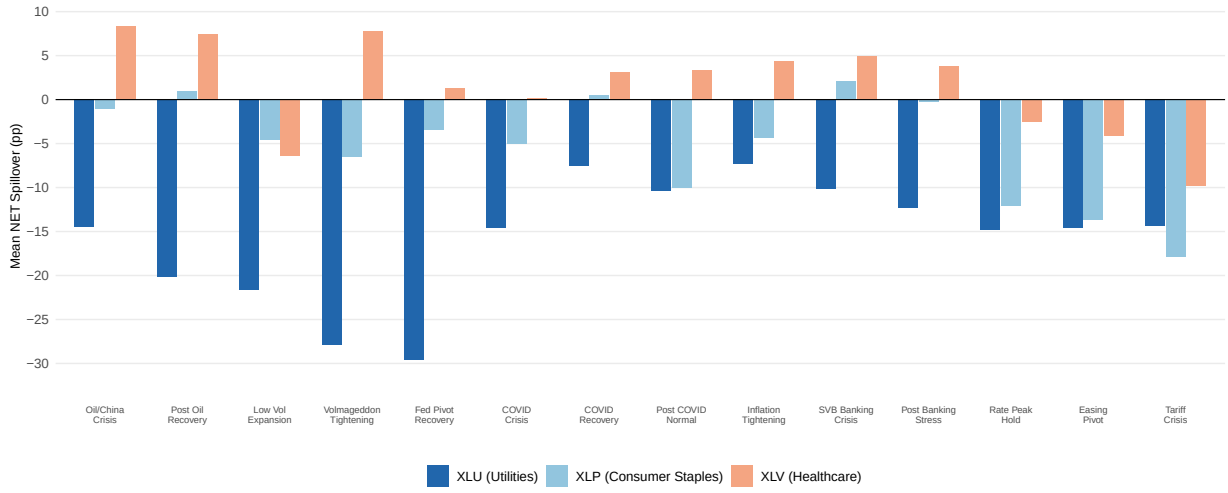

Figure 6: Defensive Sector Comparison: NET Spillovers Across Regimes

Note: Mean NET spillovers at the median quantile ($\tau = 0.5$) for three defensive sectors across all regimes in chronological order. Utilities (XLU) is negative in all fourteen regimes (14/14), Consumer Staples (XLP) is negative in eleven (11/14), and Healthcare (XLV) is negative in only four (4/14). Only XLU qualifies as a pure safe haven.

6 Conclusion

Our empirical results provide formal statistical support that the source of a crisis affects volatility spillover dynamics across equity sectors. Each type of crisis produces a unique set of sector transmitters and transmission intensities. This crisis-origin dependence operates through two channels: switching of sector leadership due to changes in rank correlation and variations in transmission intensity among sectors. We observe significant differences in transmission intensity even when we see similar ranks among the same sectors.

The observed patterns can also be explained by known macroeconomic mechanisms: monetary tightening elevates volatility transmission from cyclical sectors through increased borrowing costs; pandemics concentrate shocks in the energy sector due to the collapse of physical demand; trade policy elevates financials through trade finance and credit repricing channels; and commodity price shocks propagate through sectors exposed to international trade. As a result, crisis-origin dependence occurs because each crisis generates sector-specific fundamental shocks that then propagate through the volatility network—not through some generic "risk-off" mechanism that would lead to the same

pattern independent of the origin of the shock. The empirical results support the interpretation of Bekaert et al. (2014), suggesting that contagion intensity is related to economic fundamentals and not simply to panic.

The persistence of the cyclical-defensive divide indicates the presence of some underlying factors driving this structure. There are a few complementary explanations consistent with our empirical results. Since cyclical sectors have longer-duration cash flows, are typically more levered, and have more cyclically priced investments, they are more sensitive to interest rates (and therefore to monetary policy tightening). This increases their ability to serve as volatility transmitters. Furthermore, the growth in ETF-based strategies and option trading volumes could enhance the ability of cyclical sectors (due to greater liquidity and tighter option spreads) to transmit implied volatility to defensive sectors through the derivatives markets. Although the divide decreases during the crisis, it remains positive, confirming that the cyclical-defensive asymmetry is structural rather than a regime-dependent phenomenon.

Our findings reinforce the crisis-origin dependence in volatility transmission. In the Inflation and Volmageddon crises (high confidence), monetary tightening is characterized by cyclical sectors as the primary volatility transmitters. In the Oil/China crisis (moderate confidence), commodity-price and emerging-market shocks elevate Consumer Discretionary and Technology as primary transmitters. During the COVID-19 crisis, we see an expected demand-shock pattern in the Energy sector, replicating prior findings. The Tariff crisis saw Financials emerge as the leading transmitter. During the SVB Banking crisis, the top three transmitting sectors were Industrials, Technology, and Healthcare. Healthcare appearing as a top transmitter is consistent with SVB's life sciences exposures; however, further investigation should be done as more post-crisis data become available. Utilities' defensive characteristics across all six crises make it the most effective sector choice for crisis-period portfolio protection, offering significantly better defensive exposure than Consumer Staples or Healthcare.

For financial regulators, our evidence supports scenario-specific stress testing. Pandemic scenarios should incorporate energy-sector transmission; monetary tightening scenarios should emphasize industrial and cyclical channels; and trade-policy scenarios should account for financial-sector transmission, given the unprecedented emergence of Financials as the leading transmitter during the 2025 tariff crisis.

This paper contributes methodologically through multiple validation methods and regime purity transparency. Compared to previous descriptive studies, our approach includes formal statistical tests along with explicit assessments of data quality. This transparency enables appropriately calibrated conclusions: the Inflation and Volmageddon results, with the highest regime purity, provide the strongest basis for inference; the Oil/China and COVID results offer robust supporting evidence; and the SVB Banking and Tariff results, though based on shorter crisis durations, produce distinctive patterns.

A few limitations should be noted. We focus our analysis on U.S. sector-level aggregates and the patterns may vary internationally or at the firm level. Our crisis-origin dependence finding is based

on six crisis episodes spanning a decade. While this provides substantially greater statistical power than narrower specifications, six crises still represent a limited sample for establishing universal laws of crisis transmission. The economic regime boundaries are defined ex-post by the authors. The primary challenge to regime purity arises from combining 200-day rolling windows with the short duration of many of the crises in our sample; our strongest findings rely on strong to moderate data quality, while the Banking and Tariff findings, though based on shorter crisis durations, produce sufficiently distinctive patterns to support inference.

Some extensions would strengthen these findings. Shorter rolling windows could improve regime purity for the Banking and Tariff crises, though at the cost of increased estimation noise. International replication applied to European and Asian sector markets would establish whether crisis-origin dependence is a U.S.-specific phenomenon or a global regularity. Finally, frequency-domain decomposition of quantile connectedness would separate short-term from long-term spillover horizons, informing whether crisis-specific strategies should be tactical or strategic.

In summary, no single network structure fits all crises. Each crisis origin produces a distinct set of sector transmitters, yet the cyclical-defensive divide persists universally. Crisis-specific risk management—rather than a one-size-fits-all approach—is both empirically justified and practically necessary.

References

Akhtaruzzaman, M., Boubaker, S., and Sensoy, A. (2021). Financial contagion during COVID–19 crisis. *Finance Research Letters*, 38:101604.

Ando, T., Greenwood-Nimmo, M., and Shin, Y. (2022). Quantile connectedness: modeling tail behavior in the topology of financial networks. *Management Science*, 68(4):2401–2431.

Antonakakis, N., Chatziantoniou, I., and Gabauer, D. (2020). Refined measures of dynamic connectedness based on time-varying parameter vector autoregressions. *Journal of Risk and Financial Management*, 13(4):84.

Baruník, J., Kočenda, E., and Vácha, L. (2016). Asymmetric connectedness on the US stock market: Bad and good volatility spillovers. *Journal of Financial Markets*, 27:55–78.

Bekaert, G., Ehrmann, M., Fratzscher, M., and Mehl, A. (2014). The global crisis and equity market contagion. *The Journal of Finance*, 69(6):2597–2649.

Bollen, N. P. and Whaley, R. E. (2004). Does net buying pressure affect the shape of implied volatility functions? *The Journal of Finance*, 59(2):711–753.

Chatziantoniou, I., Gabauer, D., and Stenfors, A. (2021). Interest rate swaps and the transmission mechanism of monetary policy: A quantile connectedness approach. *Economics Letters*, 204:109891.

Chicago Board Options Exchange (2019). The CBOE volatility index – VIX. White paper, Chicago Board Options Exchange.

Diebold, F. X. and Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *The Economic Journal*, 119(534):158–171.

Diebold, F. X. and Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1):57–66.

Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton, NJ.

Hull, J. C. (2018). *Options, Futures, and Other Derivatives*. Pearson, New York, NY, 10 edition.

Jawadi, F. and Sellami, M. (2022). On the effect of oil price in the context of Covid-19. *International Journal of Finance & Economics*, 27(4):3924–3933.

Kim, D.-J., Noh, E., and Choi, S.-Y. (2025). Quantile spillover effects and sector dynamics in US stock markets: Normal vs. extreme market conditions. *Finance Research Letters*, 83:107608.

Koenker, R. and Bassett Jr, G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50.

Laborda, R. and Olmo, J. (2021). Volatility spillover between economic sectors in financial crisis prediction: Evidence spanning the great financial crisis and Covid-19 pandemic. *Research in International Business and Finance*, 57:101402.

Ngene, G. M. (2021). What drives dynamic connectedness of the US equity sectors during different business cycles? *The North American Journal of Economics and Finance*, 58:101493.

Rehman, M. U., Zeitun, R., Nautiyal, N., Vo, X. V., and Kang, S. H. (2026). How do US sectoral markets connect in calm and crisis? a quantile-based network analysis. *Applied Economics*, 58(4):680–703.

Data Availability Statement

The data that support the findings of this study are available from FirstRate Data. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors with the permission of FirstRate Data.

Figure Legends

Figure 1: Total Connectedness Index: 2015–2025. *Note*: Total Connectedness Index at median quantile (q=0.5) with economic regime shading. Red regions indicate crisis periods, green indicates calm periods, and orange indicates transition periods. TCI measures system-wide spillovers in volatility changes across all sector ETFs.

Figure 2: Cyclical vs. Defensive Sector Spillovers: 2015–2025. *Note*: Mean NET spillovers for cyclical sectors (XLI, XLY, XLK) and defensive sectors (XLU, XLP, XLV) at median quantile (q=0.5). The persistent gap between blue and red lines visualizes the cyclical-defensive divide. The gap narrows during crises and expands during calm periods throughout the sample.

Figure 3: Sector NET Spillovers Across All Economic Regimes: 2015–2025. *Note*: Heatmap shows mean NET spillovers for all sectors across all fourteen regimes with available data at median quantile (q=0.5). Cyclical sectors generally transmit more than defensive sectors. Utilities exhibits safe-haven characteristics with negative NET across all regimes. Variation in transmission patterns illustrate crisis-origin dependence.

Figure 4: Cyclical-Defensive Divide by Crisis Type: Mean NET Spillovers Across Quantiles. *Note*: Shaded regions represent bootstrap 95% CIs (1,000 replications). The gap between the lines represents the cyclical-defensive divide.

Figure 5: Directed Spillover Networks Across Crisis Types. *Note*: Directed network graphs of pairwise volatility spillovers at the median quantile ($\tau = 0.5$) for each crisis regime. Node color indicates sector classification: blue = cyclical (XLI, XLY, XLK), red = defensive (XLU, XLP, XLV), grey = other (XLE, XLB, XLF). Node size is proportional to the absolute value of NET spillovers. Node border color indicates direction: blue border = net transmitter, red border = net receiver. Edge thickness and opacity reflect the magnitude of pairwise spillovers; only edges above the 65th percentile are shown, with minimum connectivity ensured for all nodes. Layout is fixed across panels for comparability.

Figure 6: Defensive Sector Comparison: NET Spillovers Across Regimes. *Note*: Mean NET spillovers at the median quantile ($\tau = 0.5$) for three defensive sectors across all regimes in chronological order. Utilities (XLU) is negative in all fourteen regimes (14/14), Consumer Staples (XLP) is negative in eleven (11/14), and Healthcare (XLV) is negative in only four (4/14). Only XLU qualifies as a pure safe haven.