

Article

Quantile Connectedness and Downside Risk in Portfolio Construction

Konoka Hamada ¹, Yuichiro Hamada ^{2,3} and Shigeyuki Hamori ^{3,4,*} ¹ Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA; konokh@uw.edu² Department of Business Administration, Momoyama Gakuin University, Manabino, Izumi 594-1198, Japan; y_hamada@andrew.ac.jp³ Faculty of Political Science and Economics, Yamato University, Asahigaoka-Cho, Suita 564-0082, Japan⁴ Graduate School of Economics, Kobe University, Rokkodai, Kobe 657-8501, Japan

* Correspondence: hamori@econ.kobe-u.ac.jp

Abstract

This paper examines whether quantile-based connectedness measures contain useful information for portfolio risk management. Using U.S. sector equity data, we estimate connectedness measures within a quantile vector autoregression framework and construct portfolios based on the cross-sectional distribution of net connectedness. In particular, sectors identified as extreme shock transmitters receive lower portfolio weights. Our results reveal substantial asymmetries across quantiles. Portfolios constructed using lower-tail connectedness measures exhibit smaller maximum drawdowns and lower expected shortfall relative to both equal-weight benchmarks and portfolios based on upper-tail connectedness. By contrast, median connectedness measures tend to provide more stable overall portfolio performance and lower turnover. The findings also suggest that the informational content of connectedness depends critically on the quantile considered. Lower-tail connectedness becomes particularly informative during crisis periods, suggesting that downside spillovers play an important role in portfolio resilience and systemic risk transmission. Overall, the results demonstrate that quantile connectedness measures provide economically meaningful information for downside risk management and offer a simple and transparent framework for incorporating systemic risk into portfolio construction.

Keywords: quantile connectedness; downside risk; portfolio construction; systemic risk; tail spillovers; risk management

JEL Classification: G11; G01; C58; C32



Academic Editor: Thanasis Stengos

Received: 31 May 2026

Revised: 24 June 2026

Accepted: 25 June 2026

Published: 1 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Financial markets no longer operate in isolation; they function as highly sensitive, interconnected webs where a local shock rapidly echoes across asset classes, sectors, and geographies. To track these domino effects, macro-finance has leaned heavily on the connectedness framework pioneered by [Diebold and Yilmaz \(2012, 2014\)](#), and comprehensively summarized in [Diebold and Yilmaz \(2015\)](#). While this framework has become a standard tool for monitoring systemic risk, the original specification remains fundamentally linear. Recent iterations have broken this mold. Researchers have introduced time-varying parameters ([Antonakakis et al., 2020](#); [Korobilis & Yilmaz, 2018](#)) and frequency decompositions ([Baruník & Křehlík, 2018](#)) to isolate short-term noise from long-term trends. More recently, the focus has shifted entirely to market extremes. Quantile-based connectedness measures

(Ando et al., 2022) suggest that the ties binding financial networks during calm periods fundamentally mutate when panic sets in.

This structural evolution highlights a long-standing disconnect in modern finance. For decades, asset allocation has relied on the traditional variance–covariance paradigm of Markowitz (1952). It is an elegant framework, but it harbors an important limitation: diversification is often less effective during periods of stress. In a crisis, historical correlations routinely collapse, assets move in lockstep, and the perceived diversification benefits diminish (Ang et al., 2006). This systematic breakdown has intensified the search for alternative risk metrics capable of capturing genuine systemic interactions (Acharya et al., 2017; Adrian & Brunnermeier, 2016).

Despite the rapid expansion of network models, their practical implications for portfolio allocation remain largely uncharted territory. Most literature treats connectedness as a passive monitoring indicator—an indicator to monitor rather than a portfolio allocation signal. While several studies have documented the role of connectedness in measuring systemic risk and financial market interdependence (Demirer et al., 2018), relatively few studies have explored how such information can be translated into practical portfolio allocation decisions.

This paper challenges that oversight. We re-examine financial connectedness not as an abstract statistic, but as an active, state-dependent variable for downside risk management. Our core premise is straightforward: a sector acting as a central clearinghouse for shocks during routine market regimes may not be the vector that triggers a systemic meltdown during a rout. By isolating lower-tail connectedness, we aim to identify the channels through which contagion is most likely to propagate during periods of market stress.

To test this hypothesis, we estimate quantile connectedness measures within a quantile vector autoregression (QVAR) framework using U.S. sector equity data. Rather than drowning the asset allocation process in complex, unstable optimization algorithms, we design a transparent, out-of-sample portfolio rule. We simply penalize and down-weight those sectors identified as major systemic shock transmitters.

Our empirical results reveal a stark asymmetry across market regimes. First, portfolios guided by lower-tail connectedness successfully mitigate downside risk, yielding smaller maximum drawdowns and superior expected shortfall relative to both equal-weight benchmarks and upper-tail models. Second, median connectedness acts as a stabilizing force during normal market conditions, providing steadier performance and lower portfolio turnover. Third, the economic value of connectedness-based portfolio allocation is highly state-dependent, with the benefits becoming most pronounced during periods of elevated systemic risk. The results remain qualitatively unchanged under alternative quantile specifications and are further supported by additional robustness analyses based on alternative weighting schemes and regime classifications. These findings offer two methodological contributions to the connectedness and portfolio allocation literature. First, we show that the informational content of financial connectedness differs substantially across quantiles, implying that lower-tail and upper-tail connectedness convey fundamentally different information for portfolio allocation. Second, we propose a transparent, optimization-free portfolio allocation framework that incorporates quantile-specific connectedness measures into practical investment decisions. Importantly, these methodological findings remain robust across alternative quantile specifications, weighting schemes, and regime classifications. From an economic perspective, our results suggest that periods of elevated lower-tail connectedness are associated with greater systemic vulnerability. Consequently, lower-tail connectedness may provide investors with useful information for improving portfolio resilience during periods of market stress.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the data and methodology. Section 4 presents the empirical results. Section 5 discusses the implications of the findings. Section 6 concludes.

2. Related Literature

This paper straddles three distinct strands of the literature, serving as a bridge between systemic risk measurement, downside portfolio theory, and network-driven asset allocation.

The first anchor of our study is the rapidly evolving literature on financial connectedness and systemic risk. While early systemic risk measures primarily focused on tail-risk co-movements (Acharya et al., 2017; Adrian & Brunnermeier, 2016), the connectedness framework pioneered by Diebold and Yilmaz (2012, 2014), and systematically synthesized in Diebold and Yilmaz (2015), transformed systemic risk monitoring by quantifying market spillovers through forecast error variance decompositions. Earlier studies on financial dependence also emphasized dynamic correlation structures through multivariate volatility models such as the dynamic conditional correlation (DCC) framework of Engle (2002). However, these approaches primarily focus on conditional covariance dynamics rather than directional spillovers and network topology. To capture the fluid nature of these dependencies, subsequent research integrated time-varying parameters into large-scale systems (Antonakakis et al., 2020; Korobilis & Yilmaz, 2018) and utilized frequency-domain decompositions to disentangle transient noise from persistent, structural trends (Baruník & Křehlík, 2018). More recently, the field has gravitated toward extreme market regimes. Quantile-based connectedness measures (Ando et al., 2022) have exposed the nonlinear realities and stark asymmetries inherent in financial networks during market stress, consistent with the broader literature on asymmetric financial dependence (Antonakakis et al., 2020; Patton, 2006; Umar et al., 2021). We build directly on this growing literature on nonlinear connectedness and tail spillovers. However, while existing studies primarily employ these metrics as macroprudential monitoring tools, we explicitly isolate the cross-sectional differences between lower-tail and upper-tail networks to examine whether they contain distinct asset allocation value.

Second, this paper intersects with the literature on portfolio choice under asymmetric risk conditions. The classical mean–variance optimization of Markowitz (1952) treats upside volatility and downside loss symmetrically. Yet, empirical realities show that variance fails as a risk metric when return distributions exhibit fat tails and heavy skewness. Consequently, asset management has increasingly adopted downside risk metrics, such as expected shortfall and maximum drawdown, to evaluate performance when markets fracture (Ang et al., 2006). Our study inserts network architecture into this space. We propose that lower-tail connectedness acts as a real-time map of downside systemic fragility, offering a structural alternative to purely statistical tail-risk measures.

Finally, we contribute to an emerging domain that operationalizes network metrics for practical portfolio construction. Moving away from treating connectedness as a passive tracking dashboard, pioneering work by Demiret et al. (2018) and Ji et al. (2019) demonstrates that network-based variables capture structural features of market stress that traditional covariance matrices routinely miss. This perspective also relates to the broader literature questioning the practical superiority of complex portfolio optimization over simple allocation rules. In particular, DeMiguel et al. (2009) show that naive equal-weight portfolios often outperform theoretically optimal strategies in out-of-sample settings due to estimation error and parameter instability. Our approach shares a similar motivation by proposing a transparent and robust allocation framework that avoids reliance on unstable covariance matrix estimation.

Our paper carves out a distinct niche within this space through three key innovations. Rather than treating the market network as a monolith, we demonstrate that its informational value is highly localized within specific quantiles, exposing a massive structural asymmetry between bull- and bear-market spillovers. Second, we provide evidence that lower-tail network signals offer uniquely potent protection for portfolio resilience when market liquidity evaporates. Lastly, we provide institutional investors with a transparent, optimization-free allocation rule—converting complex, high-dimensional network mechanics into an intuitive, downside-hedged investment strategy.

3. Empirical Strategy

3.1. Data and Asset Universe

Our empirical analysis focuses on eight major U.S. equity sector exchange-traded funds (ETFs) from the Select Sector SPDR family: Energy (XLE), Industrials (XLI), Consumer Discretionary (XLY), Consumer Staples (XLP), Health Care (XLV), Financials (XLF), Technology (XLK), and Utilities (XLU). Daily adjusted closing prices are obtained from Yahoo Finance and span the period from January 2000 to December 2025. Returns are computed as log differences in adjusted prices.

The eight Select Sector SPDR ETFs were selected because they represent the major sectors of the U.S. economy and collectively provide broad coverage of the U.S. equity market. These sector ETFs are among the most actively traded and widely studied sectoral investment vehicles, making them particularly suitable for examining the transmission of shocks across different segments of the economy. The objective of the present study is not to evaluate sector-specific alpha generation or benchmark-relative performance, but rather to investigate how quantile-specific connectedness information can be incorporated into portfolio allocation decisions.

3.2. Quantile Connectedness Measurement

To measure connectedness across sectors under different market conditions, we employ a quantile vector autoregression (QVAR) framework following Ando et al. (2022). This approach builds on the broader literature on multivariate regression quantiles and tail dependence modeling developed by White et al. (2015), which extends conventional VAR analysis beyond mean-based dependence structures. Unlike standard VAR models, the QVAR approach allows connectedness to vary across different parts of the conditional return distribution. This framework is particularly useful for capturing asymmetric spillovers during extreme market conditions.

Let $\mathbf{r}_t = (r_{1,t}, \dots, r_{8,t})'$ denote the vector of sector returns. For a given quantile level $\tau \in (0, 1)$, the QVAR model is specified as

$$Q_\tau(\mathbf{r}_t \mid \mathcal{F}_{t-1}) = \sum_{k=1}^p A_k^{(\tau)} \mathbf{r}_{t-k}, \quad (1)$$

where $Q_\tau(\cdot \mid \mathcal{F}_{t-1})$ denotes the conditional quantile function and $A_k^{(\tau)}$ represents quantile-specific coefficient matrices.

We estimate the model using a rolling window of $W = 250$ daily observations and set the lag order to $p = 1$. We follow the quantile connectedness framework of Ando et al. (2022), which extends the Diebold and Yilmaz (2012, 2014) methodology to quantile VAR systems.

The main analysis focuses on three representative quantiles, $\tau = 0.05$, $\tau = 0.50$, and $\tau = 0.95$, corresponding to lower-tail, median, and upper-tail market conditions, respectively. The lower-tail quantile captures downside spillovers during periods of market stress,

whereas the upper-tail quantile reflects spillovers under exceptionally favorable market conditions. As a robustness check, we additionally estimate connectedness measures at $\tau = 0.10$ and $\tau = 0.90$. These additional specifications allow us to assess whether the empirical results depend on the particular choice of tail quantiles.

For a given forecast horizon H , the generalized forecast error variance decomposition element $\tilde{\theta}_{ij,t}^{(H,\tau)}$ measures the contribution of shocks in sector j to the forecast error variance of sector i at quantile τ . The forecast horizon is set to $H = 10$.

The net connectedness measure ($C_{i,t}^{net,\tau}$) is defined as

$$C_{i,t}^{net,\tau} = C_{i,t}^{\rightarrow,\tau} - C_{i,t}^{\leftarrow,\tau}, \quad (2)$$

where $C_{i,t}^{\rightarrow,\tau}$ and $C_{i,t}^{\leftarrow,\tau}$ denote directional connectedness transmitted to and received from other sectors, respectively. Positive values of $C_{i,t}^{net,\tau}$ indicate that sector i acts as a net transmitter of shocks at quantile τ , while negative values indicate a net receiver.

As a summary measure of system-wide spillovers, we compute the quantile-specific total connectedness index (TCI),

$$TCI_t^{(\tau)} = \frac{1}{8} \sum_{i=1}^8 \sum_{\substack{j=1 \\ j \neq i}}^8 \tilde{\theta}_{ij,t}^{(H,\tau)}, \quad (3)$$

which captures the overall intensity of connectedness across different parts of the return distribution.

3.3. Portfolio Construction

To incorporate connectedness into portfolio allocation, we construct a simple quantile-based portfolio that adjusts weights according to the cross-sectional distribution of net connectedness.

In the baseline specification, portfolio weights are constructed using lower-tail connectedness measures estimated at $\tau = 0.05$, reflecting the primary focus on downside risk transmission.

At each point in time, sectors are ranked based on their net connectedness $C_{i,t}^{net,\tau}$. We identify sectors in the upper quantile of the cross-sectional distribution as extreme transmitters of shocks. These sectors are assigned lower portfolio weights, while the remaining sectors receive relatively higher weights.

Formally, portfolio weights ($w_{i,t}$) are defined as

$$w_{i,t} \propto \begin{cases} \omega_L, & \text{if } C_{i,t}^{net,\tau} > Q_{0.75,t}^{(\tau)}, \\ \omega_H, & \text{otherwise,} \end{cases} \quad (4)$$

where $Q_{0.75,t}^{(\tau)}$ denotes the 75th percentile of net connectedness across sectors at time t , and $\omega_L < \omega_H$ are fixed constants. In the baseline specification, we set $\omega_L = 0.2$ and $\omega_H = 1.0$. To evaluate the sensitivity of the portfolio allocation rule, we additionally consider alternative weighting schemes in which the low-weight parameter is increased to $\omega_L = 0.3$ and $\omega_L = 0.4$, while maintaining $\omega_H = 1.0$. The results of this robustness analysis are reported in Appendix B.

Final portfolio weights are obtained by normalization,

$$w_{i,t} = \frac{w_{i,t}}{\sum_{j=1}^8 w_{j,t}}, \quad (5)$$

ensuring full investment of the portfolio.

As shown in Equation (4), the strategy penalizes strong net transmitters, while Equation (5) ensures that portfolio weights sum to one.

This rule can be interpreted as a risk-tilting strategy that reduces exposure to sectors that act as strong transmitters of shocks while maintaining diversification across the remaining sectors. The quantile-based approach provides a transparent and robust way to identify extreme spillover contributors without requiring complex optimization techniques or additional tuning parameters.

Portfolio weights are updated at a monthly frequency using only information available at time t , ensuring a strictly out-of-sample implementation.

We also report portfolio turnover to assess the stability of the proposed strategy and to provide an indication of potential transaction costs in practical applications.

3.4. Regime Classification

To examine whether the effectiveness of connectedness-based allocation depends on market conditions, we classify the sample into high- and low-connectedness regimes.

Specifically, we define the high-connectedness regime as periods in which the lower-tail total connectedness index, $TCI_t^{(0.05)}$, exceeds its sample median, and the low-connectedness regime otherwise. This classification allows us to evaluate portfolio performance separately across different systemic risk environments.

We use the sample median of $TCI_t^{(0.05)}$ as the baseline classification threshold. To address potential concerns regarding the use of full-sample information, we additionally employ an expanding-window median threshold that relies only on information available up to each point in time. The results of this alternative regime classification are reported in Appendix C.

This regime-based approach is motivated by the observation that financial connectedness tends to increase during episodes of financial turmoil, potentially altering the effectiveness of diversification strategies.

3.5. Downside Risk Measures

Given our focus on systemic risk exposure, we complement standard performance metrics with measures of downside risk.

In addition to the Sharpe ratio, we consider the following indicators:

- Maximum drawdown, capturing the largest peak-to-trough loss;
- Sortino ratio, focusing on downside volatility;
- Expected shortfall (ES), measuring tail risk beyond a specified quantile;
- Calmar ratio, evaluating portfolio returns relative to maximum drawdown.

Expected shortfall is computed at the 5% level as the average return conditional on returns falling below the corresponding Value-at-Risk threshold.

These measures allow us to assess whether connectedness-based allocation enhances portfolio resilience to systemic shocks, particularly during adverse market conditions.

Such risk indicators are particularly relevant in environments characterized by elevated systemic risk, as extreme losses tend to cluster during periods of elevated connectedness.

This approach allows us to directly link portfolio performance to quantile-specific systemic risk conditions.

3.6. Benchmark Portfolios and Evaluation

We evaluate the performance of the connectedness-based portfolio against a standard equal-weight benchmark. Performance is assessed using both standard and downside-focused metrics. In addition, we report results separately across high-

and low-connectedness regimes to examine the state-dependent effectiveness of the proposed strategy.

We also conduct statistical tests of performance differences to assess whether the observed improvements are economically and statistically significant.

4. Empirical Results

4.1. Quantile Connectedness Dynamics

Figure 1 presents the time-varying total connectedness indices estimated from the quantile VAR framework at $\tau = 0.05$, $\tau = 0.50$, and $\tau = 0.95$. Several important features emerge from the figure.




Figure 1. Time-varying total connectedness indices across quantiles. The figure reports the dynamics of total connectedness estimated from quantile VAR models at the lower-tail ($\tau = 0.05$), median ($\tau = 0.50$), and upper-tail ($\tau = 0.95$) quantiles. Vertical dotted lines indicate major financial events, including the Global Financial Crisis (2008–2009) and the COVID-19 pandemic (2020), and are included to facilitate interpretation of the dynamics.

First, connectedness exhibits substantial time variation across all quantiles, with pronounced spikes during major episodes of financial stress, including the Global Financial Crisis (2008–2009) and the COVID-19 shock (2020). This finding is consistent with the existing literature emphasizing the countercyclical nature of financial connectedness and systemic risk.

Second, the dynamics of connectedness differ substantially across quantiles. In particular, lower-tail connectedness ($\tau = 0.05$) displays sharper and more persistent increases during crisis periods relative to median and upper-tail connectedness. This result suggests that downside spillovers intensify considerably under adverse market conditions, highlighting the asymmetric nature of systemic risk transmission.

By contrast, upper-tail connectedness ($\tau = 0.95$) exhibits smoother dynamics and weaker responses during crisis episodes. This finding indicates that spillovers associated with exceptionally favorable market conditions differ fundamentally from downside contagion mechanisms.

Overall, Figure 1 demonstrates that the informational content of connectedness depends critically on the quantile considered, motivating the use of quantile-specific connectedness measures for portfolio allocation and downside risk management.

4.2. Portfolio Performance

Table 1 reports the performance of portfolios constructed using connectedness measures estimated at alternative quantiles, together with the equal-weight benchmark.

Several important findings emerge. First, portfolios based on lower-tail connectedness ($\tau = 0.05$) provide the strongest downside risk protection. The lower-tail portfolio exhibits the smallest maximum drawdown among all strategies, indicating improved resilience during periods of severe market stress. In addition, the lower-tail portfolio achieves one of the lowest expected shortfall values and among the highest Calmar ratios, suggesting that downside connectedness contains economically meaningful information about tail risk exposure and drawdown management.

Second, portfolios based on median connectedness ($\tau = 0.50$) deliver the strongest overall risk-adjusted performance. In particular, the median-connectedness portfolio achieves the highest Sharpe ratio while simultaneously exhibiting the lowest portfolio turnover. These results suggest that median connectedness captures relatively stable systemic relationships that are useful for portfolio allocation under normal market conditions.

Third, portfolios constructed from upper-tail connectedness ($\tau = 0.95$) perform relatively poorly across most evaluation criteria. In particular, the upper-tail portfolio exhibits larger drawdowns and weaker risk-adjusted performance relative to both the lower-tail and median-connectedness portfolios. This finding indicates that upper-tail spillovers contain less useful information for downside risk management.

Table 1. Portfolio Performance across Quantile Connectedness Strategies.

	Q05	Q10	Q50	Q90	Q95	EW
Annualized Return	0.0444	0.0427	0.0455	0.0445	0.0417	0.0451
Annualized Std. Dev.	0.1799	0.1781	0.1778	0.1785	0.1789	0.1806
Sharpe Ratio	0.2468	0.2400	0.2561	0.2492	0.2330	0.2499
Maximum Drawdown	0.5762	0.5846	0.5916	0.5864	0.6013	0.5936
Sortino Ratio	0.0287	0.0280	0.0294	0.0288	0.0274	0.0290
Expected Shortfall (5%)	−0.0278	−0.0276	−0.0276	−0.0277	−0.0278	−0.0281
Calmar Ratio	0.0771	0.0731	0.0770	0.0758	0.0693	0.0760
Turnover	0.0452	0.0313	0.0123	0.0271	0.0383	N/A

Notes: Q05, Q10, Q50, Q90, and Q95 denote portfolios constructed using connectedness measures estimated at the corresponding quantiles ($\tau = 0.05$, $\tau = 0.10$, $\tau = 0.50$, $\tau = 0.90$, and $\tau = 0.95$, respectively). EW denotes the equal-weight benchmark. Expected shortfall is computed at the 5% level. The Calmar ratio measures annualized returns relative to maximum drawdown. Turnover denotes the average absolute portfolio weight adjustment across periods. N/A indicates that turnover is not applicable to the equal-weight benchmark because its portfolio weights remain constant throughout the sample period.

Fourth, the portfolio performance results remain remarkably stable across alternative quantile specifications. Portfolios constructed using $\tau = 0.10$ and $\tau = 0.90$ generate performance measures that are highly similar to those obtained under $\tau = 0.05$ and $\tau = 0.95$, respectively. In particular, annualized returns, Sharpe ratios, Sortino ratios, and Calmar ratios exhibit only minor variation across neighboring quantiles. This finding suggests that the main conclusions are not driven by the specific choice of tail quantiles.

Taken together, these results reveal substantial asymmetries across quantiles. Lower-tail connectedness appears particularly informative for downside protection, whereas upper-tail connectedness does not provide comparable portfolio benefits.

An additional practical consideration is portfolio turnover. Across all specifications, turnover remains relatively modest, ranging from approximately 1.2% to 4.5% per rebalancing period. The median-connectedness portfolio exhibits the lowest turnover, indicating a particularly stable allocation rule. More broadly, the turnover levels suggest that transaction costs are unlikely to overturn the main conclusions, supporting the practical feasibility of the proposed strategy.

Figure 2 presents the cumulative portfolio performance across strategies. The figure confirms that the lower-tail connectedness strategy experiences smaller drawdowns during major crisis periods, whereas the median-connectedness strategy exhibits relatively stable long-run performance. These patterns are broadly consistent with the results reported in Table 1 and remain robust across alternative quantile specifications.




Figure 2. Portfolio performance across quantile connectedness strategies. The figure reports cumulative portfolio values for portfolios constructed using lower-tail ($\tau = 0.05$), median ($\tau = 0.50$), and upper-tail ($\tau = 0.95$) connectedness measures, together with the equal-weight benchmark.

4.3. Drawdown Analysis

Figure 3 compares drawdown dynamics across portfolio strategies. The lower-tail connectedness portfolio consistently exhibits smaller drawdowns during periods of elevated systemic risk, particularly during the Global Financial Crisis and the COVID-19 shock. By contrast, the upper-tail connectedness portfolio experiences deeper and more persistent losses.




Figure 3. Drawdown comparison across quantile connectedness strategies. The lower-tail connectedness strategy ($\tau = 0.05$) exhibits smaller drawdowns during major market stress episodes, indicating improved downside risk protection.

These findings reinforce the view that lower-tail spillovers contain particularly useful information for downside risk management. The results also indicate that the informational value of connectedness is highly asymmetric across quantiles.

4.4. Regime-Dependent Performance

To further examine whether the effectiveness of connectedness-based allocation depends on systemic risk conditions, we evaluate portfolio performance separately across high- and low-connectedness regimes based on the lower-tail total connectedness index.

Table 2 reports the corresponding results.

Table 2. Portfolio Performance across Connectedness Regimes.

	Q05	Q50	Q95	EW
High-Connectedness Regime				
Annualized Return	0.0767	0.0785	0.0728	0.0784
Annualized Std. Dev.	0.2086	0.2054	0.2065	0.2087
Sharpe Ratio	0.3678	0.3824	0.3524	0.3759
Low-Connectedness Regime				
Annualized Return	0.0130	0.0138	0.0117	0.0127
Annualized Std. Dev.	0.1462	0.1456	0.1467	0.1477
Sharpe Ratio	0.0889	0.0945	0.0799	0.0861

Notes: Q05, Q50, and Q95 denote portfolios constructed using connectedness measures estimated at $\tau = 0.05$, $\tau = 0.50$, and $\tau = 0.95$, respectively. EW denotes the equal-weight benchmark. The table reports portfolio performance separately across high- and low-connectedness regimes defined using the lower-tail total connectedness index, $TCI_t^{(0.05)}$. The baseline classification uses the sample median as the threshold. Results based on an expanding-window median threshold are reported in Appendix C.

Several findings are noteworthy. First, the benefits of connectedness-based allocation become substantially stronger during high-connectedness regimes. In these periods, both the lower-tail and median-connectedness portfolios outperform the equal-weight benchmark in terms of risk-adjusted performance.

Second, the lower-tail portfolio continues to provide relatively strong downside protection during high-connectedness periods, supporting the interpretation that downside spillovers become particularly important during episodes of elevated systemic risk.

Third, performance differences across strategies become considerably smaller during low-connectedness regimes. This result suggests that the usefulness of connectedness-based allocation is state-dependent and concentrated primarily during adverse market environments.

Overall, the regime analysis indicates that connectedness measures are most informative precisely when systemic risk becomes elevated and traditional diversification benefits weaken.

These conclusions remain broadly unchanged when regimes are alternatively classified using an expanding-window median threshold, as reported in Appendix C. Although the exact regime composition varies across classification schemes, the overall evidence continues to support the state-dependent nature of connectedness-based portfolio allocation.

5. Discussion

The empirical findings uncovered in this study offer a fresh perspective on how non-linear network architecture interacts with practical portfolio risk management.

Our core insight centers on a stark reality: financial networks are fundamentally asymmetric, and treating them as a homogeneous monolith obscures the most critical data signals. Isolating lower-tail connectedness ($\tau = 0.05$) yields a profoundly more potent signal for capital preservation than tracking its upper-tail counterpart ($\tau = 0.95$). Portfolios anchored in lower-tail metrics successfully insulate wealth—driving down maximum drawdowns and curbing expected shortfall—while upper-tail metrics fail to offer any meaningful downside protection.

This divergence is rooted in the behavioral economics of market stress. Lower-tail connectedness captures the network when the plumbing fractures: it maps the system during waves of panic selling, forced institutional deleveraging, margin calls, and sudden liquidity freezes. During these structural dislocations, sectors positioned as net shock transmitters act as systemic contagion vectors. Systematically pruning exposure to these specific nodes acts as a structural antidote, reinforcing portfolio resilience precisely when market correlations begin to step in dangerous unison.

Conversely, upper-tail connectedness ($\tau = 0.95$) reflects a completely different psychological regime. These spillovers capture periods of synchronized market exuberance and rising tides, where broad-based optimism lifts all sectors simultaneously. While these upside dynamics are valuable for tracking momentum or market rallies, they are fundamentally uninformative for capital preservation strategies; they tell us how the market runs, not how it breaks.

A particularly compelling nuance emerges when we look at the center of the distribution. While the lower tail dominates in mitigating catastrophe, portfolios built around median connectedness ($\tau = 0.50$) secure the highest overall Sharpe ratio and the lowest turnover. This suggests that the median network captures the persistent, fundamental economic linkages between sectors during routine market operations. For institutional managers, this uncovers a distinct tactical trade-off: median connectedness serves as an excellent guide for long-term, fair-weather compounding, whereas lower-tail connectedness is an indispensable tool for active crisis management.

Our regime analysis underscores this state-dependent behavior. The outperformance of connectedness-based asset allocation scales up dramatically as systemic risk intensifies—the exact moments when classic asset diversification leaves investors exposed. In these high-connectedness regimes, lower-tail spillovers sharpen, proving that the signal becomes loudest precisely when the cost of ignoring it is highest.

Ultimately, these dynamics demonstrate that connectedness cannot be compressed into a single, average indicator. Different quantiles map entirely different dimensions of market psychology and structural fragility. Lower-tail connectedness serves as a pure metric of systemic vulnerability and downside contagion, whereas upper-tail connectedness tracks macro-momentum.

From an operational standpoint, our allocation framework offers immediate institutional utility. The strategy sidesteps the standard pitfalls of modern portfolio optimization—namely, extreme sensitivity to input estimation errors and parameter instability. It is transparent, computationally lean, and requires no complex matrix inversions. Because it operates strictly in a rolling, out-of-sample framework using real-time information, it offers a practical, executable playbook for risk managers tasked with safeguarding capital under fire.

Inevitably, certain boundaries define the scope of these insights. Our empirical exploration focuses tightly on U.S. sector equities; these non-linear network dynamics may express themselves differently across international equity matrices or distinct asset classes. Additionally, our baseline rules do not explicitly friction-model transaction costs, though the low turnover of the median strategy suggests inherent structural stability. Finally, our QVAR approach captures reduced-form dependencies; it maps the velocity and path of contagion rather than explicitly isolating the underlying structural or regulatory triggers.

These limitations naturally pave the way for future research. A logical next step is to expand the quantile matrix into cross-asset networks, analyzing the hidden systemic channels running between sovereign bonds, commodities, cryptocurrencies, and foreign exchange markets. Furthermore, blending these quantile connectedness signals with dynamic machine learning architectures or robust optimization techniques could refine the

allocation rules. More broadly, pinning down the exact macroeconomic imbalances and behavioral triggers that cause lower-tail connectedness to explode remains a vital frontier for understanding systemic risk.

6. Conclusions

This paper evaluates the role of quantile connectedness in portfolio construction, utilizing U.S. sector equity portfolios as our empirical testing ground. Operating within a quantile vector autoregression framework, we mapped the shifting landscape of financial spillovers across lower-tail, median, and upper-tail market regimes. By translating these cross-sectional network dynamics into a set of direct asset allocation rules, we moved beyond mere risk measurement into active portfolio mitigation.

Our empirical findings reveal that financial connectedness is far from uniform; it exhibits massive structural heterogeneity across quantiles. Lower-tail connectedness, in particular, undergoes sharp, persistent spikes during major market dislocations, most notably during the Global Financial Crisis and the COVID-19 shock. This asymmetry underscores a critical reality of systemic risk: the network structure of financial networks tightens drastically during a severe market downturn, making downside spillovers the primary vector of contagion.

Crucially, the economic value of these network metrics for asset allocation hinges entirely on the specific quantile an investor tracks. Portfolios insulated by lower-tail connectedness metrics consistently deliver superior downside protection, yielding smaller maximum drawdowns and lower expected shortfall than both equal-weight baselines and upper-tail counterstrategies. Conversely, median connectedness functions as a stabilizer for ordinary market regimes, optimizing overall stability and minimizing portfolio turnover, while upper-tail connectedness offers virtually no actionable utility for risk mitigation.

Furthermore, the efficacy of this connectedness-based framework is deeply state-dependent. The outperformance of our proposed strategy scales up precisely when systemic risk peaks—the exact moments when market-wide spillovers accelerate and traditional diversification strategies fail. This confirms that network metrics provide the steepest diversification benefits exactly when downside protection matters most.

These results offer a vital refinement to the literature. While our findings align with the foundational work of [Diebold and Yilmaz \(2012, 2014\)](#) and [Demirer et al. \(2018\)](#) in confirming that asset interconnections intensify during crises, we extend this narrative by demonstrating that the informational value of these networks is profoundly asymmetric. Lower-tail connectedness holds a distinct, economically meaningful signal for building portfolio resilience that the rest of the distribution simply cannot replicate.

This calls for a more nuanced interpretation of network metrics in asset management. Connectedness should not be misconstrued as a direct predictor of asset returns. Instead, it operates as a clear proxy for structural fragility and systemic vulnerability. Viewed through this lens, quantile-specific networks offer institutional investors a robust diagnostic tool to pinpoint the specific economic sectors poised to amplify market panics.

From a practical perspective, our allocation framework bypasses the computational fragility of standard asset optimization. It is transparent, structurally simple, and executes seamlessly in an out-of-sample environment without relying on complex, unstable matrix inversions. This implementation simplicity makes it highly viable for institutional asset managers who require reliable drawdown protection during periods of market stress.

An additional contribution of this study is the robustness of the main findings across alternative empirical specifications. The superiority of lower-tail connectedness for downside risk management remains intact when connectedness is estimated at neighboring tail quantiles. Likewise, the portfolio performance results are largely unaffected by alter-

native weighting schemes and alternative regime classification procedures based on an expanding-window median threshold. These robustness checks strengthen the conclusion that the economic value of lower-tail connectedness is not an artifact of a particular model specification, but rather reflects a stable feature of systemic risk transmission in financial markets.

Naturally, certain boundaries pin this study down and chart the course for future inquiry. First, our empirical focus rests exclusively on U.S. sector equities; these network dynamics may shift when applied to international markets or distinct asset classes. Second, our baseline strategy does not explicitly model transaction costs or leverage constraints. Finally, our structural framework captures reduced-form dependencies rather than identifying explicit, structural causal channels of contagion.

To build on these foundations, future research should expand these quantile networks to encompass cross-asset dynamics—integrating sovereign bonds, commodities, cryptocurrencies, and foreign exchange markets into a broader global matrix. Another promising avenue lies in marrying quantile connectedness signals with dynamic machine learning algorithms or portfolio optimization techniques. More fundamentally, uncovering the underlying macroeconomic drivers and behavioral biases that trigger these lower-tail network explosions remains a vital frontier in the study of systemic risk.

Ultimately, this paper demonstrates that ignoring the specific quantile of a financial network may lead investors to overlook valuable information about systemic risk transmission. Lower-tail spillovers contain economically meaningful signals for enhancing portfolio resilience and provide a practical framework for managing downside risk during periods of market stress.

Author Contributions: Conceptualization, Y.H. and S.H.; methodology, K.H. and S.H.; formal analysis, K.H. and S.H.; investigation, Y.H. and S.H.; writing—original draft preparation, K.H. and S.H.; writing—review and editing, Y.H. and S.H.; supervision, Y.H. and S.H.; project administration, S.H.; funding acquisition, S.H. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by JSPS KAKENHI, Grant Number JP25K05043.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data analyzed in this study are publicly available from Yahoo Finance at <https://finance.yahoo.com/>.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Average Portfolio Weights

This appendix reports the average portfolio weights assigned to each sector under the alternative quantile connectedness strategies. The table illustrates how the proposed allocation rule systematically reduces exposure to sectors identified as major net transmitters of shocks while increasing exposure to sectors with relatively lower systemic influence.

Table A1. Average Portfolio Weights across Quantile Connectedness Strategies.

Sector	Q05	Q50	Q95	EW
XLE	0.1519	0.1559	0.1523	0.1250
XLI	0.0712	0.0382	0.0624	0.1250
XLY	0.1062	0.1052	0.1082	0.1250

Table A1. *Cont.*

Sector	Q05	Q50	Q95	EW
XLP	0.1451	0.1535	0.1461	0.1250
XLV	0.1316	0.1414	0.1407	0.1250
XLF	0.1126	0.1058	0.1104	0.1250
XLK	0.1284	0.1437	0.1243	0.1250
XLU	0.1530	0.1562	0.1558	0.1250

Notes: Q05, Q50, and Q95 denote portfolios constructed using connectedness measures estimated at $\tau = 0.05$, $\tau = 0.50$, and $\tau = 0.95$, respectively. EW denotes the equal-weight benchmark. Reported values represent average portfolio weights over the full sample period.

Appendix B. Robustness to Alternative Weighting Schemes

This appendix examines whether the portfolio performance results are sensitive to the choice of the weighting rule. In the baseline specification, sectors identified as extreme net transmitters receive a weight of 0.2, while the remaining sectors receive a weight of 1.0 prior to normalization. We consider two alternative specifications in which the low-weight parameter is increased to 0.3 and 0.4, respectively, while keeping the high-weight parameter fixed at 1.0.

Table A2. Performance Robustness under Alternative Weighting Schemes.

Weight Rule	Annualized Return	Sharpe Ratio	Calmar Ratio	Turnover
0.2/1.0	0.0444	0.2468	0.0771	0.0452
0.3/1.0	0.0445	0.2475	0.0770	0.0384
0.4/1.0	0.0447	0.2481	0.0768	0.0319

Notes: The table reports portfolio performance under alternative weighting schemes. The first number denotes the weight assigned to sectors classified as extreme net transmitters, while the second number denotes the weight assigned to all remaining sectors prior to normalization. Reported results are based on the lower-tail connectedness portfolio ($\tau = 0.05$).

Appendix C. Robustness to Alternative Regime Classification

This appendix examines the robustness of the regime-dependent results to an alternative classification procedure. While the baseline analysis defines high- and low-connectedness regimes using the full-sample median of the lower-tail total connectedness index, this approach may raise concerns regarding the use of future information. To address this issue, we additionally classify regimes using an expanding-window median threshold that relies only on information available up to each point in time.

Table A3. Robustness to Alternative Regime Classification.

	Q05	Q10	Q50	Q90	Q95	EW
High-Connectedness Regime						
Full Median	0.0767	0.0737	0.0785	0.0732	0.0727	0.0784
Expanding Median	0.0508	0.0500	0.0523	0.0508	0.0496	0.0508
Low-Connectedness Regime						
Full Median	0.0131	0.0127	0.0135	0.0165	0.0115	0.0129
Expanding Median	0.0361	0.0334	0.0367	0.0362	0.0315	0.0377

Notes: Q05, Q10, Q50, Q90, and Q95 denote portfolios constructed using connectedness measures estimated at the corresponding quantiles. The table reports annualized portfolio returns across high- and low-connectedness regimes under alternative classification schemes. The baseline classification uses the full-sample median of the lower-tail total connectedness index, while the alternative classification uses an expanding-window median threshold.

The results remain broadly consistent across the two classification schemes. Although the exact composition of the regimes differs when an expanding-window threshold is employed, the qualitative conclusions remain unchanged. In particular, connectedness-based portfolio performance continues to exhibit strong state dependence, with materially different outcomes across high- and low-connectedness environments. Moreover, the relative ranking of the portfolio strategies remains largely intact across the alternative regime classifications. These findings suggest that the regime-dependent results reported in the main text are not driven by the use of full-sample information.

References

- Acharya, V. V., Pedersen, L. H., Philippon, T., & Richardson, M. (2017). Measuring systemic risk. *Review of Financial Studies*, 30(1), 2–47.
- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741. [\[CrossRef\]](#)
- Ando, T., Greenwood-Nimmo, M., & Shin, Y. (2022). Quantile connectedness: Modeling tail behavior in the topology of financial networks. *Management Science*, 68(4), 2401–2431. [\[CrossRef\]](#)
- Ang, A., Chen, J., & Xing, Y. (2006). Downside risk. *Review of Financial Studies*, 19(4), 1191–1239. [\[CrossRef\]](#)
- Antonakakis, N., Chatziantoniou, I., & Gabauer, D. (2020). Refined measures of dynamic connectedness based on TVP-VAR. *Journal of Risk and Financial Management*, 13(4), 84.
- Baruník, J., & Křehlík, T. (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, 16(2), 271–296. [\[CrossRef\]](#)
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22(5), 1915–1953.
- Demirer, M., Diebold, F. X., Liu, L., & Yilmaz, K. (2018). Estimating global bank network connectedness. *Journal of Applied Econometrics*, 33(1), 1–15.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66. [\[CrossRef\]](#)
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134. [\[CrossRef\]](#)
- Diebold, F. X., & Yilmaz, K. (2015). *Financial and macroeconomic connectedness: A network approach to measurement and monitoring*. Oxford University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. *Journal of Business & Economic Statistics*, 20(3), 339–350. [\[CrossRef\]](#)
- Ji, Q., Bouri, E., Lau, C. K. M., & Roubaud, D. (2019). Dynamic connectedness and integration in cryptocurrency markets. *International Review of Financial Analysis*, 63, 257–272. [\[CrossRef\]](#)
- Korobilis, D., & Yilmaz, K. (2018). Measuring dynamic connectedness with large Bayesian VAR models. In *Koç University-TUSIAD economic research forum working papers 1802*. Koç University.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91. [\[CrossRef\]](#)
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556. [\[CrossRef\]](#)
- Umar, Z., Gubareva, M., Teplova, T., & Tran, D. K. (2022). COVID-19 impact on NFTs and major asset classes interrelations: Insights from the wavelet coherence analysis. *Finance Research Letters*, 47, 102725. [\[CrossRef\]](#)
- White, H., Kim, T.-H., & Manganelli, S. (2015). VAR for VaR: Measuring tail dependence using multivariate regression quantiles. *Journal of Econometrics*, 187(1), 169–188. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.