

A Machine Learning application to signal
financial risk during setback periods

Vito Ciciretti

Independent Researcher
BorsigStr. 3, Berlin - 10115, Germany
Email: vciciretti8@gmail.com

Monomita Nandy

Brunel University
Brunel Business School
Kingston Lane, Uxbridge - Middlesex UB8 3PH, UK
Email: Monomita.nandy@brunel.ac.uk

Alberto Pallotta

Teaching fellow
Middlesex University
The Burroughs, Hendon, London - NW4 4BT, UK
Email: a.pallotta@mdx.ac.uk

Suman Lodh

Middlesex University
Middlesex University Business School
The Burroughs, Hendon, London - NW4 4BT, UK
Email: s.lodh@mdx.ac.uk

Jekaterina Kartasova

Middlesex University
Middlesex University Business School
The Burroughs, Hendon, London - NW4 4BT, UK
Email: j.kartasova@mdx.ac.uk

A Machine Learning application to signal financial risk during setback periods

Abstract

The aim of this research is to offer empirical evidence on the usage of eigenvalues and eigenvectors as risk signal able to capture exogenous changes in financial markets' sentiment. Starting from the covariance matrix of asset returns and their minimum spanning tree topological representation, we compare the 2008 Global Financial Crisis with the 2020 financial setback triggered by the COVID-19 health crisis. We show that a hierarchical clustering algorithm on them is capable to successfully separate calm periods from volatile ones. Moreover, we show that improved performances are linked to filtering a momentum investment strategy with the probability of being in a volatile market resulting from a linear logistic model on the metrics. Finally, we validate our findings using out-of-sample (ten synthetic) datasets generated using a Variational Auto-Encoder neural network. We contribute to the literature by showcasing the prominent forward-looking view offered by eigenvalues-eigenvectors compared to traditional volatility measures. Our findings are important to investors, risk managers and regulators in timely signaling risk changes and creating profitable investment strategies.

JEL classification: G01, G11

Keywords: clustering, COVID-19, minimum spanning tree, eigenvalues, momentum, variational autoencoder

1 Introduction

In 2020, the COVID-19 health crisis triggered a sudden shift in lives influencing the global economy and financial markets (Ding et al. 2020). Unexpectedly emerged, COVID-19 resulted in complete social isolation, causing a variety of adverse economic effects. The health crisis significantly influenced both the global economy and financial markets more than the Global Financial Crisis (GFC) in 2008. Moreover, financial markets are often characterized by less pronounced, yet more frequent exogenous events – such as Brexit (Belke et al. 2018, Hohlmeier & Fahrholz 2018), US-China trade war (Xu & Lien 2020), European debt crisis (Kousenidis et al. 2013), China Real Estate bubble (Glaeser et al. 2017, Jiang et al. 2021) – which represent a capital management challenge to both portfolio managers and risk managers. Thus, in this paper we ask the following question: how efficiently eigenvalues and eigenvectors capture systematic risk during periods of financial setbacks?

Extensive research has documented that traditional risk measures - such as those based on volatility - are mainly reflective of past markets' behavior rather than being forward-looking (Duan & Zhang 2014). Eigenvalues and eigenvectors, on the other hand, being the output of multidimensional covariance matrices, consistently embed future investors' views and their sentiments. As such, their application to financial problems can support practitioners to timely react to exogenous systematic shocks. The aim of this research is to offer an empirical evidence on the usage of eigenvalues and eigenvectors as risk signal able to capture exogenous changes in financial market conditions implied in covariance matrices of sentiment assets ¹ returns. In fact, eigenvalues-eigenvectors are geometrically seen as the factor by which a matrix changes when it is stretched or sheered (Avellaneda & Lee 2010). Since covariance matrices and volatilities are used to measure financial risk (Andersen et al. 2013), in financial applications the set of eigenvalues-eigenvectors represent the factor by which financial risk changes in response to asset returns movements - an empirically observed behavior of the financial market during economic uncertainties. However, we must point out that the aim is not to predict exogenous events, rather to offer a risk signal that allows to quickly react to changes in systematic risk. From our empirical analysis, eigenvectors emerge as a useful tool that can benefit both portfolio managers (via effective portfolio rebalancing) and risk managers (via improved Value at Risk methodologies (Gençay et al. 2003)). Thus, we showcase how the use of eigenvalues and eigenvectors incorporates better predicted risk factors for systematic risk compared to traditional risk models. The findings are useful to practitioners to incorporate in their daily financial management. In the forecasting literature there is an already established link between covariance matrices and eigenvalues (Papailias & Thomakos 2017). However, in the finance literature there is not enough evidence that covariance matrices perform efficiently during periods of market distress. Thus, to address the above gap in the finance literature, we include the graph

¹we define them as widely exchanged assets that can be used as a proxy for describing the ongoing investors' beliefs in a specific market, such as S&P500, US treasury yields, gold, etc.

theory equivalent of the covariance matrices to enhance the applicability of the proposed signal, especially during an uncertain market condition. In fact, we convert the covariance matrices into adjacency matrices to topologically arrange the assets into Minimum Spanning Trees (MST). We show that the eigenvector centrality of the MST is nothing else than the eigenvector of the adjacency matrix. In fact, [Simon \(1991\)](#) argued that complex systems, such as financial markets, can be arranged in a natural hierarchy comprising nested substructures. Furthermore, in the graph theory literature, systemic risk shocks are referred to as abrupt changes in the density of cross-security connectivity ([Billio et al. 2012](#)), which is proxied by centrality measures such as eigenvector centrality.

To do so, in this study we use a time-series data of hourly, open, future prices of sentiment assets derived from Bloomberg during January 2008 to September 2021. Given the dataset of hourly asset returns, we first analyze the behavior of covariance matrices, minimum spanning trees and their associated sets of eigenvalues-eigenvectors during the two most recent financial setbacks, namely the 2008 GFC and the financial setback triggered by the 2020 COVID-19 health crisis. Supported by this empirical evidence, we draw attention to four metrics: the variance explained by the largest eigenvalue, the difference of variance explained by the largest and the five largest eigenvalues, the mean eigenvector centrality of the MST and its standard deviation.

We empirically validate the choice of the metrics with a two-fold framework. First, we apply hierarchical clustering – an unsupervised machine learning algorithm ([Raffinot 2017](#)) - on the four metrics showcasing their ability to neatly separate the observations recorded during periods of financial crisis from those occurred in non-crisis periods. Second, we filter a momentum investment strategy with the probability returned by a linear logistic model on the four metrics. Finally, we run an out-of-sample validation by generating ten synthetic datasets using a Variational Auto-Encoder (VAE) ([Kingma & Welling 2013](#)), a neural network algorithm having an output layer of the same dimensionality as the input one and whose objective is to reconstruct the distribution of the input data. Since the number of units in the hidden layer is typically fewer than those in the input layer, the reconstruction of the input data is inherently lossy, hence the output data is different from the sample data, yet sharing the same underlying distribution. The results - as measured by the Dunn index - point to the four metrics being able to separate high volatility periods from calm ones by means of the hierarchical clustering algorithm. This result is consistent across the ten out-of-sample synthetic datasets. Moreover, the filtered momentum investment strategy exhibits superior Sharpe ratios across all the sub-samples. As such, we contribute to the literature by showcasing the more prominent forward-looking view offered by eigenvalues-eigenvectors compared to traditional volatility measures.

The rest of the paper is organized as follows. Section 2 introduces the dataset used in this study. Section 3 draws the four eigenvalue-eigenvector based metrics from the empirical evidence offered by an analysis of the GFC

and COVID-19 crises. Section 4 describes the methodologies used in the validation framework. Section 5 illustrates the validation and out-of-sample results. Section 6 concludes and addresses future research.

2 Data

We use a time-series dataset of hourly, open, future prices of sentiment assets sourced from Bloomberg for the period January 2008 to September 2021. The assets are selected based on generally regarded empirical evidence referred to the overall market psychology affecting the demand and supply implied in investors' views (Verma & Soydemir 2009).

Table 1 lists the sentiment assets used in this study and the Ticker codes.

Table 1: List of the sentiment assets composing the datasets

Name	Ticker code
S&P500	ES
US treasury rate 2y	TU
US treasury rate 10y	TY
Gold	GC
EUR / USD	EURUSD
WTI	CL

Source: Bloomberg

We calculate the matrix of linear returns as:

$$r_t = \frac{r_t}{r_{t-1}} - 1 \quad \forall t = 1, \dots, T.$$

Successively, we estimate the realized covariance matrix through non-parametric realized covariances RC_t . (Shephard & Barndorff-Nielsen 2004) show that RC_t converges in probability to the quadratic variation of the price process. Thus, we calculate weekly realized covariances, RC_t , as the aggregation of cross-products of hourly returns, such that:

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

where N_t is the number of daily observations in the t -week, for $t = 1, \dots, T$ and r_τ 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 mouldable with any time series model.

Next, as base of the graph theory application, we build the adjacency matrix to represent the links between assets in the tree. Each entry indicates whether pairs of assets are adjacent. However, since clustering (which will

be applied in the first step of the validation process) requires a distance measure that satisfies the three axioms of a Euclidean metric, before forming the adjacency matrix we shall transform the covariances into metrics. As described in (Kayo & Kimura 2011), we first transform the covariances in correlations as $\rho_{i,j} = \sigma_{i,j}\sigma_j\sigma_i, \forall i, j$ and then we apply the following transformation:

$$d_{i,j} = 1 - \rho_{i,j}^2.$$

Moreover, the diagonal of the adjacency matrix is set equal to zero to avoid self-loops ². A non-zero entry in the adjacency matrix indicates the existence of a financial relation between pairs of assets wherein the strength of the link is measured by $d_{i,j}$. On the adjacency matrix, we estimate the minimum spanning tree (MST) ³ with the algorithm by (Kruskal 1956) ⁴.

The MST representation allows to quantify the influence structure among the assets by means of a centrality measure. As such, peripheral nodes that have limited impact on the dynamics of the network have lower centrality, while higher centrality nodes play a major role. 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. In our methodology, we use the eigenvector centrality measure, 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 fact, the eigenvector centrality ζ of each node x is given by the weighted average of the centrality values of its neighbours:

$$\zeta(x) := \frac{1}{\lambda} \sum_{(x,x^T)} W(x,x^T)\zeta(x^T)$$

where λ is a constant. In terms of the squared adjacency matrix A , we can re-write the eigenvector centrality as:

$$\zeta(x) := \frac{1}{\lambda} \sum_j A_{ij}\zeta(x_j)$$

where $j \in \mathbb{N}$ represents the neighbour nodes and in matrix form yields to:

$$\lambda\zeta = A\zeta$$

²a 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

³For a general introduction to graph theory, see (Bollobás 1998, Bollobás & Béla 2001).

⁴Starting from having a tree in each vertex, the Kruskal's algorithm then iterates throughout a process wherein any link with a minimum weight between the vertices is removed and if the removed link connects two different trees, then these are combined into a single tree.

from which it is immediate to see that ζ is the eigenvector of the adjacency matrix A and λ its eigenvalue. The solution is not uniquely determined since every pair of (λ, ζ) solves the equation. Nevertheless, for connected graphs, the Perron-Frobenius theorem (see [AS 1998](#), Chapter 8, Theorem 5) ensures that the eigenvector corresponding to the maximal eigenvalue contains all positive components. Hence, ζ is defined as the normalized dominant eigenvector of A .

3 Parallels between COVID-19 and GFC: a case study

In this section, we use a sub-sample of the entire dataset consisting of two non-continuous 52 weeks of observations concerning the October 2008 Global Financial Crisis and the 2020 financial setback triggered by the COVID-19 health crisis. For both periods, the start date is set at eight weeks before the day that the S&P500 recorded the minimum value during the crisis period, while the end date is fifty-two weeks after the start date. As a result, the two datasets span the periods first week of August 2008 - July 2009 and third week of January 2020 - January 2021 respectively.

Table 2 lists the descriptive statistics of the linear returns in the two sub-samples.

Table 2: Descriptive statistics of the returns in the two sub-samples

First week of August 2008 - July 2009					
Name	Mean	St. Dev.	Median	Skewness	Kurtosis
S&P500	10.28%	18%	25.22%	0.18	2.19
US Treasury 10y	0.02%	0.03%	0.07%	0.54	3.75
US Treasury 2y	0.01%	0%	0.01%	-2.11	12.44
Gold	0.14%	0.13%	0.17%	0.62	2.27
WTI	0.15%	-0.04%	0.46%	0.45	1.24
EURUSD	0.07%	0%	0.1%	0.52	0.81
Third week of January 2020 - January 2021					
Name	Mean	St. Dev.	Median	Skewness	Kurtosis
S&P500	8%	15.9%	13.1%	-0.83	1.08
US treasury 2y	0.2%	0.0%	1.3%	6.28	7.66
US treasury 10y	0.05%	0.0%	1.9%	11.74	
Gold	6.1%	2.6%	5.2%	0.04	0.66
Crude oil	-0.4%	-0.1%	36.8%	-9.87	15.92
EUR / USD	0.0%	0.0%	0.1%	0.57	0.5

Source: Bloomberg

Having calculated realized covariances and eigenvector centralities, we extract four metrics, namely: the variance explained by the largest eigenvalue and the difference of variance explained by the largest and the five largest eigenvalues from the realized covariance matrix as well as the mean eigenvector centrality and the standard deviation of the eigenvector centrality from the MST.

The variance explained by the largest eigenvalue is calculated as:

$$\frac{\lambda_1}{\sum_{j=1}^n \lambda_j} \tag{2}$$

where λ_1 is the largest of the n eigenvalues associated with the realized covariance matrix. Accordingly, the variance explained by the first five eigenvalues is calculated as:

$$\frac{\sum_{j=1}^5 \lambda_j}{\sum_{j=1}^n \lambda_j} \tag{3}$$

where $\sum_{j=1}^5 \lambda_j$ represents the sum of the first five eigenvalues and the difference in variance explained is obtained by simply subtracting (2) from (3). The mean and standard deviation of the eigenvector centrality as simply obtained as the arithmetic average and the sample standard deviation respectively.

Having obtained the four eigenvalue-eigenvector metrics from the realized covariance matrices and the minimum spanning trees, we now showcase their behavior to support their choice as risk signals. To do so, using the two sub-samples reflecting both volatile market observations and calm ones, we highlight a similar behavior during the 2008 GCF and the 2020 financial setback. In particular, figures 1 to 4 allow a comparison of the distribution density functions of the four metrics. The read shaded area in the densities marks the difference in distribution between calm weeks and volatile weeks, where we defined the volatile ones as the four weeks before and after the S&P 500 recording its lowest value in both sub-samples. To improve comparability, we horizontally stack the 2008 and 2020 datasets.

Figure 1

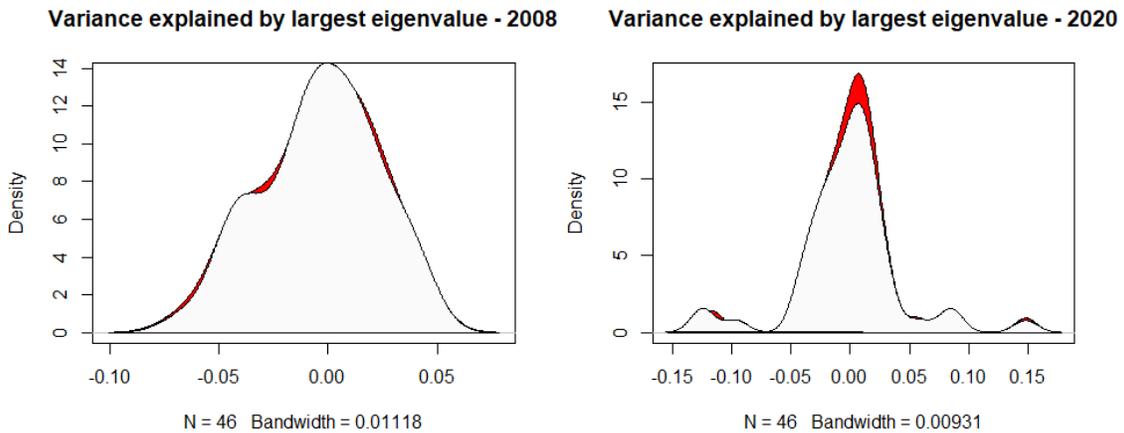


Figure 2

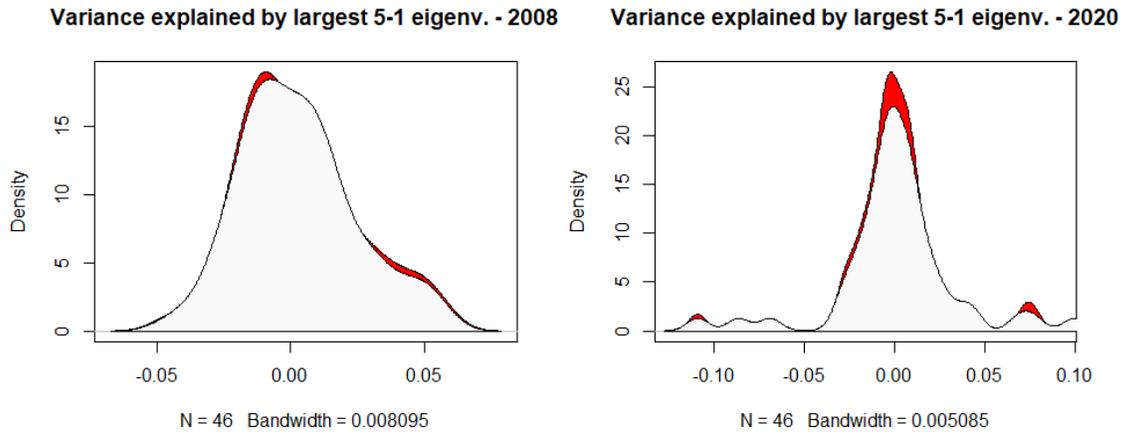


Figure 3

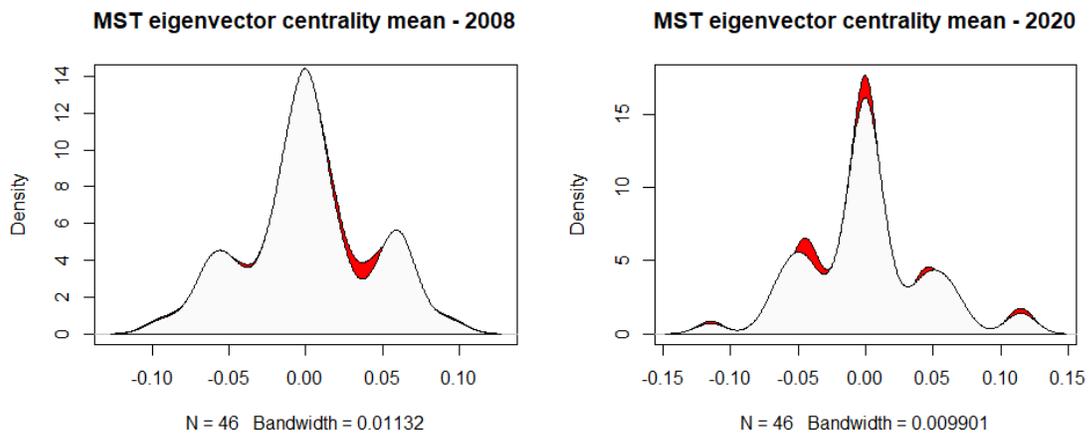
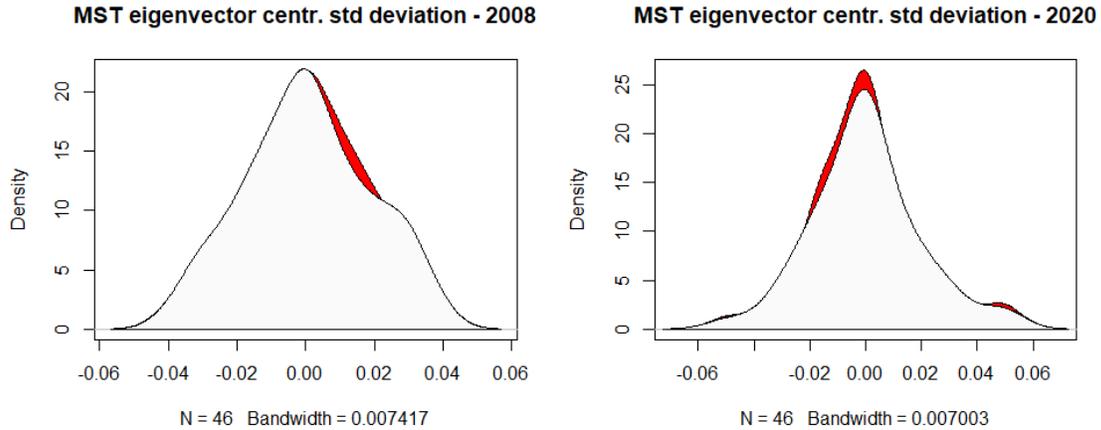


Figure 4

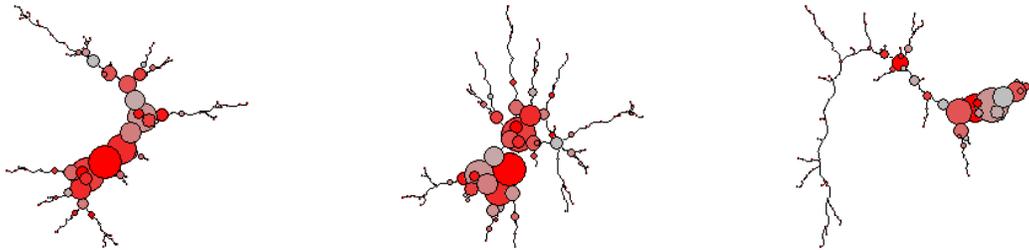


Starting from the variance explained by the largest eigenvalue and the difference of variance explained by the largest five eigenvalues and the largest eigenvalue, the density functions exhibit a similar jump in mean during the volatile weeks. Moreover, there is higher density in both tails of the distributions. As highlighted by [Laloux et al. \(1999\)](#), the largest eigenvalues of a covariance matrix reflect more systematic sources of risk. [Potters et al. \(2005\)](#) link the largest eigenvalues with systematic risk and the next three to four largest eigenvalues to industry momentum. The structural change in the shape of the distributions allows us to conclude that systematic risk becomes more predominant during highly volatile weeks. From a Capital Asset Pricing Model ([Fama & French 2004](#)) perspective, assets become more correlated within each other, resulting in a higher portion of variance being explained by movements of the market portfolio. From an investment strategy perspective, this implies that rebalancing the degree of correlation with the market portfolio during high volatile weeks is likely to shield from market downturns proxied by changes in the relative amount of variance explained by the largest eigenvalues. In fact, [Haugen et al. \(1991\)](#) reports a negative correlation between asset returns and volatility levels in equity and fixed income markets. From a risk management perspective, instead, this result suggests that risk measures such as Value at Risk should incorporate higher tail-risk when structural changes in covariance matrices are detected.

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 have shown how the eigenvector centrality is nothing else than the eigenvector of the adjacency matrix, which, in turn, is built out of covariance matrices. [Figure 5](#) allows a direct comparison of the minimum spanning tree generated in August 2008 during the 2008 GFC, the one generated in March 2020 during the financial setback triggered by COVID-19 and the one generated in a calm trading week (first week of

August 2021). The size of the vertices represent each asset’s eigenvector centrality, wherein each node represents a single asset. As it is visible, the MSTs at the onset of the two volatile periods present a similar denser structure where the connections between the assets is stronger. The interconnectivity among securities, together with the eigenvalues of the covariance matrix on which the adjacency matrix is built, bounced higher as a consequence of increased market risk. During the calm period, instead, the structure of the tree appears stretched on four 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 centre of the graph are capable to explain a larger portion of the comovements and result in higher eigenvector centralities.

Figure 5: The MST representation during two volatile weeks - August 2008 (left) and March 2020 (centre) - and a calm week - August 2021 (right)



4 Validation methodology

In this section we describe the methodologies used for the two-fold validation framework.

4.1 Clustering

Clustering analysis 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 to objects assigned to different clusters. As such, clusters are formed upon a degree of similarity often proxied by a distance measure. In this work, we employ the agglomerative hierarchical clustering via the AGNES algorithm described in (Kaufman & Rousseeuw 1990) as it allows us to find groupings with well-defined hierarchies within relationships.

This clustering algorithm starts with assuming that each object belongs to a different cluster, to then merge

at each of the $N - 1$ steps the two least dissimilar clusters in a single cluster by minimizing the inter-cluster dissimilarity. We consider the metrics $x_{t,j}$ for $t = 1, \dots, T$ and $j = 1, \dots, N$ extracted from the realized correlation matrixes as the base for measuring *intra-cluster* dissimilarities. We define the pairwise dissimilarities by means of the Manhattan function in which the distance between two points is the sum of the absolute differences of their Cartesian coordinates. As at each of the $N - 1$ steps one cluster is dropped and merged into a different one, we need to introduce a second measure of dissimilarity *between-clusters*, which is computed from the set of pairwise dissimilarities by means of a linkage function. We employed the Ward linkage function (Ward Jr 1963) which analyses the variance of the clusters via the sum of square distances from the centroids instead of measuring the distance directly.

The process of merging objects into clusters can be represented by a dendrogram, a rooted binary tree where the nodes represent clusters, the root node represents the whole dataset and the terminal nodes represent each individual object. Each non-terminal node has two daughter nodes, which represent the two objects that are merged into the same cluster in the case of the agglomerative approach. The final step is to cut the dendrogram at a particular height to partition the objects into a pre-defined number of clusters.

4.2 Momentum investment strategy

As second validation tool, we consider a daily momentum investment strategy on the SPY, a tradable ETF on the S&P500. We assume the SPY follows an Ornstein-Uhlenbeck process⁵ of the form:

$$dy_t = \kappa (\theta_t - y_{t-1}) dt + \sigma_t dW_t$$

where y_t is the asset return at week t , θ_t is the long-term mean, σ_t is the volatility, W_t is a Brownian motion and κ is a scalar parameter governing the speed of the reversion towards the mean θ_t . As such, the return generating process converges towards the long-term mean parameter θ over time.

We use the following Euler discretization for numerical implementation:

$$\hat{y}_t = \kappa \left(\hat{\theta}_t - y_{t-1} \right) \Delta t + \hat{\sigma}_t \sqrt{\Delta t} Z, \quad Z \sim N(0, 1).$$

We use the sample standard deviation of the returns on a rolling window of the last twenty weeks as σ_t estimate. θ_t is estimated via the moving average of the last fifty returns. Given a real value for the mean-reversion parameter κ , a Monte Carlo simulation sampling from the standard normal distribution estimates y_t . The confidence interval

⁵See (Finch 2004)

for the point estimate is given by:

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

where T is the number of time observations and α is arbitrarily set equal to 0.02. Finally, to find the optimal mean-reversion parameter κ , we employ a grid-search algorithm to find the parameter that minimizes the Mean Squared Error (MSE):

$$MSE = \frac{1}{T} \sum_{t=1}^T (\hat{y}_t - y_t)^2.$$

We found the optimal κ being equal to 0.85. The momentum investment strategy is defined as follows:

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

We introduce a filtered momentum strategy based on the four eigenvalues-eigenvectors based metrics by introducing the following linear logistic model:

$$\omega_t = \beta_{0,t} + \sum_{j=1}^4 \beta_{j,t} x_{j,t} + u_t$$

where $x_{j,t}, \forall t \in T = 1, \dots, n$ and $\forall j = 1, \dots, 4$ are the four metrics, $\beta_{0,t}$ and $\beta_{j,t}$ are the coefficients of the logistic regression and $u_t \sim N(0, 1)$. The logistic model outputs a probability $\omega_t \in [0, 1]$ capable to signal a structural change in the realized covariances or in the adjacency matrix. In the logistic model, to deal with multicollinearity, in each period we apply a stepwise variable selection methodology (Hocking 1976).

To filter the strategy with the four metrics, we scale the mean reversion parameter κ as follows:

$$\kappa_t = \kappa_0 [1 + \varepsilon \cdot I(\omega > 0.5)]$$

where $I(h_t = 1)$ is an index function which assumes a value equal to 1 when the probability ω_t resulting from the logistic model is higher than 0.5. In this case, the mean reversion scaling parameter ε is triggered. The optimal ε , obtained via grid search, is set at 1.2. One must notice that the mean reversion parameter, κ_t , becomes time-dependent in the filtered strategy.

4.3 Variational Autoencoder

In financial applications, analyzing the out-of-sample performance allows avoiding in-sample overfitting. For this reason, to validate out-of-sample the usage of the four metrics, we generate ten synthetic datasets via a Variational

Autoencoder (VAE) (Kingma & Welling 2013).

Consider a dataset $X = \{X^{(i)}\}_{i=1}^n$ composed of n *i.i.d.* samples coming from a random variable x . Let's assume that the data is 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 generate a value x^i from the conditional distribution $p_\theta(x|z)$. Let's 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 to mimic the hidden random process and generate a synthetic dataset that resembles the real data. To do so, we employ a Variational Autoencoder. Assume the prior over the latent variables to be a centred 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 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:

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

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 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:

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

where $\log p_\theta(x^i | z^{i,l})$ is a Gaussian fully connected neural network decoding term.

5 Validation results

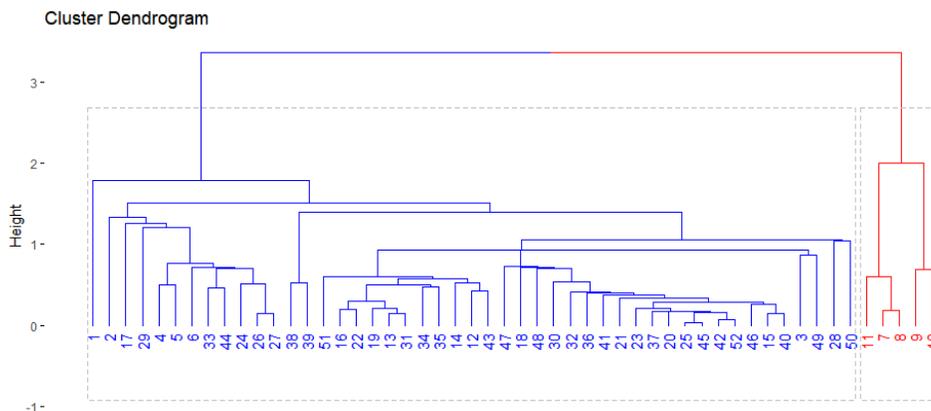
5.1 Clustering

Since clustering techniques are scale sensitive due to the reliance on distance functions, we first normalize the time series of the four metrics extracted from the covariance matrices and the MST. Afterwards, we assess the clustering

tendency of the data by calculating the Hopkins' statistic on the spatial randomness of data, which compares the distances between a sample of data points and their nearest neighbours to the distances from a sample of pseudo points and their nearest neighbours. Under the null hypothesis that the data is un-clusterable, the test statistic follows a beta distribution with both parameters equal to the number of points selected to sample the observations. The closer the Hopkins' statistics is to 1 and far above 0.5, the more the object have higher tendency to cluster. Our test shows a value of 0.91, hence rejecting the null hypothesis of un-clusterable data.

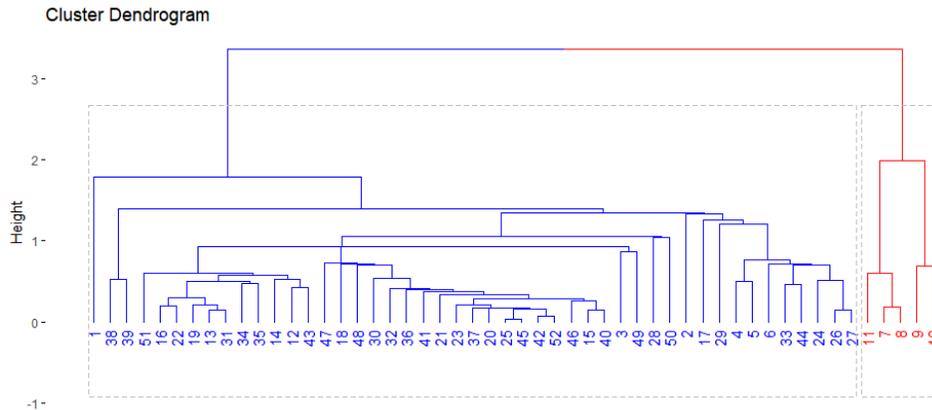
Having assessed the tendency to cluster, we run the AGNES algorithm (Kaufman & Rousseeuw 1990) on the four metrics. The results are graphically presented by the dendrograms in Figure 6 and Figure 7. In both cases, the weeks around the market plunge pinned at week 8 are grouped together in a neatly separated cluster. In January 2020, weeks of elevated volatility are recorded even before the COVID-19 was declared a pandemic, amid concerns over the overvaluation of the equity markets, the China-USA trade war and the first news of the health emergency counting the first infections in Europe and USA. Moreover, as statistical measure to validate the clusters, we use the Dunn index (Dunn 1973, 1974), a metric used to evaluate clustering algorithms calculated as the ratio between the minimum between-cluster distance and the maximal within-cluster distance. A higher value implies a better clustering. We normalize the Dunn Index to let it assume values between 0 and 1. Our procedures obtain a Dunn index of 0.93 and 0.85 respectively for the GFC and the COVID-19 sub-samples. This allows us to conclude that the metrics are able to capture a different hierarchy in response to structural changes in the covariance matrices.

Figure 6: The dendrogram produced by AGNES on the 2008 dataset



Note: the dendrogram shows that the AGNES algorithm is able to separate the weeks around the one when the S&P 500 recorded the lowest value during the 2008 GFC (in red) from the others. The 2008 sub-sample is drawn from the original dataset so that the first observation is set at 8 weeks before the lowest value touched by the S&P 500.

Figure 7: The dendrogram produced by AGNES on the 2020 dataset



Note: as for the 2008 GFC, the 2020 dendrogram also shows that the AGNES algorithm is able to separate the weeks around the one when the S&P 500 recorded the lowest value during the 2020 financial setback triggered by the COVID-19 health crisis (in red).

5.2 Momentum investment strategy

In this section, we compare the performance of the filtered momentum investment strategy ⁶ with the unfiltered version.

Table 3 reports performance metrics for both strategies.

Table 3: Performance evaluation of the two investment strategies

	Filtered strategy	Unfiltered strategy
Sharpe ratio	1.25	1.03
Sortino ratio	0.93	0.75
Kurtosis	3.1	4.7

The evaluation metrics show that superior performance is linked to the usage of the eigenvalue-eigenvector based metrics for filtering the momentum investment strategy. In particular, the improvement is mostly in terms of risk strategy. In fact, the filtering mechanism allows to quickly react to structural changes in covariance matrices and minimum spanning trees. As $\epsilon > 1$, the filtered strategy is capable to increase the speed of mean reversion when systemic risk is detected by the sheering of the covariance matrix. As such, the filtered strategy positions itself closer to the longer term mean, which is more likely to proxy the fair value of the underlying asset. In fact, the unfiltered strategy incurred in large losses during the COVID-19 and 2008 Global Financial Crisis. During calm markets, on the other hand, the filtered strategy allows to rise with the tide in the same way as the unfiltered

⁶We are not analyzing the momentum strategy using the classical winners minus losers as it may exhibit sasonality especially with monthly or quarterly frequencies.

strategy. As a result, the filtered strategy exhibits larger Sharpe and Sortino ratios. Moreover, the conservative shifting to quicker convergence to the mean is also mirrored in lower kurtosis.

5.3 Out-of-sample analysis

In financial applications, focusing on the out-of-sample performance allows to avoid in-sample overfitting. For this reason, we re-run the two empirical applications on ten synthetic datasets. The VAE model specification is fine-tuned via grid search. We use two hidden layers with 5 hidden units activated by means of the leaky 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 is targeting the representation error via the consistency of the mean squared error.

The main idea of this section is to ensure consistency across various datasets and discard the hypothesis that the results are sample-dependent. Table 4 reports the Dunn index and the Sharpe ratio across the ten synthetic datasets. As it is visible, results are consistent across the simulations. In particular, the normalized Dunn index averages 0.91, hinting that the AGNES algorithm on the four eigenvalue-eigenvector metrics results in small within-cluster variance and well separated (low between cluster variance) clusters. The Sharpe ratio averages 1.21 above the 0.98 of the unfiltered one.

Table 4: Evaluation of clustering and filtered momentum strategy performance across the ten out-of-sample synthetic datasets

Sample	Dunn index	Sharpe filtered	Sharpe unfiltered
S1	0.86	1.20	1.08
S2	0.94	1.22	0.95
S3	0.84	1.21	0.95
S4	0.88	1.23	1.04
S5	0.94	1.26	0.99
S6	0.98	1.16	0.98
S7	0.84	1.20	0.96
S8	0.95	1.30	0.96
S9	0.91	1.18	0.96
S10	0.88	1.19	0.97
Average	0.91	1.21	0.98

6 Conclusions and future work

The aim of this research is to offer empirical evidence on the usage of eigenvalues and eigenvectors as risk signal able to capture exogenous changes in financial markets' sentiment.

Using a datasets composed of hourly observations of sentiment assets' returns, we first calculate the realized covariance matrices and arrange the securities in a minimum spanning tree. Successively, we split the dataset in two sub-samples with each one being 52 weeks long and starting eight weeks before the deepest point recorded by the S&P500 during the 2008 Global Financial Crisis and the 2020 financial setback triggered by the COVID-19 health crisis respectively. We offer a case study on the behavior of four eigenvalue-eigenvector based metrics extracted from the covariance matrices and the minimum spanning trees, namely the variance explained by the largest eigenvalue, the difference of variance explained by the five largest and largest eigenvalue, the mean eigenvector centrality of the minimum spanning tree and its standard deviation.

We analyse the distributions of the four metrics during both setbacks. The distribution of the variance explained by the largest eigenvalue and by the difference in variance explained by the five largest and largest eigenvalue exhibit jumps during the crisis suggesting that the systematic source of risk becomes more predominant during these periods. In terms of the mean and standard deviation of the eigenvector centrality of the minimum spanning tree, their distributions also exhibit higher means during weeks of crises. In this case too, the reason is related to the fact that systematic risk explains a larger portion of asset movements and empirical correlations tend to jump higher. The minimum spanning tree topological arrangement, moreover, allows us to conclude that during volatile markets securities concentrate in the core of the tree wherein fewer assets are capable to explain co-movements among several more peripheral assets. Hence, in both cases, changes in eigenvalues-eigenvectors are likely triggered by changes in systematic risk.

To validate the choice of the metrics, we first show that an agglomerative hierarchical clustering algorithm based on the metrics can isolate the weeks of crisis from those of calm trading weeks. Second, we show that improved performances are linked to filtering a momentum investment strategy with the probability of being in a volatile market returned by a linear logistic model on the metrics. Finally, we validate out-of-sample our findings on ten synthetic datasets generated using a Variational Auto-Encoder neural network.

Future work should focus on exploring the effects of other eigenvalue-eigenvector based metrics. Moreover, we suggest that major research should be conducted on the role of eigenvalues and eigenvectors to measure financial risk, such as Value at Risk, portfolio construction and asset pricing.

References

- Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2013), Financial risk measurement for financial risk management, *in* ‘Handbook of the Economics of Finance’, Vol. 2, Elsevier, pp. 1127–1220.
- AS, B. B. (1998), ‘Modern graph theory. volume 184 of graduate texts in mathematics’, *Springer, New York NY* **2**, 4.
- Avellaneda, M. & Lee, J.-H. (2010), ‘Statistical arbitrage in the us equities market’, *Quantitative Finance* **10**(7), 761–782.
- Belke, A., Dubova, I. & Osowski, T. (2018), ‘Policy uncertainty and international financial markets: the case of brexit’, *Applied Economics* **50**(34-35), 3752–3770.
- Billio, M., Getmansky, M., Lo, A. W. & Pelizzon, L. (2012), ‘Econometric measures of connectedness and systemic risk in the finance and insurance sectors’, *Journal of Financial Economics* **104**(3), 535–559.
- Bollobás, B. (1998), Random graphs, *in* ‘Modern graph theory’, Springer, pp. 215–252.
- Bollobás, B. & Béla, B. (2001), *Random graphs*, number 73, Cambridge university press.
- Ding, W., Levine, R., Lin, C. & Xie, W. (2020), Corporate immunity to the covid-19 pandemic, Technical report, National Bureau of Economic Research.
- Duan, J.-C. & Zhang, W. (2014), ‘Forward-looking market risk premium’, *Management Science* **60**(2), 521–538.
- Dunn, J. C. (1973), ‘A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters’.
- Dunn, J. C. (1974), ‘Well-separated clusters and optimal fuzzy partitions’, *Journal of cybernetics* **4**(1), 95–104.
- Fama, E. F. & French, K. R. (2004), ‘The capital asset pricing model: Theory and evidence’, *Journal of economic perspectives* **18**(3), 25–46.
- Finch, S. (2004), ‘Ornstein-uhlenbeck process’.
- Gençay, R., Selçuk, F. & Ulugülyağci, A. (2003), ‘High volatility, thick tails and extreme value theory in value-at-risk estimation’, *Insurance: Mathematics and Economics* **33**(2), 337–356.
- Glaeser, E., Huang, W., Ma, Y. & Shleifer, A. (2017), ‘A real estate boom with chinese characteristics’, *Journal of Economic Perspectives* **31**(1), 93–116.

- Haugen, R. A., Talmor, E. & Torous, W. N. (1991), ‘The effect of volatility changes on the level of stock prices and subsequent expected returns’, *The Journal of Finance* **46**(3), 985–1007.
- Hocking, R. R. (1976), ‘A biometrics invited paper. the analysis and selection of variables in linear regression’, *Biometrics* pp. 1–49.
- Hohlmeier, M. & Fahrholz, C. (2018), ‘The impact of brexit on financial markets—taking stock’, *International Journal of Financial Studies* **6**(3), 65.
- Jiang, Y., Zheng, L. & Wang, J. (2021), ‘Research on external financial risk measurement of china real estate’, *International Journal of Finance & Economics* **26**(4), 5472–5484.
- Kaufman, L. & Rousseeuw, P. J. (1990), *Finding Groups in Data: An Introduction to Cluster Analysis.*, John Wiley.
- Kayo, E. K. & Kimura, H. (2011), ‘Hierarchical determinants of capital structure’, *Journal of banking & finance* **35**(2), 358–371.
- Kingma, D. P. & Welling, M. (2013), ‘Auto-encoding variational bayes’, *arXiv preprint arXiv:1312.6114* .
- Kousenidis, D. V., Ladas, A. C. & Negakis, C. I. (2013), ‘The effects of the european debt crisis on earnings quality’, *International Review of Financial Analysis* **30**, 351–362.
- Kruskal, J. B. (1956), ‘On the shortest spanning subtree of a graph and the traveling salesman problem’, *Proceedings of the American Mathematical Society* **7**, 48–50.
- Laloux, L., Cizeau, P., Bouchaud, J.-P. & Potters, M. (1999), ‘Noise dressing of financial correlation matrices’, *Physical review letters* **83**(7), 1467.
- Papailias, F. & Thomakos, D. (2017), ‘Exssa: Ssa-based reconstruction of time series via exponential smoothing of covariance eigenvalues’, *International Journal of Forecasting* **33**(1), 214–229.
- Potters, M., Bouchaud, J.-P. & Laloux, L. (2005), ‘Financial applications of random matrix theory: Old laces and new pieces’, *arXiv preprint physics/0507111* .
- Raffinot, T. (2017), ‘Hierarchical clustering-based asset allocation’, *The Journal of Portfolio Management* **44**(2), 89–99.
- Shephard, N. & Barndorff-Nielsen, O. (2004), ‘Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics’, *Econometrica* **72**, 885–925.

- Simon, H. A. (1991), The architecture of complexity, *in* 'Facets of systems science', Springer, pp. 457–476.
- Verma, R. & Soydemir, G. (2009), 'The impact of individual and institutional investor sentiment on the market price of risk', *The Quarterly Review of Economics and Finance* **49**(3), 1129–1145.
- Ward Jr, J. H. (1963), 'Hierarchical grouping to optimize an objective function', *Journal of the American statistical association* **58**(301), 236–244.
- Xu, Y. & Lien, D. (2020), 'Dynamic exchange rate dependences: The effect of the us-china trade war', *Journal of International Financial Markets, Institutions and Money* **68**, 101238.