



An early-warning risk signals framework to capture systematic risk in financial markets

Vito Ciciretti, Monomita Nandy, Alberto Pallotta, Suman Lodh, P. K. Senyo & Jekaterina Kartasova

To cite this article: Vito Ciciretti, Monomita Nandy, Alberto Pallotta, Suman Lodh, P. K. Senyo & Jekaterina Kartasova (14 Apr 2025): An early-warning risk signals framework to capture systematic risk in financial markets, Quantitative Finance, DOI: [10.1080/14697688.2025.2482637](https://doi.org/10.1080/14697688.2025.2482637)

To link to this article: <https://doi.org/10.1080/14697688.2025.2482637>



© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 14 Apr 2025.



[Submit your article to this journal](#)



Article views: 206








[View related articles](#)



[View Crossmark data](#)

An early-warning risk signals framework to capture systematic risk in financial markets

VITO CICIRETTI [†], MONOMITA NANDY ^{‡*}, ALBERTO PALLOTTA [§],
SUMAN LODH [¶], P. K. SENYO ^{||} and JEKATERINA KARTASOVA ^{††}

[†]Independent Researcher, Berlin, Germany

[‡]Brunel Business School, Brunel University London, Middlesex, UK

[§]Accounting, Finance & Economics, Middlesex University, London, UK

[¶]Kingston Business School, Kingston University London, Surrey, UK

^{||}Accounting, Finance & Economics, University of Southampton, Southampton, UK

^{††}Middlesex University Business School, Middlesex University, London, UK

(Received 11 March 2023; accepted 14 March 2025)

Despite extensive research on the relationship between systematic risk and expected returns, there exists limited knowledge of how early-warning risk signals could capture investors' expectations about changes in systematic risk. Leveraging on graph theory and covariance matrices, this study proposes a novel framework to develop risk signal patterns. Our approach not only discerns high-risk periods from calmer ones but also elucidates the pivotal role of interconnections among securities as indicators of systematic risk. The findings offer actionable insights for timely portfolio management and risk management responses in periods of transitions towards higher systematic risk. Moreover, by leveraging on graph theory, regulators can take timely decisions about how much liquidity to inject into the markets during periods of uncertainty. This study contributes to the literature by establishing a novel framework on linking investors' expectations and expected changes in systematic risk.

Keywords: Risk signaling; Covariance matrix; Graph theory; Systematic risk; Financial networks

1. Introduction

In 2020, the COVID-19 health crisis triggered a sudden shift in the global economy and financial markets, with financial impact comparable to the 2008 Global Financial Crisis (Ding et al. 2020, Lustig and Mariscal 2020). During periods of downturn and uncertainty, as financial markets grow in interconnectivity and complexity, it becomes highly challenging to develop early-warning risk signals (Gencay et al. 2005, Aven and Zio 2021). The history of financial markets is characterized by sudden shifts in systematic risk due to economic and political events. We define systematic risk as the vulnerability of financial markets to events which affect aggregate outcomes such as broad market returns (Pelger 2020). Recent examples are Brexit (Belke et al. 2018, Hohlmeier and Fahrholz 2018), the US–China trade war (Xu and Lien 2020), the European debt crisis (Kousenidis et al. 2013), the

China Real Estate bubble (Glaeser et al. 2017, Jiang et al. 2021), and the impact of major monetary policies in response to post-pandemic and Russo-Ukrainian conflict.

The concept of systematic risk was first introduced by Sharpe (1964) with the Capital Asset Pricing Model (CAPM), proposing a novel relationship between systematic risk and expected returns. Following this pioneering step, measures of systematic risk have mostly revolved around the concept of beta, the linear regression slope coefficient quantifying the relationship between the risk of an individual asset and risk measures of the broader market. For instance, Bowman (1979) quantified systematic risk as the beta with respect to accounting variables such as firms' leverage. For a long time, studies on systematic risk relied on an approach proposed by Bollerslev and Zhang (2003). The measure was based on Fama-French 3-factor model on high-frequency data. Recently, Ball et al. (2022) challenged the asset pricing theory in relation to systematic risk (Roll 1977) and used firm's earnings and macro-economic variables to measure firm-level systematic

*Corresponding author. Email: monomita.nandy@brunel.ac.uk

risk. Pelger (2020) implemented a factor model based on a high-frequency dataset to explain individual stock returns on sector portfolios, such as oil, financial, and energy. However, as financial markets evolved, financial theorists and practitioners alike began to grapple with the strong assumptions of CAPM and beta—such as efficient markets, normal distribution and linear relationship between risk and return. In response, seminal studies began addressing the limitations of beta estimation by borrowing on diverse mathematical fields. For instance, some approaches employ the systematic risk estimation on a wavelet approach that decomposes a given time series on a scale-by-scale basis (e.g. Gencay et al. 2005). Mestre (2023) indicated that a systematic risk varies in time and frequency so wavelet approach should be used to estimate betas of capital asset pricing model. However, most recently by criticizing the non-parametric estimation of betas based on high-frequency data, Liao and Todorov (2024) introduced ‘bias-mimicking statistic’ in relation to analyzing systematic risk. Stock and Watson (2002) use a principal component analysis (PCA) to extract common trends and movements in a large dataset of asset prices, isolating the systematic components of risk. Thus, it is evident from the above literature that there is a strong interconnection between return of securities and systematic risks. However, volatility forecasting methods remain the essential part of identifying and managing market risk. Engle and Bollerslev (1986) introduce GARCH models that still represent a cornerstone in volatility forecasting methods. On the other hand, Guo and Wohar (2006) provide evidence to time-varying volatility, while Bucci and Ciciretti (2022) offer a comparison of econometrics methods and clustering ones for detecting periods of high volatility. In recent years, the financial literature has shown a growing interest in explaining systematic risk with the interconnection between securities exemplified by networks. For instance, Onnela et al. (2003) showed how the minimum spanning tree length shrinks during periods of enhanced systematic risk—with a focus on Black Monday. Coelho et al. (2007) show how the topology of trees changes following market movements. Ciciretti and Bucci (2023) propose a graph-theory based portfolio construction methodology based on reflecting periods of high systematic risk in portfolio composition by means of eigenvector centrality.

To this day, the financial literature has fallen short in using the interconnections between securities as proxies for evaluating systematic risk, especially during financial setbacks. In few instances, the use of eigenvalues and eigenvectors[‡] to explain the complex correlation between financial assets has offered useful insights (e.g. Noh 2000; Garcia-Jorcano and Sanchis-Marco 2021). Yet, there exists a gap in their application as systematic risk signals. In a quest to address the limited research around early-warning systematic risk signals, our study poses the following research question: *How can early-warning risk signals explain investors’ expectations about changes in systematic risk?*

[‡] An eigenvector of a matrix is a nonzero vector that points to the direction of change of the matrix when a linear transformation is applied to it. The corresponding eigenvalue is the factor by which the eigenvector is scaled.

To address this, our study utilizes graph theory[‡] and matrix theory—especially regarding eigenvalues and eigenvectors—to develop proxies that not only capture but also provide early warnings of changes in systematic risk. By leveraging on graph theory, we represent the complex relationships among securities. Specifically, minimum spanning trees[§] (MST) allow us to quantify the meaningful interactions (Di Cerbo and Taylor 2021) by means of centrality measures. Despite its great capability in explaining complex systems such as financial markets, the application of graph theory to finance remains yet to be fully explored.

Risk-averse investors reflect their expectations of systematic risk on portfolio compositions. For instance, when investors expect a transition to a period of high systematic risk, higher-risk securities—such as equities—are divested to reposition the portfolio in a more defensive setup. The rebalancing involves selling risky securities, thereby creating downward pressure on prices due to lower demand. Amid high systematic risk, as investors release portfolio risk, correlations among assets jump higher. As such, the eigenvector and eigenvalues of the covariance matrix also grow higher, as the matrix stretches to account for higher levels of systematic risk. This observable market behavior indicates that the covariance matrix—which encapsulates asset return correlations—serves as a viable metric for gauging systematic risk. In financial terms, eigenvalues are the factor by which asset returns jointly change in response to changes in systematic risk. As investors’ demand changes based on their expectation of systematic risk, it alters the covariances among assets; this is reflected in the eigenvalues and eigenvectors of the covariance matrix, which effectively encapsulate the market expectation of systematic risk.

These aspects are incorporated in this paper to produce early-warning risk signals based in covariances and minimum spanning trees. Specifically, we build four metrics that act as early-warning risk signals; namely, the variance explained by the largest eigenvalue, the variance explained by the first five largest eigenvalues in excess of the first, the mean of the eigenvector centralities, and their standard deviations. The risk signals can neatly separate periods of high systematic risks from calm ones. As such, we showcase how their application can improve a mean-reverting investment strategy.

The novelty of our study is associated with its twofold contribution. First, we find that eigenvectors and eigenvalues can signal the expected changes in systematic risk with higher precision, especially during financial setbacks.[¶] Specifically, we show that the changes in systematic risk can be captured through graph centrality measures, such as eigenvector

[‡] Graph theory is the study of graphs, which are structures used to model pairwise relations between objects. A graph in this context is made up of vertices which are connected by edges. See Bollobás (1998) and Bollobás and Bela (2001) for a general introduction to graph theory.

[§] A minimum spanning tree is a subset of the edges of a connected, edge-weighted undirected graph that connects all the vertices together, without any cycles and with the minimum possible total edge weight. That is, it is a spanning tree whose sum of edge weights is as small as possible.

[¶] Estimated risk signals in this study are validated in an out-of-sample setup.

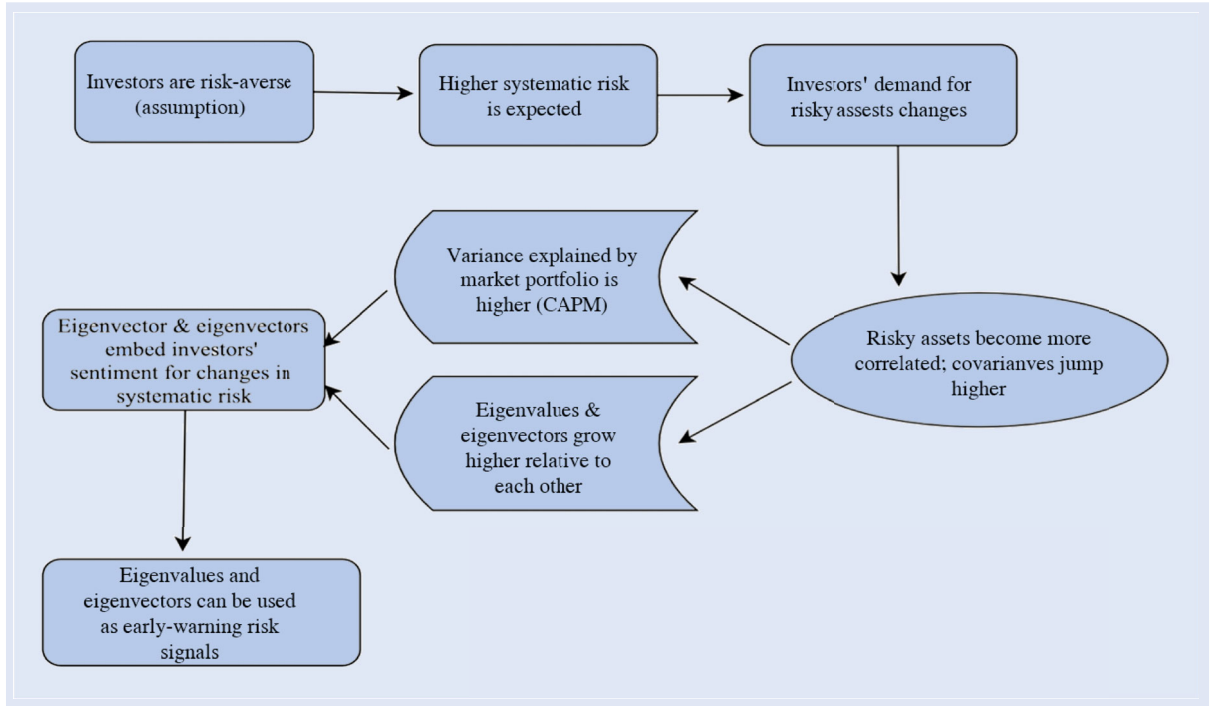


Figure 1. The link between investors' risk aversion, systematic risk, and eigenvalue-eigenvector based early-warning risk signals.

centrality. Our early-warning risk signal framework can distinguish periods of high volatility from calm ones and, as such, can improve the performance of investment strategies. Secondly, we demonstrate how graph theory can be a prosperous research field to represent the hierarchical structure of financial markets. Our graph theory approach extends existing research (e.g. Savona 2014) and encourages scholars to examine other critical questions related to systematic risk.

The societal and economic implications of detecting systematic risk changes in highly volatile periods are far-reaching. Early signals allow investors, risk managers, and regulators to respond in a timely manner to any adverse impact of crises on economic stability. Our research implies that investors and risk managers as well as regulators may incorporate the above-mentioned four metrics in their practices, to build resilient, regime-switching portfolios, to allow a more precise evaluation of risk, and to address the needed regulation adaptations in terms of monetary policy or banking regulations in response to early-spotted changes in systematic risk.

The rest of the paper is organized as follows. Section 2 illustrates the theoretical background, while section 3 introduces the methodology to build and validate the early-warning risk signals. Section 4 introduces the dataset. Section 5 shows the results. Section 6 offers a discussion and section 7 concludes.

2. Identifying the early-warning risk signals

Figure 1 outlines the economic theory underpinning our risk signals. Investors' appetite for risky assets is embedded into the covariance matrix of asset returns (Agrawal et al. 2022). The covariances of asset returns jump higher

during periods of transition towards high systematic risk (Haugen et al. 1991). Asset prices move ahead of market events because of changes in investors' implied risk expectations. Thus, changes in covariance matrices embed changes in investors' expectations—assumed to be risk-averse—due to changes in the demand for risky assets. As such, we posit that changes in investors' expectations are mirrored by changes in eigenvalues and eigenvectors. Laloux et al. (1999) and Potters et al. (2005) show that the largest eigenvalues of a covariance matrix reflect more systematic sources of risk.

When financial markets are adversely affected by a downturn in economy, all assets become more correlated with each other (Junior and Franca 2011). As a result, the covariances among the assets jump higher. Hence, a large portion of variance is explained by the change in the systematic risk. As such, the eigenvalues and eigenvectors grow higher compared to each other, becoming a proxy that capture investors' expectations about systematic risk. Therefore, we use the variance explained by the largest eigenvalue and the variance explained by the five largest eigenvalues in excess of the first as early-warning risk signals.

Moving to the graph theory side, Mantegna (1999) was the first to apply this discipline to explain the dynamic relationships among assets in financial markets. We show the capabilities of graph theory in explaining the systematic risk through minimum spanning trees, which is a collection of nodes (in our study, different securities in the market) connected to each other by links. The value of any link is proportional to the strength of the connections between nodes tied in that link. The connections are exemplified by distance measures between each pair of stocks, which are inversely proportional to linear correlations. In other words, then two securities display a high linear correlation, then their distance

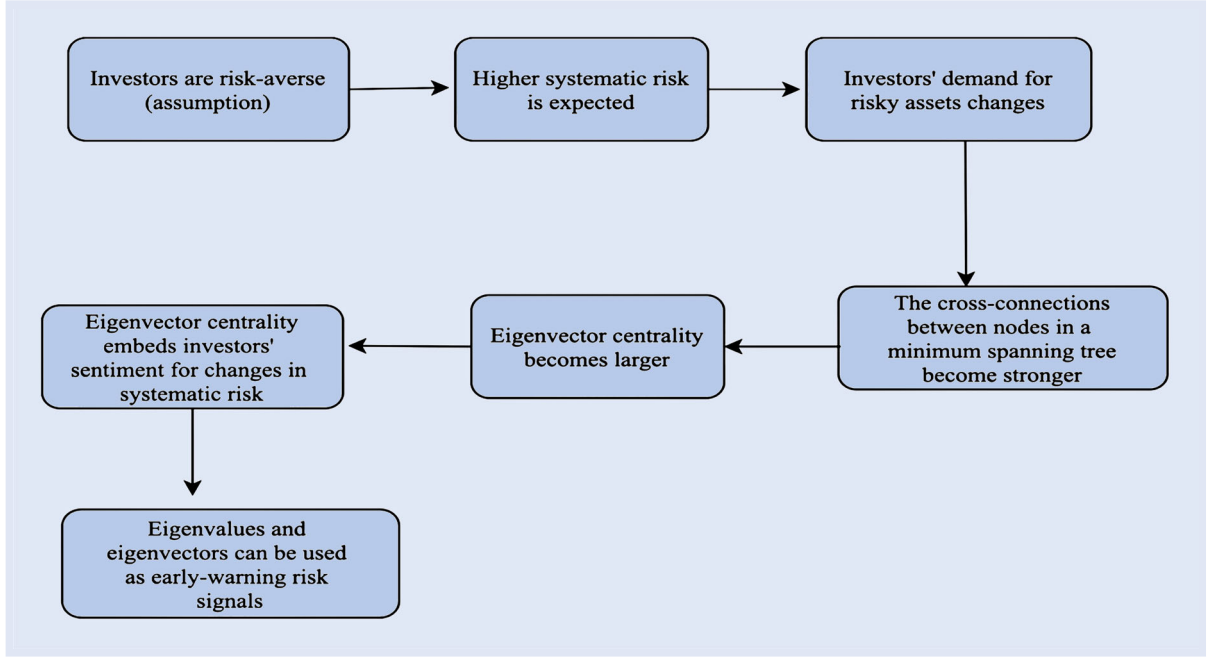


Figure 2. The link between investors' risk aversion, systematic risk, and graph theory based early-warning risk signals.

in the spanning tree is lower, hence are placed closer to each other. As during financial setbacks investors become more risk-averse, the changes in the investors' demand for risky assets can be captured by changes in the distances between pairs of nodes in the minimum spanning tree. Graph centrality measures, such as eigenvector centrality, are proxies for the degree of cross-connectivity across securities in a financial market, which can capture the expectation of the investors about changes in systematic risk. In fact, when the connections of a security to others become stronger, we observe a higher centrality value. At the same time, securities that are poorly connected show lower centrality values are placed at the peripheries of the tree—showing lesser connections to the rest of the market. As a consequence, securities that are placed at the center of the tree are better proxies of systematic risk, due to the stronger connections to the rest of the securities in the market.

We illustrate the underlying assumption of minimum spanning trees and their application in this study in figure 2. A minimum spanning tree is built from a covariance matrix. Under graph theory, the eigenvector centrality—explaining degree of connectivity among securities—becomes the eigenvector of a transformed covariance matrix, which is known as an adjacency matrix.[†] Hence, the same conclusions regarding the covariance matrix with respect to eigenvalues can be extended to the minimum spanning trees for the eigenvector centrality. In other words, when a transition to high systematic risk is expected by risk-averse investors, their muted demand for risky assets is mirrored in stronger connections between nodes—hinting at higher correlations between securities, which is reflected in higher eigenvector centrality values. Thus, graph theory allows us to derive the mean

of the eigenvector centralities and their standard deviations. Appendix A illustrates the different shapes of minimum spanning trees in periods of high systematic risk against calm periods.

3. Methodology

We split the methodology into three phases. Phase 1 calculates the early-warning risk signals. Phase 2 validates the choice of the risk signals in an in-sample setup. Phase 3 deals with the out-of-sample validation and performance analysis.

The following sections delve into the greater details of each phase.

3.1. Phase 1—Extracting early-warning risk signals

We start by calculating the linear returns of the futures asset prices as:

$$r_{t,i} = \frac{p_{t,i}}{p_{t-1,i}} - 1, \forall t = 1, \dots, T; \forall i = 1, \dots, n, \quad (1)$$

where, $r_{t,i}$ is the linear return of asset i at hour t and $p_{t,i}$ is the hourly open price of the future contract on the asset i . Successively, we estimate the realized covariance matrix through non-parametric realized covariances. As realized covariances converge in probability to the quadratic variation of the price process (Shephard and Barndorff-Nielsen 2004), we calculate weekly realized covariances, RC_t , as the aggregation of cross-products of hourly returns, such that:

$$RC_t = \sum_{\tau=1}^{N_t} r_{\tau} r'_{\tau} \quad (2)$$

[†] An adjacency matrix is a square matrix used to represent a graph. The elements of the matrix indicate whether pairs of vertices are adjacent or not in the graph.

Phase 1: Extract early-warning risk signals

1. Calculate the returns
2. Estimate the realized covariance matrix
3. Extract the first two early-warning risk signals:
 - a. Variance explained by the largest eigenvalue,
 - b. Variance explained by the largest five eigenvalues in excess of the first,
4. Transform the covariances into distances
5. Calculate the adjacency matrix
6. Calibrate the Minimum Spanning Tree (MST)
7. Calculate the eigenvector centrality of the MST
8. Extract two additional early-warning risk signals:
 - c. Arithmetic average of eigenvector centrality
 - d. Standard deviation of eigenvector centrality

Phase 2: In-sample validation

1. Extract two sub-datasets from whole dataset:
 - a. 52 consecutive weeks starting 8 weeks before the second week of October 2008 (when S&P 500 recorded the lowest value during the Global Financial Crisis)
 - b. 52 consecutive weeks starting 8 weeks before the second week of March 2020 (when S&P 500 recorded the lowest value during the COVID-19-triggered financial setback)
2. *Validation 1: Hierarchical clustering:*
 - a. Run AGNES hierarchical clustering algorithm on the two datasets
 - b. Plot dendrograms and check clustering quality,
 - c. Calculate Dunn index,
3. *Validation 2: Mean-reversion investment strategy:*
 - a. Build a naïve mean-reversion investment strategy,
 - b. Compare three filtering rules to improve the naïve strategy:
 - i. Sample realized volatility,
 - ii. eGARCH volatility,
 - iii. our four early-warning risk signals,
 - c. Calculate the Sharpe ratio, Sortino ratio and (Value-at-Risk) VaR of the naïve and each strategy filtered strategy.

Phase 3: Out-of-sample validation and performance analysis

1. Use a Variational Autoencoder (VAE) to generate one thousand synthetic datasets
2. Re-run Phase 2 on each synthetic dataset

where N_t is the number of weekly observations in the t -week, $\forall t = 1, \dots, T$, and r_t is the $n \times 1$ vector of asset returns for the τ -th observation. This ensures positive definite realized covariance matrices and makes volatility fully observable and moldable with any time series model. The realized covariance matrix RC_t is the first source from which we extract two metrics:

- The variance explained by the largest eigenvalue: $\frac{\lambda_1}{\sum_{j=1}^n \lambda_j}$
- The variance explained by the first five eigenvalues in excess of the first: $\frac{\sum_{j=2}^5 \lambda_j}{\sum_{j=1}^n \lambda_j}$

where λ_i is the eigenvalue in position i in the sorted eigenvalues vector λ extracted from the realized covariance matrix RC_t .

As Laloux et al. (1999) describe, the largest eigenvalue predominantly captures the market-wide factor of systematic

risk. Subsequently, the following three to four largest eigenvalues predominantly capture sector-specific risks, which are not solely explained by the market factor (Potters et al. 2005). These eigenvalues tend to increase as systematic risk increases. This is often a result of risk-averse investors shifting their portfolios towards more defensive strategies, leading to higher correlations, and thereby mirroring the increased systematic risk in the eigenvalues of the correlation matrix. Eigenvalues beyond the fifth are not used in this study as they mostly represent idiosyncratic risk and tend to capture firm-specific risk factors (Avellaneda 2020). As such, these are more difficult to interpret and tend to be unstable over time.

The next two metrics are extracted by leveraging on graph theory. To do so, we first build the adjacency matrix A_t from the realized covariance matrix RC_t . As constructing an adjacency matrix requires that the covariances are transformed into Euclidean metrics, following Kayo and Kimura (2011), we first convert the covariances into linear correlations:

$$\rho_{ij} = \sigma_{ij} \sigma_i \sigma_j, \forall i, j \quad (3)$$

and then calculate the Euclidean distances as:

$$d_{ij} = 1 - \rho_{ij}^2. \quad (4)$$

The distances are used to calculate the adjacency matrix A_t , whose diagonal is set equal to zero to avoid self-loops.[†] A non-zero entry in the adjacency matrix indicates the existence of a financial relationship between pairs of assets wherein the strength of the link is measured by d_{ij} . Given the adjacency matrix, we estimate the minimum spanning tree (MST) with the algorithm devised by Kruskal (1956). Kruskal's algorithm constructs a minimum spanning tree by initially sorting all edges of the graph in ascending order of their weights, then iteratively adding the shortest edge to the MST that does not form a cycle, and continuing this process until the MST contains $V-1$ edges, with V being the number of vertices in the graph. The MST is a graph theory topological representation to quantify the influence structure among the assets by means of centrality measures (Brookfield et al. 2013). A centrality measure is a function that assigns a non-negative value to each node such that the higher the value, the more the node is central. According to these, it is possible to distinguish peripheral nodes that have a limited impact on the dynamics of the network. As such, a node with higher centrality plays a major role in risk-return determination, while a peripheral node influences systematic risk to a lower degree.

In our methodology, we use the eigenvector centrality, according to which an asset displays a high centrality either by direct links to other assets or by being connected to other securities that are themselves highly central. In mathematical terms, the eigenvector centrality ζ of node x is given by the weighted average of the centrality values of its neighbors y :

$$\zeta(x) := \frac{1}{\lambda} \sum_{y \in N(x)} W(x, y) \zeta(y) \quad (5)$$

[†] A self-loop is a link that connects a vertex to itself. We eliminate self-loops since we deem self-relations as insignificant for the purpose of this study.

where λ is a constant, $W(x, y)$ is the weight of the edge between x and neighboring node y , $N(x)$ is the set of all nodes directly connected to x and the summation iterates over all such nodes y . In terms of the squared adjacency matrix A_t , we can re-write the eigenvector centrality as:

$$\zeta(x) := \frac{1}{\lambda} \sum_j A_{i,j} \zeta(y) \quad (6)$$

where $j \in N$ represents the neighbor nodes. In matrix form, this is equivalent to:

$$\lambda \zeta = A \zeta, \quad (7)$$

from where it is visible that ζ is the eigenvector of the adjacency matrix A_t and λ is its eigenvalue.

Minimum spanning trees connect nodes in a way that the total weight (proportional to the correlation between assets) is minimized across all edges. The minimal connectivity structure makes MSTs efficient in capturing the most significant relationships between securities without the noise of less significant connections. During periods of high systematic risk, correlations between securities often increase, which is reflected in larger eigenvalues and eigenvectors. As eigenvector centrality is the eigenvector of the adjacency matrix connected to the largest eigenvalue—which is representative of systematic risk (Laloux et al. 1999)—this centrality measure jumps higher too during periods of financial distress, hence capturing systematic risk (Ciciretti and Bucci 2023).

The eigenvector centrality ζ is used to calculate the last two metrics for early-warning risk signaling, namely:

- the mean eigenvector centrality $\bar{\zeta}$: $\bar{\zeta} = \frac{\sum_i \zeta_{i,t}}{n}$
- the standard deviation of the eigenvector centralities $\sigma(\zeta)$:

$$\sigma(\zeta) = \frac{\sum_i (\zeta_{i,t} - \bar{\zeta})^2}{n - 1}$$

$\forall i = 1, \dots, n$ and $\forall t = 1, \dots, T$.

3.2. Phase 2—In-sample validation

To assess the adequacy of the four metrics as early-warning risk signals, we employ a twofold in-sample validation framework. The first validation method aims at assessing whether the metrics can be used to signal a change in systematic risk. To do so, we check whether a clustering algorithm that uses the metrics can correctly label high volatility periods and calm ones. The second validation method aims at assessing whether the metrics can be employed to improve the profitability of an investment strategy.

3.2.1. Clustering. Clustering is an unsupervised learning technique aimed at grouping a collection of objects into subsets such that the objects within each cluster are more closely related than they are to objects assigned to different clusters. As such, clusters are formed upon a degree of similarity often proxied for by a distance measure. By feeding the four metrics to a clustering algorithm, we question if it is possible to separate weeks of downturns from calm weeks.

For the clustering algorithm, we choose the agglomerative hierarchical clustering (AGNES) algorithm introduced by Kaufman and Rousseeuw (1990) due to three reasons. First, its results are reproducible, as opposed to other clustering methodologies (Garge et al. 2005). In fact, partitioning methods, such as k-means, produce clusters dependent on an initial random seed. Second, financial markets have an embedded nature of hierarchy (Simon 1991) which can be reproduced through hierarchical methods, such as AGNES. Finally, among the set of hierarchical clustering methods, the AGNES algorithm is the least computationally expensive while still granting an elevated clustering precision.

Starting with each object placed in a different cluster, at each step AGNES merges the two least dissimilar clusters with the objective of minimizing the inter-cluster dissimilarity. The dissimilarity is defined upon some metrics inputted in the clustering process. In our case, these are the four metrics described in the previous section. The pairwise dissimilarities are calculated by means of the Manhattan distance, which is best suited for clustering problems with higher dimensionality (Sinwar and Rahul 2014). Since AGNES reduces the number of clusters in each step by merging them, it is essential to redefine the dissimilarities between these newly formed clusters and the remaining ones. This process involves dropping one cluster at each step and merging it into another, necessitating a reassessment of cluster distances to maintain the integrity of the clustering algorithm. For this, a linkage function is used that recalculates dissimilarities after each merge. So, we use the Ward linkage function (Ward 1963) due to its focus on minimizing the variance within clusters. Unlike other linkage functions that directly measure distances, Ward's approach identifies and merges cluster pairs that result in the least increase in total within-cluster variance, leading to more homogeneous clusters.

3.2.1.1. Mean-reversion investment strategy. **Mean-reversion investment strategy:** As a second validation method, we assess whether the metrics can be employed to improve the profitability of an investment strategy. We choose a univariate[†] mean-reverting investment strategy as this class of strategies is aimed at capturing upward directional moves—as the momentum factor—while controlling for sudden breaks in trends and reversions to the conditional mean. First, we assess the baseline performance of a naïve mean-reversion strategy based on a mean-reverting stochastic process. We compare the baseline performance to similar strategies enhanced by the inclusion of a probability of imminent transitioning to a period of high systematic risk. To calculate such probability, we employ a logit model based on the four metrics. Additionally, we add to the comparison two models wherein the probability of transition to a period of high systematic risk is calculated with traditional early-warning risk signals, such as

[†] Momentum and mean-reversion investment strategies are inherently univariate, not aligned with portfolio construction purposes. Evaluating persistent and anti-persistent behaviours should be asset-specific, as varying persistency patterns across different assets do not readily lend themselves to creating well-diversified portfolios. While cointegration strategies are an exception, the rarity of finding significantly cointegrated assets limits their practicality for robust portfolio formation. Therefore, our focus remains on the univariate case.

the sample volatility and the forecast volatility by means of an eGARCH model. The performance of the four investment strategies is assessed in terms of Sharpe ratio (Sharpe 1998), Sortino ratio (Sortino and Price 1994) and VaR.

We consider a mean-reversion investment strategy on the following assets: SPY, Nasdaq, Apple, Google, JP Morgan, Eurostoxx 50, EUR/GBP, German Bund 10 years and Italian BTP 10 years. We assume that each asset conditionally follows the following (Euler discretized) Ornstein-Uhlenbeck process (Uhlenbeck and Ornstein 1930):

$$\hat{y}_t = \kappa_0(\vartheta_t - y_{t-1}) + \hat{\sigma}_t Z_t \quad (8)$$

where \hat{y}_t is the asset return at period t , ϑ_t is the conditional long-term mean, $\hat{\sigma}$ is the volatility, $Z_t \sim N(0, 1)$, and κ_0 is a scalar parameter governing the speed of the reversion toward the conditional mean ϑ_t . A Monte Carlo simulation of \hat{y} allows to calculate the confidence interval of the estimate as:

$$CI(y_t)_\alpha \sim \hat{y}_t \pm Z_{1-\alpha/2} \frac{\hat{\sigma}_t}{\sqrt{T}} \quad (9)$$

where T is the number of observations in the simulation. To find the optimal mean-reversion parameter κ_0 , we employ a grid-search algorithm[†] to minimize the Mean Squared Error (MSE), $MSE = \frac{1}{T} \sum_{t=1}^T (\hat{y}_t - y_t)^2$, across the Monte Carlo simulations. Finally, the investment strategy is defined as follows:

$$\begin{cases} \text{buy } y_t \text{ and sell } y_{t+1} & \text{if } y_{t-1} \geq \hat{y}_t + Z_{1-\frac{\alpha}{2}} \frac{\hat{\sigma}_t}{\sqrt{T}} \\ \text{sell } y_t \text{ and buy } y_{t+1} & \text{if } y_{t-1} \leq \hat{y}_t - Z_{1-\frac{\alpha}{2}} \frac{\hat{\sigma}_t}{\sqrt{T}} \end{cases}$$

To filter the naïve baseline strategy with the early-warning signals, we estimate the probability of imminent transitioning to a high systematic risk, ω_t , with the following linear logistic model:

$$\omega_t = \beta^T x_t + u_t \quad (10)$$

where x_t represents the vector of early-warning signals, β is the vector of coefficients of the logistic regression and $u_t \sim N(0, 1)$. To deal with multicollinearity, in each period,

[†] The grid to estimate κ_0 ranges from 0 to +1.5 with size steps of 0.05. The reasons for using grid-search rather than least square are threefold:

- Grid search explores a range of parameter combinations, which is particularly useful in applications to financial markets, characterized by complex, non-linear relationships. Unlike least squares, which is tailored for linear regression with a convex loss function, grid search does not assume a convex objective function and as such works better when multiple local minima/maxima are present (James *et al.* 2013).
- Financial time series often exhibit heteroskedasticity, while methods like least squares assume homoscedasticity (constant variance). In contrast, grid search, as a non-parametric method, is not constrained by these assumptions.
- Practical financial applications demand strategies that are both effective and interpretable. Grid search provides a transparent and easily understandable approach to parameter selection.

we apply a stepwise variable selection methodology (Hocking 1976). In each application, the early-warning signals are represented by sample volatility, eGARCH volatility forecast and our eigenvalue—eigenvector-based risk framework.

In the filtered strategies, we substitute the mean-reverting parameter κ_0 with:

$$\kappa_t = \kappa_0[\varepsilon \cdot I(\omega_t > 0.5)] \quad (11)$$

where $I(\omega_t > 0.5)$ is an index function that assumes a value equal to 1 when the probability ω_t resulting from the logistic model is higher than 0.5, or 0 conversely. The optimal probability-scaling parameter, ε , is obtained via grid search.[‡] One must notice that the mean-reversion parameter, κ_t , becomes time-dependent in the filtered strategies, while κ_0 is a scalar in the baseline model.

3.3. Phase 3—out-of-sample validation of the early-warning risk signals

In this section, we introduce the methodology for the out-of-sample validation. This is particularly relevant to ensure that the calibrated models avoid overfitting and exhibit a similar behavior once applied in the real world (Bailey *et al.* 2014). As out-of-sample datasets, following literature, we use one thousand synthetic datasets generated via a Variational Autoencoder (VAE) (Kingma and Welling 2013). Successively, the two in-sample validation methods are re-applied on the one thousand synthetic datasets for the out-of-sample validation.

Consider a dataset $X = \{x^{(i)}\}_{i=1}^N$ composed of N *i.i.d.* (independent and identically distributed) samples coming from a random variable X . Let us assume that the data are generated by a random process involving an unobserved continuous random variable z . The process consists of first generating a value z from some prior distribution $p_\theta(z)$ to then generating a value x^i from the conditional distribution $p_\theta(x|z)$. Let us assume that the Probability Density Functions (PDFs) of $p_\theta(z)$ and $p_\theta(x|z)$ are differentiable almost everywhere with respect to z, θ . However, the true parameters θ and the values of the latent variable z are unknown. The objective is to find an efficient neural network approximation for the latent variable z as this would allow us to mimic the hidden random process and generate a synthetic dataset that resembles the real data. To do so, we employ a VAE. Assume the prior over the latent variables to be a centered isotropic multivariate Gaussian $p_\theta(z) = N(z, 0, I)$. Let $p_\theta(x|z)$ be a multivariate Gaussian with θ estimated via a fully connected neural network with a single hidden layer. The true posterior is intractable but, assuming that it is approximated by a Gaussian distribution with an approximately diagonal covariance, then the variational approximate posterior is a multivariate Gaussian with a diagonal covariance structure:

$$\log q_\theta(z|x^i) \log N(z, \mu^i, \sigma^i I) \quad (12)$$

where $q_\theta(z|x^i)$ is based on an alternative technique for sampling z such as Monte Carlo and (μ^i, σ^i) are the mean and

[‡] The grid to estimate ε ranges from 0.05 to 1 with size steps of 0.05.

Table 1. List of the assets composing the dataset.

Name	Ticker code	Description
S&P 500	ES	The equity index composed of the largest ~ 500 US stocks by market capitalization
US treasury rate 2y	TU	The yield of a US treasury bond with 2 years maturity
US treasury rate 10y	TY	The yield of a US treasury bond with 10 years maturity
Gold	GC	The future contract on gold
EUR / USD	EURUSD	The Euro/US Dollar exchange rate
WTI	CL	West Texas Intermediate, a future contract on oil produced in the USA

standard deviation of the approximate posterior which are outputted by the neural network as nonlinear functions of x^i and the variational parameters ϕ .

Afterwards, one simply needs to sample from the posterior $z^{i,l} \sim q_\theta(z|x^i)$ with $z^{i,l} = g_\theta(x^i, \epsilon^l) = \mu^i + \sigma^i \epsilon^l$ where $\epsilon^l \sim N(0, I)$. It can be proven that the Kullback-Leibler divergence can be computed without estimation and the resulting estimator for the datapoint x^i is given by:

$$\mathcal{L}(\theta, \phi, x^i) \cong \frac{1}{2} \sum_{j=1}^J (1 + \log((\sigma_j^{(i)})^2) - (\mu_j^{(i)})^2 - (\sigma_j^{(i)})^2) + \frac{1}{L} \sum_{l=1}^L \log p_\theta(x^i|z^{i,l}) \quad (13)$$

where $\log p_\theta(z^{i,l})$ is a Gaussian fully connected neural network decoding term, j runs over the J dimension of the latent variable z , and l refers to the individual samples drawn from the posterior $q_\theta(z|x^i)$, where L represents the number of samples in the Monte Carlo approximation.

The robustness of the VAE used in our study is coherent with other studies (Camuto et al. 2021). With this VAE application, new data can be used to validate an out-of-sample testable hypothesis which resembles the real-life characteristics of the original data distributions. Further, research has shown how VAE outperforms traditional synthetic sampling methods under various evaluation metrics (Wan et al. 2017).

The VAE model specification is fine-tuned via grid search on the whole continuous dataset. We use two hidden layers with five hidden units activated by means of the leaky Rectified Linear Unit (RELU) function, initialized with the Glorot kernel initializer, with a Ridge regularizer of 0.02 and a Lasso regularizer of 0.01. The training is done via Stochastic Gradient Descent with a Nesterov momentum of 0.6. The loss function targets the representation error via the consistency of the mean squared error.

4. Data collection

We used a time-series dataset of hourly prices of different asset classes assets sourced from Bloomberg for the period of January 2008 to November 2023. The final dataset consists of 4009 daily observations (40 010 hourly observations). This period is selected to isolate the low interest rate environment characterizing financial markets since the 2008 Global Financial Crisis. The time series refers to open prices of futures contracts on the selected assets, due to the higher liquidity

of future contracts. Table 1 lists the assets used in this study and the respective Ticker codes. The assets selected are generally monitored by investors to gauge the current sentiment of financial markets. In fact, while the S&P 500 is regarded as the most important equity index in the world, the US treasury rates are a gauge of economic risk as investors trade-off equity positions for bonds during periods of financial uncertainty (Smales 2014). Moreover, the difference between the 2 years rate and the 10 years rate is closely monitored by investors, as periods of inverted yield curve (i.e. when this difference becomes positive) are empirically associated with higher volatility markets (Ang et al. 2006). In addition, gold is generally regarded as a safe asset due to its store of value property (Baur and McDermott 2010) and the West Texas Intermediate (WTI) is a gauge of political uncertainty (Bittlingmayer 1998). Finally, the US dollar historically exhibits positive returns in highly volatile periods (Wen et al. 2018).

5. Results

This section reports the in-sample and out-of-sample validation results.

5.1. Clustering results

As the clustering validation is aimed at assessing whether the four metrics can be used as an early-warning risk signal for periods of high systematic risk, in this section we only use a sub-sample of the entire dataset centered around 2008 and 2020, the years where the two most recent financial crisis (2008 Global Financial Crisis and 2020 COVID-19 financial setback) were recorded. For both periods, the start date is set at eight weeks before the day that the S&P 500 recorded the minimum value during the downturn period, while the end date is 52 weeks after the start date. As a result, the two datasets span the 52 consecutive weeks from the first week of August 2008 to July 2009 and the third week of January 2020 to January 2021, respectively. The choice of the week numbers to define a week as highly volatile comes from direct observation of the Volatility Index (VIX), a tradable index whose value is proportional to the overall implied volatility in the market. Appendix B reports a plot of the VIX.

The validation results from the clustering are graphically presented by the dendrograms in figures 3 and 4. In both cases, the weeks around the market plunge—weeks 7–11 by construction of both samples—are grouped together in a neatly separated cluster. In January 2020, weeks of elevated volatility are recorded even before COVID-19 was declared




Figure 3. The dendrogram in the Global Financial Crisis of 2008 database.




Figure 4. The dendrogram in the 2020 COVID-19-related dataset.

a pandemic, amid concerns over the overvaluation of the equity markets, the China-USA trade war, and the first news of the health emergency spreading outside of China. The high volatility in November is not separately classified by the clustering algorithm as the transition to high systematic risk already occurred in the previous month.

Moreover, as a statistical measure to validate the clusters, we use the Dunn index (Dunn 1973, 1974) calculated as the ratio between the minimum between-clusters distance and the maximal within-cluster distance. A higher value implies better clustering. We normalize the Dunn Index between 0 and 1 for better interpretation. Our procedures obtain an in-sample Dunn index of 0.93 and 0.85 for the Global Financial Crisis and the COVID-19 sub-samples, respectively. For the out-of-sample setup, the average out-of-sample Dunn index across the one thousand samples is 0.89. This allows us to conclude that the metrics can capture a different hierarchy in response to structural changes in the covariance matrices.

Finally, figure 5 illustrates the in-sample behavior of the four metrics during periods of high systematic risk against calmer periods in the two subsets. The red shaded area in the densities marks the difference in distribution between calm

weeks and weeks of high systematic risk, where we defined the high systematic risk weeks as the four weeks before and after the S&P 500 recording its lowest value in both sub-samples. We conduct both a t-test for mean comparison and a Kolmogorov-Smirnov test, both of which reject the null hypothesis of identical distributions between the two periods at a 95% confidence level.

Starting from the variance explained by the largest eigenvalue and the difference of variance explained by the largest five eigenvalues in excess of the first, the density functions exhibit a jump in mean during the volatile weeks, allowing us to conclude that systematic risk becomes more predominant during highly volatile weeks. In terms of eigenvector centrality mean and standard deviation, their distributions also jump higher in mean and kurtosis during volatile weeks. Once again, this is due to the systematic source of risk explaining a larger portion of asset movements. In fact, we show how the eigenvector centrality is the eigenvector of the adjacency matrix, which, in turn, is built out of covariance matrices. Figures A1 in Appendix A reports a direct comparison of the minimum spanning tree generated by the companies included in the S&P 500 in August 2008 with the one in March 2020




Figure 5. The distributions of the four early-warning metrics in the two datasets.

and the one in a calm trading week (first week of August 2021). The size of the vertices represents each stock's eigenvector centrality, wherein each node represents a single stock. As is visible, the MSTs at the onset of the two volatile periods present a similar denser structure where the connections

between the assets are stronger. The interconnectivity among securities, together with the eigenvalues of the covariance matrix on which the adjacency matrix is built, bounced higher because of increased market risk. During the calm period, instead, the structure of the tree appears stretched on four

Table 2. The average Sharpe ratios of the mean-reversion strategies across the one thousand synthetic datasets on the different assets.

	Buy & hold	Baseline mean reversion	Sample volatility	eGARCH	Eigenvalue-Eigenvector based signals
SPY	0.199	0.151	0.283	0.105	1.067
Nasdaq	0.154	0.174	0.149	0.219	1.138
Apple	0.291	0.430	0.351	0.215	0.884
Google	0.225	0.239	0.237	0.307	0.921
JP Morgan	0.309	0.205	0.253	0.307	0.852
Eurostoxx 50	0.212	0.212	0.360	0.284	1.141
EUR/GBP	0.716	0.109	0.517	0.541	0.982
German 10y	0.111	0.704	0.269	0.671	0.863
Italy 10y	0.190	0.135	0.215	0.259	1.173

Table 3. The average Sortino ratio of the mean-reversion strategies across the one thousand synthetic datasets on the different assets.

	Buy & hold	Baseline mean reversion	Sample volatility	eGARCH	Eigenvalue-Eigenvector based signals
SPY	0.920	0.790	0.840	0.940	2.520
Nasdaq	1.300	0.720	0.600	1.270	2.340
Apple	0.790	0.900	0.980	0.620	2.330
Google	0.700	0.640	0.840	0.550	2.080
JP Morgan	1.690	1.320	0.730	1.090	2.330
Eurostoxx 50	0.780	0.700	0.770	0.600	2.900
EUR/GBP	0.800	0.620	0.550	0.840	2.530
German 10y	1.370	1.620	1.280	1.070	2.680
Italy 10y	0.510	1.110	1.280	0.900	1.490

major sections, with the core of the tree placed at the top. As such, we conclude that financial markets exhibit higher concentration during markets' crises as fewer concentrated assets in the center of the graph are capable to explain a larger portion of the co-movements, resulting in higher eigenvector centralities.

5.2. Investment strategy results

In this section, we compare the performance of the baseline mean-reversion investment strategy with the filtered versions, namely, the one filtered with sample volatility, the one filtered with the eGARCH volatility forecast, and the one where the four risk-signal metrics are employed. Moreover, we add to the comparison a buy and hold strategy on each asset as a benchmark.

For the parameters in equation (8), we use the sample standard deviation of the returns on a rolling window of the last 20 weeks as σ_t estimate, while ϑ_t is estimated via the moving average of the last 50 returns. Optimal ε is estimated at 1.2, 1.15, and 1.2, respectively, for the three above-mentioned filtered strategies, while optimal κ_0 is estimated at 0.9, 0.75, and 0.8, respectively.

Tables 2–4 report the average Sharpe ratio, Sortino ratio and VaR at 95% confidence level for all strategies on all the assets in the out-of-sample setting of the one thousand synthetic datasets generated by the VAE model. The evaluation metrics show that superior performance and lower downside risk is linked to the usage of the eigenvalue-eigenvector based metrics for filtering the momentum investment strategy.

The other two competing filtered strategies, instead, do not manage to consistently outperform the baseline model across the out-of-sample datasets. A key aspect of this success lies in the strategy's reduced volatility, contributing significantly to an improved Sharpe ratio. By adapting to structural shifts in covariance matrices and minimum spanning trees, the eigenvalue-eigenvector approach demonstrates agility in responding to market changes. As $\varepsilon > 1$, this filtered strategy can increase the speed of reversion to the conditional mean[†] when a transition to a high systematic risk is detected. This is consistent with financial markets as prices are empirically observed to revert in periods of high volatility.

6. Discussion

In this section, we summarize the main findings of the study and discuss the implications of our studies for research and practice. Firstly, we find the set of eigenvalues and eigenvectors of the covariance matrix of asset returns can be used to measure expected changes in systematic risks with higher precision, especially during financial setbacks. We prove that the expectation of investors about change in systematic risk and their appetite for risky assets are embedded into the covariance matrix of asset returns. Thus, the first two matrices we

[†] In fact, according to equation (11), when $\varepsilon > 1$ and $\omega_t > 0.5$, then $\kappa_t > \kappa_0$, resulting in a higher speed of reversion towards the conditional mean.

Table 4. The average Value at Risk at 95% confidence level of the mean-reversion strategies across the one thousand synthetic datasets on the different assets.

	Buy & hold	Baseline mean reversion	Sample volatility	eGARCH	Eigenvalue-Eigenvector based signals
SPY	− 10.669	− 12.106	− 13.705	− 12.886	− 8.576
Nasdaq	− 14.948	− 14.641	− 14.563	− 13.183	− 8.771
Apple	− 5.327	− 6.002	− 8.969	− 14.576	− 8.240
Google	− 12.427	− 13.428	− 14.926	− 10.419	− 9.085
JP Morgan	− 10.475	− 9.722	− 14.357	− 6.575	− 8.432
Eurostoxx 50	− 6.459	− 11.230	− 7.124	− 12.793	− 7.220
EUR/GBP	− 12.999	− 6.396	− 11.602	− 10.711	− 5.242
German 10y	− 14.944	− 7.465	− 9.467	− 9.372	− 4.440
Italy 10y	− 9.802	− 12.518	− 10.970	− 8.594	− 5.617

built and validated in this paper can be used to construct early-warning risk signals and can be applied to an efficient risk management tool to develop successful investment strategies during financial setbacks.

Secondly, so far, prior research shows that the covariances of asset returns jump higher during periods of transition towards high systematic risk (Haugen et al. 1991, Bollerslev et al. 2018). However, these studies are unable to capture relevant systematic risk information when the financial market is highly volatile during the financial crisis. We demonstrate that, in volatile situations in financial markets, graph theory can be highly efficient in explaining the theoretical basis for the empirical analysis of the hierarchical structure of financial markets. Based on the graph theory, the risk signals developed in this paper can capture relevant information required to develop an early-risk signals framework. Graph theory—through its minimum spanning tree tool—offers useful representations of financial markets. Centrality measures, such as the eigenvector centrality, can be used to summarize the degree of cross-connectivity across securities in a financial market and hence could produce early-risk warning signals. Moreover, hierarchical representation of financial markets—obtained by application of graph theory—can exemplify the propagation of systematic risk in financial markets, which can serve to set up risk management strategies by practitioners. In practical terms, for a portfolio manager, an early portfolio rebalancing against changes in systematic risk is likely to shield the portfolio from market downturns. From a risk management perspective, instead, risk measures, such as value at risk, can discount a higher probability of tail-risk events when structural changes in covariance matrices are detected. Thus, based on graph theory and covariance matrices, we explain the theoretical implication of our early-warning risk signals framework.

Finally, as systematic risk explains a larger portion of asset movements during periods of enhanced volatility, empirical correlations between assets tend to jump higher. The covariances jump higher, together with their eigenvalues and eigenvectors. From a graph theory standpoint, instead, during periods of high systematic risks, the securities concentrate in the core of the tree. Hence, the graph becomes denser as the securities turn out to be more interconnected with each other (Ciciretti and Bucci 2023). This structural change of the tree is captured by the eigenvector of the adjacency matrix.

For this reason, the mean and standard deviation change with the expectations of systematic risk and can be used as early-warning risk signals.

The early-warning risk signals framework has broad practical implications. The above findings are novel as the framework establishes a link between investors' expectations, expected changes in systematic risk, and changes in eigenvectors and eigenvalues. Such a model can lead to a better understanding of the systematic risk during financial setbacks among practitioners as well as can provide new evidence for academic scholars about which factors to consider in systematic risk determination. From a practical standpoint, investors and risk managers can benefit by re-balancing their portfolios in a timely manner and accurately measuring market risk. Moreover, regulators may benefit by taking timely decisions about how much liquidity to inject into the market during the uncertain period to stabilize the economy as well as to set up fiscal recovery plans.

The reason why investment banks, traders, hedge funds, and others care about systematic risk is simply explained by the Capital Asset Pricing Model (CAPM). CAPM shows how returns can be linearly decomposed into two components: one dependent on systematic risk and one on idiosyncratic risk. Moreover, while the latter can be diversified away, the former exposure cannot be curtailed even with an asymptotically large number of securities in the portfolio. As such, the systematic risk component affects the return generation process for such practitioners and cannot be diversified away.

7. Conclusion

The overarching aim of this paper is to develop an early-warning risk signals framework to explain the investors' expectations about changes in systematic risk in the financial markets. To address this aim, we draw on graph theory and covariance matrices. Our framework, which is based on four eigenvalues and eigenvector metrics, shows how the interconnection between securities in financial markets can serve as proxies for evaluating systematic risk. In addition, our framework allows the embedding of investors' market sentiment with a probability of transitioning towards a high systematic

risk period. Thus, investors and risk managers could incorporate such metrics to build resilient, regime-switching portfolios, which enables a more precise evaluation of risk. From practitioners' standpoint, during financial setbacks, portfolio managers are concerned with moves in systematic risk as the returns of the risky assets in the portfolios are affected by the market systematic component. Investment portfolios often lack swift reactivity to protect against periods of high systematic risk and have the potential of benefiting from early-warning risk signals to reposition the portfolio in a timely manner. In a similar way, risk managers are generally concerned with tail-risk events, which may generate systematic or idiosyncratic risks as explained in the CAPM theory. Based on our findings, risk managers may greatly modify risk measures by including early-warning risk signals based on eigenvalue-eigenvector matrices. The study is not free from limitations. First, we focused on the global financial crisis and the crisis generated by COVID-19. It is worth investigating whether our framework will be applicable in times of other types of crises, e.g. geopolitical crises. Second, our study mainly focuses on systematic risk. Thus, it might be interesting for future studies to validate our framework to detect shifts to systematic risk. Moreover, future studies can focus on multivariate investment applications—such as portfolio construction—employing our early-warning systematic risk framework.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data sharing

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

ORCID

Vito Ciciretti  <http://orcid.org/0009-0004-9571-5179>
 Monomita Nandy  <http://orcid.org/0000-0001-8191-2412>
 Alberto Pallotta  <http://orcid.org/0000-0001-7828-3715>
 Suman Lodh  <http://orcid.org/0000-0002-4513-1480>

References

- Agrawal, R., Roy, U. and Uhler, C., Covariance matrix estimation under total positivity for portfolio selection. *J. Financ. Econ.*, 2022, **20**(2), 367–389.
- Ang, A., Piazzesi, M. and Wei, M., What does the yield curve tell us about GDP growth? *J. Econom.*, 2006, **131**(1-2), 359–403.
- Avellaneda, M., Hierarchical PCA and applications to portfolio management. *Rev. mex. econ. finanz.*, 2020, **15**(1), 1–16.
- Aven, T. and Zio, E., Globalization and global risk: How risk analysis needs to be enhanced to be effective in confronting current threats. *Reliab. Eng. Syst. Saf.*, 2021, **205**, 107270.
- Bailey, D. H., Borwein, J., Lopez de Prado, M. and Zhu, Q. J., Pseudo-mathematics and financial charlatanism: The effects of backtest overfitting on out-of-sample performance. *Notices Amer. Math. Soc.*, 2014, **61**(5), 458–471.
- Ball, R., Sadka, G. and Tseng, A., Using accounting earnings and aggregate economic indicators to estimate firm-level systematic risk. *Rev. Acc. Stud.*, 2022, **27**(2), 607–646.
- Baur, D. G. and McDermott, T. K., Is gold a safe haven? International evidence. *J. Bank. Finance*, 2010, **34**(8), 1886–1898.
- Belke, A., Dubova, I. and Osowski, T., Policy uncertainty and international financial markets: The case of Brexit. *Appl. Econ.*, 2018, **50**(34-35), 3752–3770.
- Bittlingmayer, G., Output, stock volatility, and political uncertainty in a natural experiment: Germany, 1880–1940. *J. Finance*, 1998, **53**(6), 2243–2257.
- Bollerslev, T. and Zhang, B. Y., Measuring and modeling systematic risk in factor pricing models using high-frequency data. *J. Empir. Finance*, 2003, **10**(5), 533–558.
- Bollerslev, T., Hood, B., Huss, J. and Pedersen, L. H., Risk everywhere: Modeling and managing volatility. *Rev. Financ. Stud.*, 2018, **31**(7), 2729–2773.
- Bollobás, B., Random graphs. In *Modern Graph Theory* (Graduate Texts in Mathematics, Vol. 184). pp. 215–252, 1998 (Springer: New York, NY).
- Bollobás, B. and Bela, B., *Random Graphs*, 73, pp. 1–20, 2001 (Cambridge University Press).
- Bowman, R. G., The theoretical relationship between systematic risk and financial (accounting) variables. *J. Finance*, 1979, **34**(3), 617–630.
- Brookfield, D., Boussabaine, H. and Su, C., Identifying reference companies using the book-to-market ratio: A minimum spanning tree approach. *Eur. J. Fin.*, 2013, **19**(6), 466–490.
- Bucci, A. and Ciciretti, V., Market regime detection via realized covariances. *Econ. Model.*, 2022, **111**, 105832.
- Camuto, A., Willetts, M., Roberts, S., Holmes, C. and Rainforth, T., Towards a theoretical understanding of the robustness of variational autoencoders. In *International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 3565–3573, 2021.
- Coelho, R., Gilmore, C. G., Lucey, B., Richmond, P. and Hutzler, S., The evolution of interdependence in world equity markets—evidence from minimum spanning trees. *Physica A. Stat. Mech. Appl.*, 2007, **376**, 455–466.
- Ciciretti, V. and Bucci, A., Building optimal regime-switching portfolios. *N. Am. J. Econ. Finance*, 2023, **64**, 101837.
- Di Cerbo, L. F. and Taylor, S., Graph theoretical representations of equity indices and their centrality measures. *Quant. Finance*, 2021, **21**(4), 523–537.
- Ding, W., Levine, R., Lin, C. and Xie, W., *Corporate immunity to the COVID-19 pandemic*, Technical report, 2020 (National Bureau of Economic Research).
- Dunn, J. C. A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters, pp. 32–57, 1973 (Taylor & Francis).
- Dunn, J. C., Well-separated clusters and optimal fuzzy partitions. *J. Cybern.*, 1974, **4**(1), 95–104.
- Engle, R. F. and Bollerslev, T., Modelling the persistence of conditional variances. *Econom. Rev.*, 1986, **5**(1), 1–50.
- Garcia-Jorcano, L. and Sanchis-Marco, L., Systemic-systematic risk in financial system: A dynamic ranking based on expectiles. *Int. Rev. Econ. Financ.*, 2021, **75**, 330–365.
- Garge, N. R., Grier, P. P., Alan, P. S., Bernard, S. G. and David, B. A., Reproducible clusters from microarray research: Whither? *BMC Bioinformatics*, 2005, **6**(2), 1–11.
- Gencay, R., Selcuk, F. and Whitcher, B., Multiscale systematic risk. *J. Int. Money Finance*, 2005, **24**(1), 55–70.
- Glaeser, E., Huang, W., Ma, Y. and Shleifer, A., A real estate boom with Chinese characteristics. *J. Econ. Perspect.*, 2017, **31**(1), 93–116.

- Guo, W. and Wohar, M. E., Identifying regime changes in market volatility. *J. Financ. Res.*, 2006, **29**(1), 79–93.
- Haugen, R. A., Talmor, E. and Torous, W. N., The effect of volatility changes on the level of stock prices and subsequent expected returns. *J. Finance*, 1991, **46**(3), 985–1007.
- Hocking, R. R., A biometrics invited paper. The analysis and selection of variables in linear regression. *Biometrics*, 1976, **32**, 1–49.
- Hohlmeier, M. and Fahrholz, C., The impact of Brexit on financial markets—taking stock. *Int. J. Financ. Stud.*, 2018, **6**(3), 65.
- James, G., Witten, D., Hastie, T. and Tibshirani, R., *An Introduction to Statistical Learning*, 2013 (Springer).
- Jiang, Y., Zheng, L. and Wang, J., Research on external financial risk measurement of China real estate. *Int. J. Finance Econ.*, 2021, **26**(4), 5472–5484.
- Junior, L. and Franca, Í., Correlation of financial markets in times of crisis. *Physica A Stat. Mech. Appl.*, 2011, **391**, 187–208. doi:[10.1016/j.physa.2011.07.023](https://doi.org/10.1016/j.physa.2011.07.023).
- Kaufman, L. and Rousseeuw, P. J., *Finding Groups in Data: An Introduction to Cluster Analysis*, 1990 (John Wiley & Sons).
- Kayo, E. K. and Kimura, H., Hierarchical determinants of capital structure. *J. Bank. Finance*, 2011, **35**(2), 358–371.
- Kingma, D. P. and Welling, M., Autoencoding variational bayes, arXiv preprint arXiv:1312.6114, 2013.
- Kousenidis, D. V., Ladas, A. C. and Negakis, C. I., The effects of the European debt crisis on earnings quality. *Int. Rev. Financ. Anal.*, 2013, **30**, 351–362.
- Kruskal, J. B., On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Am. Math. Soc.*, 1956, **7**, 48–50.
- Laloux, L., Cizeau, P., Bouchaud, J.-P. and Potters, M., Noise dressing of financial correlation matrices. *Phys. Rev. Lett.*, 1999, **83**(7), 1467.
- Liao, Y. and Todorov, V., Changes in the span of systematic risk exposures. *Quant. Econ.*, 2024, **15**(3), 817–847.
- Lustig, N. and Mariscal, J., How COVID-19 could be like the global financial crisis (or worse). In *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever it Takes*, edited by R. Baldwin and B. Weder di Mauro, pp. 185–190, 2020 (CEPR Press).
- Mantegna, R. N., Hierarchical structure in financial markets. *Eur. Phys. J. B.*, 1999, **11**(1), 193–197.
- Mestre, R., Stock profiling using time–frequency-varying systematic risk measure. *Financ. Innov.*, 2023, **9**(1), 52.
- Noh, J. D., Model for correlations in stock markets. *Phys. Rev. E*, 2000, **61**(5), 5981.
- Onnela, J. P., Chakraborti, A., Kaski, K. and Kertesz, J., Dynamic asset trees and Black Monday. *Physica A Stat. Mech. Appl.*, 2003, **324**(1-2), 247–252.
- Pelger, M., Understanding systematic risk: A high-frequency approach. *J. Finance*, 2020, **75**(4), 2179–2220.
- Potters, M., Bouchaud, J.-P. and Laloux, L., Financial applications of random matrix theory: Old laces and new pieces, arXiv preprint physics/050711, 2005.
- Roll, R., A critique of the asset pricing theory's tests—Part I: On past and potential testability of the theory. *J. Financ. Econ.*, 1977, **4**, 129–176.
- Savona, R., Hedge fund systemic risk signals. *Eur. J. Oper. Res.*, 2014, **236**(1), 282–291.
- Sharpe, W. F., Capital asset prices: A theory of market equilibrium under conditions of risk. *J. Finance*, 1964, **19**(3), 425–442.
- Sharpe, W. F., The sharpe ratio. In *Streetwise—the Best of the Journal of Portfolio Management*, edited by P. L. Bernstein and F. J. Fabozzi, Vol. 3, pp. 169–185, 1998 (Princeton University Press: New York).
- Shephard, N. and Barndorff-Nielsen, O., Econometric analysis of realized covariation: High frequency-based covariance, regression, and correlation in financial economics. *Econometrica*, 2004, **72**, 885–925.
- Simon, H. A., The architecture of complexity. In *Facets of Systems Science*, edited by G. J. Klir, pp. 457–476, 1991 (Plenum Press).
- Sinwar, D. and Rahul, K., Study of Euclidean and Manhattan distance metrics using simple k-means clustering. *Int. J. Res. Appl. Sci., Eng. Technol.*, 2014, **2**(5), 270–274.
- Smales, L. A., News sentiment and the investor fear gauge. *Fin. Res. Lett.*, 2014, **11**(2), 122–130.
- Sortino, F. A. and Price, L. N., Performance measurement in a downside risk framework. *J. Invest.*, 1994, **3**(3), 59–64.
- Stock, J. H. and Watson, M. W., Forecasting output and inflation: The role of asset prices. *J. Econ. Lit.*, 2002, **40**(3), 788–829.
- Uhlenbeck, G. E. and Ornstein, L. S., On the theory of the Brownian motion. *Phys. Rev.*, 1930, **36**(5), 823.
- Wan, Z., Zhang, Y. and He, H., Variational autoencoder based synthetic data generation for imbalanced learning. In *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1–7, 2017 (IEEE).
- Ward, J. H., Jr., Hierarchical grouping to optimize an objective function. *J. Am. Stat. Assoc.*, 1963, **58**(301), 236–244.
- Wen, F., Xiao, J., Huang, C. and Xia, X., Interaction between oil and us dollar exchange rate: Nonlinear causality, time-varying influence and structural breaks in volatility. *Appl. Econ.*, 2018, **50**(3), 319–334.
- Xu, Y. and Lien, D., Dynamic exchange rate dependences: The effect of the US-China trade war. *J. Int. Financ. Markets Inst. Money*, 2020, **68**, 101238.

Appendices

Appendix A. Plot of Minimum Spanning Trees




Figure A1. Minimum Spanning Trees representation of S&P500 in three periods: the first two are in a high volatility environment, i.e. during October 2008 (left) and during March 2020 (center). The one on the right is the MST during a calm period.

Appendix B. Plot of the VIX during 2008 and its 10 days moving average

Figure B1. The time-series plot of the VIX index together with its 10 days moving average.