



An eigenvalue distribution derived 'Stability Measure' for evaluating Minimum Variance portfolios

Smyth, W., & Broby, D. (2022). An eigenvalue distribution derived 'Stability Measure' for evaluating Minimum Variance portfolios. *Quantitative Finance*, 23(3), 521-537. Advance online publication. <https://doi.org/10.1080/14697688.2022.2149420>

[Link to publication record in Ulster University Research Portal](#)

Published in:
Quantitative Finance

Publication Status:
Published online: 02/12/2022

DOI:
[10.1080/14697688.2022.2149420](https://doi.org/10.1080/14697688.2022.2149420)

Document Version
Publisher's PDF, also known as Version of record

Document Licence:
CC BY-NC-ND

General rights

The copyright and moral rights to the output are retained by the output author(s), unless otherwise stated by the document licence.

Unless otherwise stated, users are permitted to download a copy of the output for personal study or non-commercial research and are permitted to freely distribute the URL of the output. They are not permitted to alter, reproduce, distribute or make any commercial use of the output without obtaining the permission of the author(s).

If the document is licenced under Creative Commons, the rights of users of the documents can be found at <https://creativecommons.org/share-your-work/licenses/>.

Take down policy

The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk



An eigenvalue distribution derived 'Stability Measure' for evaluating Minimum Variance portfolios

William Smyth & Daniel Broby

To cite this article: William Smyth & Daniel Broby (2022): An eigenvalue distribution derived 'Stability Measure' for evaluating Minimum Variance portfolios, Quantitative Finance, DOI: [10.1080/14697688.2022.2149420](https://doi.org/10.1080/14697688.2022.2149420)

To link to this article: <https://doi.org/10.1080/14697688.2022.2149420>



© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 02 Dec 2022.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

An eigenvalue distribution derived ‘Stability Measure’ for evaluating Minimum Variance portfolios

WILLIAM SMYTH * and DANIEL BROBY 

Ulster University Business School, Cathedral Quarter, Belfast, Northern Ireland, UK

(Received 28 March 2022; accepted 14 November 2022; published online 2 December 2022)

The Minimum Variance portfolio is subject to varying degrees of stability and robustness. We, therefore, propose a theoretical measure of its stability relative to a Marchenko–Pastur derived random correlation matrix. We demonstrate its practical use on the S&P 400, the S&P 500, the S&P 600 and the Russell 1000. Using historic market data from 2002 to 2021, we perform an optimisation on the empirical correlation matrix eigenvalue distribution to determine the implied variance $\nu(t)$ for the underlying data-generating process. Through monitoring its change over time $\Delta\nu(t)$, we provide a Stability Measure for the Minimum Variance portfolio and thereby help researchers measure changes to estimation risk and manage rebalancing regimes.

Keywords: Stability Measure; Minimum Variance portfolio; Modern portfolio theory; Covariance matrix; Marchenko–Pastur

1. Introduction

The Minimum Variance portfolio is that set of assets that has the least risk. It is constructed by estimating an asset return covariance matrix. This is easier than constructing a Mean Variance portfolio which requires a forecast of mean returns. We observe that properties of the correlation matrix eigenvalue distribution can be used to measure the stability of the Minimum Variance portfolio. We derive an implied process variance value $\nu(t)$ from these properties which can be monitored over time. It is, in effect, an eigenvalue distribution derived ‘signal-to-noise’ ratio at time t . We use the Marchenko–Pastur theorem to separate eigenvalues for ‘signal’ from those for ‘noise’. Focusing on eigenvalues for the noise we perform a similarity optimisation between the empirical distribution and the analytic Marchenko–Pastur distribution. We demonstrate how the result of this optimisation, ν , can be used as a Minimum Variance portfolio evaluation tool by observing how it changes over time. This has practical usage in determining rebalancing decisions.

The Minimum Variance portfolio lies on the Efficient Frontier and was first identified by Haugen and Baker (1991). A number of academics have subsequently conducted studies on it (Merton 1980, Chopra and Ziemba 2013). They observe

that its construction is subject to sampling error. Klein and Bawa (1976) point out that such estimation risk has important considerations for optimal portfolio choice. As a result, the asset weights are noisy and unstable over time. That said, it has an important quality that allows us to evaluate its stability. Unlike the optimal portfolio, its covariance matrix is the only unknown parameter. The former requires an estimate of the mean of the risky asset returns, which is more difficult than the estimation of the covariance matrix (Merton 1980). Although the Minimum Variance portfolio has less estimation risk, it still has some instability. As such, we argue that changes in implied process variance ν provide a useful measure of stability over time. Our Stability Measure is therefore based on the concept of portfolio-specific stationarity, and determined by monitoring the temporal properties of the signal-to-noise ratio of a portfolio’s empirical correlation matrix.

The Stability Measure builds on the seminal work of Marchenko and Pastur (1967). They derived an analytic form for the probability density function of the eigenvalue distribution of a random covariance matrix. This is formed from a data matrix whose elements are independent identically distributed random variables drawn from a zero-mean process with finite variance. In the case that the data-generating process has a unit variance, the covariance and correlation matrices are one and the same. The Marchenko–Pastur eigenvalue distribution

*Corresponding author. Email: w.smyth@ulster.ac.uk

is therefore well suited to be used as a basis for the calculation of ν .

We demonstrate the efficacy of the Stability Measure using historic market data and the associated empirical correlation matrix eigenvalue distribution, based on the S&P 400, the S&P 500, the S&P 600 and the Russell 1000 over the 20-year period 2002–2021 inclusive. Our results show the evolution of $\nu(t)$ over time $\Delta\nu(t)$. They identify periods of, and breaks in, portfolio-specific stationarity. Our findings have implications for the efficacy of portfolio management generally. They have methodological implications for scholars investigating Modern Portfolio Theory, the Efficient Frontier and the estimation error. They also have practical relevance and implications for the management of Minimum Variance index portfolios.

2. Minimum Variance portfolio

The Minimum Variance portfolio, unlike the Mean Variance portfolio, is estimated without reference to forecast returns. It amounts to the solution of a constrained optimisation which can be represented as follows,

$$\min_w \sigma_{p,w}^2 = w' C w \quad (1)$$

such that

$$w' 1 = 1. \quad (2)$$

$\sigma_{p,w}^2$ represents the variance of portfolio p for a given weights vector w , and C is the covariance matrix of the N portfolio components. The Minimum Variance Portfolio is therefore the portfolio on the efficient frontier which has the lowest possible risk.

Generating a random covariance matrix consistent with the Marchenko–Pastur theorem involves populating a data matrix with a series of independent identically distributed random variables drawn from a zero-mean random process of finite variance. By contrast, the daily stock returns which populate an empirical data matrix (T time-series observations for each of N stocks) are not independent (diversification depends upon this) or identically distributed (the underlying data generation process is unknown and may not even be well defined). Additionally, the derivation of the Marchenko–Pastur theorem incorporates the use of finite fourth-order moments which may not exist for the kind of fatter-tailed distributions habitually encountered when dealing with asset returns. Having said this, financial data matrices generally have very low signal-to-noise ratios rendering them potentially suitable for comparison with Marchenko–Pastur derived matrices. Also, through Figure 2(c,d) we demonstrate that for an appropriate choice of the Marchenko–Pastur quality parameter q (defined in Section 4), the analytic probability density function is a good fit for fatter-tailed distributions. We also note increased instances of the application of this theory in the financial literature (Laloux *et al.* 2000, Sharifi *et al.* 2004, Bun *et al.* 2017, Lopez de Prado 2020). However, our Stability Measure method is the first such approach to use Marchenko–Pastur in this way in portfolio evaluation.

The analysis of portfolio constituents using the Marchenko–Pastur approach delivers an asymptotic probability density function for the eigenvalue distribution of a random correlation matrix. This property is central in our generation of a signal-to-noise ratio for empirical stock data. The Marchenko–Pastur theorem facilitates a threshold separation between those eigenvalues associated with noise and those associated with signal. We apply this threshold to empirical data and focus exclusively on the distribution of eigenvalues associated with noise. We then compare the latter discrete distribution with the analytic Marchenko–Pastur probability density function using a kernel density fit. The rationale for this is to minimise the distance between the two using the sum of squared differences as the distance metric. This amounts to an optimisation over the Marchenko–Pastur process variance parameter σ^2 where the distance-minimising value σ_{opt}^2 may be interpreted as the implied variance of the data-generating process for the stock components comprising the portfolio. We use the term implied variance because, as we are optimising over the M–P process variance parameter σ^2 , the optimum value tells us the equivalent M–P process variance value which would produce the closest possible theoretical distribution to our empirical distribution. Hence, we can interpret the optimum value as an implied variance value for the underlying data-generating process we have passed to the process. σ_{opt}^2 is the parameter we refer to as ν .

Our results evaluating $\Delta\nu(t)$ show that, during periods of portfolio stability (which we make think of a period during which the portfolio-specific data-generating process is stationary), ν changes in a predictable way. This observation supports the case for monitoring Minimum Variance portfolio stability in this way.

3. Estimation errors of the mean and covariance matrix

The literature on estimation errors in the Efficient Frontier is extensive. In line with our thinking, Rao (1971) suggested that variance should be minimised to obtain the best unbiased estimators. Modern Portfolio Theory suggests a portfolio manager can construct an optimal portfolio w from the covariance matrix W and vector α of its expected excess returns. The stocks in the investment universe are assumed to be stationary and Gaussian. The covariance matrix W is estimated from the historical time series. This makes the matrix noisy. The optimisation process magnifies this noise. The traditional way to approach this is to increase the number of observations T (Ledoit and Wolf 2003).

Gennotte (1986) pointed out that estimation error becomes more of a problem in the case when distributional properties are not observable. Bun, Bun *et al.* (2017) provide further support for the use of random matrices with an emphasis on the Marchenko–Pastur theorem. Shepard (2009) argues that by using factors, N/T is substituted by K/T , where K is the number eigenvalues. This makes the derived random portfolio more robust, as K is a lower number than N .

DeMiguel and Nogales (2009) addressed the Minimum Variance portfolio with a single non-linear solution to reduce estimation error. This was due to its instability. This resulted

in their portfolio weights being more stable than those of the traditional portfolios. We suggest the Stability Measure could be used to further investigate this.

It was shown by Ledoit and Wolf (2003) that the covariance matrix of two existing estimators, the sample covariance matrix and single-index covariance matrix, can be addressed by shrinkage in the following way,

$$\Sigma_{\text{shrink}} = \alpha F + (1 - \alpha)\Sigma_{\text{SCM}}, \quad 0 \leq \alpha \leq 1 \quad (3)$$

where the convex combination Σ_{shrink} has a shrinkage target F and shrinkage estimator Σ_{SCM} , with α being the shrinkage intensity.

Statistical methods are required to deal with the large dimensional datasets now routinely being generated in many fields including finance. In many studies, the rationale for using data in matrix form is to explain the joint dynamics of a collection of N observables. In this paper that dynamic is the connection between the daily returns of the various stocks that make up an index portfolio such as the S&P 400, S&P 500, S&P 600 and Russell 1000. A very natural way to quantify the similarities between N observables is an $N \times N$ correlation matrix. Its eigenvalues and eigenvectors may be used to represent the most important dynamical modes of the system. In principal component analysis (PCA), this amounts to those linear combinations of the original variables with the greatest variance.

The objective is to estimate the population correlation matrix whose elements are the correlations between pairs of observables. In our case, pairs of stocks. The major difficulty in practice is that this matrix is rarely computable because the underlying data generation process is unknown and is often what the study is attempting to determine. Empirically, the approach used is often to collect a large number T of realisations of the process for each of the N observables. For instance, a large number of daily returns for all the components of the S&P 500. For sufficiently large T , it appears natural to use the sample correlation matrix to estimate the population correlation matrix. In the case where $N \ll T$, it is well known that the sample correlation matrix E (almost surely) converges to the population correlation matrix C (Van der Vaart 2000).

A common measure of estimation risk in high-dimensional problems such as Minimum Variance portfolio selection for index portfolios is given by the quotient $\text{Tr}E^{-1}/\text{Tr}C^{-1}$, where Tr is the matrix trace operator representing the sum of the elements along the main diagonal and E^{-1} is the inverse of E . This quotient takes a value very close to unity when T is sufficiently large for a given N , i.e. when $q = T/N \rightarrow \infty$. However, when the number of observables N is sufficiently large relative to T so that q is not large, this relationship becomes (El Karoui 2010)

$$\frac{\text{Tr}E^{-1}}{\text{Tr}C^{-1}} = \frac{q}{q-1} \quad (4)$$

which holds for a wide class of processes. As can be seen, the out-of-sample risk $\text{Tr}E^{-1}$ can far exceed the true optimal risk $\text{Tr}C^{-1}$ when q is not very large, diverging as $q \rightarrow 1$. So, the estimation process becomes problematic if T is not very large

compared to N . This is precisely the region of parameter space that defines many financial matrices. In particular, correlation matrices for index portfolios. In the case of the S&P 500 for instance, with $N = 500$, setting $T = 2500$ would correspond to ten years of daily data. At $q = 5$ this clearly does not satisfy the requirement that $N \ll T$, and yet 10 years of daily data represents a time period over which it would seem unrealistic to make assumptions of stationarity.

The first result on the behaviour of sample covariance matrices in the large dimensional limit came from the seminal work of Marchenko and Pastur (1967) in which they obtained a self-consistent equation for the spectrum of E given C as $N \rightarrow \infty$. The influence of the quality ratio q appears explicitly. Anderson (1963) showed that in the classical limit $T \rightarrow \infty$ and N fixed that the sample eigenvalues converge to the population eigenvalues, a result recovered by the Marchenko–Pastur formula in the limit $q \rightarrow \infty$. However, when $q = O(1)$, the same formula shows that all the sample eigenvalues become noisy estimators of the true (population) counterparts no matter how large T is. Such a situation is referred to as the curse of dimensionality. More precisely, the distortion of the spectrum of E compared to the true one becomes increasingly pronounced as $q \rightarrow 1^+$ (see Figure 1(a)). This phenomenon may be explained by observing that when the sample size T is very large, each individual coefficient of the covariance matrix C can be estimated with negligible error (provided one can assume that C itself does not vary with time (i.e. the underlying data-generating process is stationary)). However, if N is also large and of the order of T , as is the case in many situations, the sample estimator E becomes inadmissible (Bun *et al.* 2017). The large number of noisy variables creates systematic errors in the computation of the correlation matrix eigenvalues.

4. The role of the Marchenko–Pastur theorem in determining portfolio stability

The application of random matrix theory to the study of financial correlation matrices dates back to Laloux *et al.* (1999) and Plerou *et al.* (1999). Since then, the Marchenko–Pastur theorem (Marchenko and Pastur 1967) has been used in several investment and portfolio optimisation studies, including: Sharifi *et al.* (2004), Urama *et al.* (2017) and Bruno *et al.* (2018).

For instance, Bruno *et al.* (2018) use a Marchenko–Pastur adjustment to reduce the effect of portfolio similarity between different fund managers, reducing what they call crowding by between 14% and 60%. Sharifi *et al.* (2004) approach the problem of estimation with specific regard to portfolio stability. They determine the noise percentage present in the correlation matrix, C , and propose a technique for filtering C which takes account of stability more precisely than a standard cleaning technique. These publications present results that are consistent with and supportive of our approach. We are not aware of any studies which implement a rolling monitor for portfolio stability in the high-dimension regime in the way we have presented in this paper.




Figure 1. (a) Marchenko–Pastur eigenvalue probability density function for $\sigma^2 = 1$ and a range of values of the quality parameter q . (b) as (a) but for fixed $q = 1.5$ and a range of values of σ^2 .

The Marchenko–Pastur theorem can be expressed as follows. Let X be a matrix with dimensions $T \times N$ whose elements are observations from a zero-mean random process with variance $\sigma^2 < \infty$. In all that follows the quotient $q = T/N$ is an important parameter. In the limit $N, T \rightarrow \infty$ such that $1 < q < \infty$, matrix

$$M = \frac{1}{T} X'X \quad (5)$$

has eigenvalues whose distribution converges to the analytic Marchenko–Pastur probability density function

$$p(\lambda) = \begin{cases} q \frac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{2\pi\lambda\sigma^2}, & \lambda \in [\lambda_-, \lambda_+] \\ 0, & \lambda \notin [\lambda_-, \lambda_+] \end{cases} \quad (6)$$

where $\lambda_{\pm} = \sigma^2(1 \pm \sqrt{q})^2$. Figure 1(a) shows the form this probability density function takes for a variety of values of q for fixed $\sigma^2 = 1$. In the limit that q gets very large, the eigenvalue distribution gets extremely narrow as shown in the image indicating a distribution where all eigenvalues are 1. In terms of population dynamics, this is equivalent to the form of the correlation matrix approaching that of the identity matrix. Within the context of portfolio analysis, the focus of this paper, the $T \times N$ data matrix for which the correlation matrix is determined is interpreted as N representing the number of stocks or portfolio components and T representing the length of a time series of stock returns. In this study N will account for the number of components in an index portfolio such as the S&P 500 and so will be of the order of 500, while the returns data will be daily returns and so T , the number of observations or rows in the data matrix, will represent a series of consecutive trading days.

With reference to Figure 1(a), as we will be operating with fixed values for N , q is varied by varying T . When we later deal with actual stock data we will, unless otherwise stated,

be using $q = 1.5$. Figure 1(b) depicts the effect of altering the process variance for fixed $q = 1.5$.

Key to our approach is the need to quantify the similarity between the analytic Marchenko–Pastur pdf and the discrete distribution of eigenvalues coming from the eigen-decomposition of an empirical correlation matrix. To that end, we utilise a Gaussian kernel density fit on the empirical eigenvalue distribution and then measure the distance between this KDE and the analytic pdf.

Based on our central premise that the Marchenko–Pastur is a good fit for the distribution of those eigenvalues associated with noise, distinct from those associated with signal, we perform an optimisation over the Marchenko–Pastur process variance parameter σ^2 . The value $\nu = \sigma_{\text{opt}}^2$ which maximises the similarity (equivalently which minimises the distance) between the analytic pdf and the empirical KDE is the key parameter in our portfolio stability management proposal. As we move through time ν can be measured at a suitable frequency. In this paper, we measure it on a weekly basis. Portfolio stability is determined by the change that occurs in ν . We find that portfolio stability is characterised by periods of time during which weekly relative change in ν occurs within a well-defined boundary ($\pm 5\%$).

5. Random data (no signal)

In order to provide some context for our analysis we begin by covering some basic facts and features of the situation. We start by considering a basic process whereby the $T \times N$ data matrix is populated using TN independent draws from a known stationary data-generating process such as the standard normal distribution $N(0, 1)$. Next, we standardise the empirical data so that each column of the data matrix has zero mean and unit variance. This will be standard procedure for any empirical data matrix we utilise in this study. The

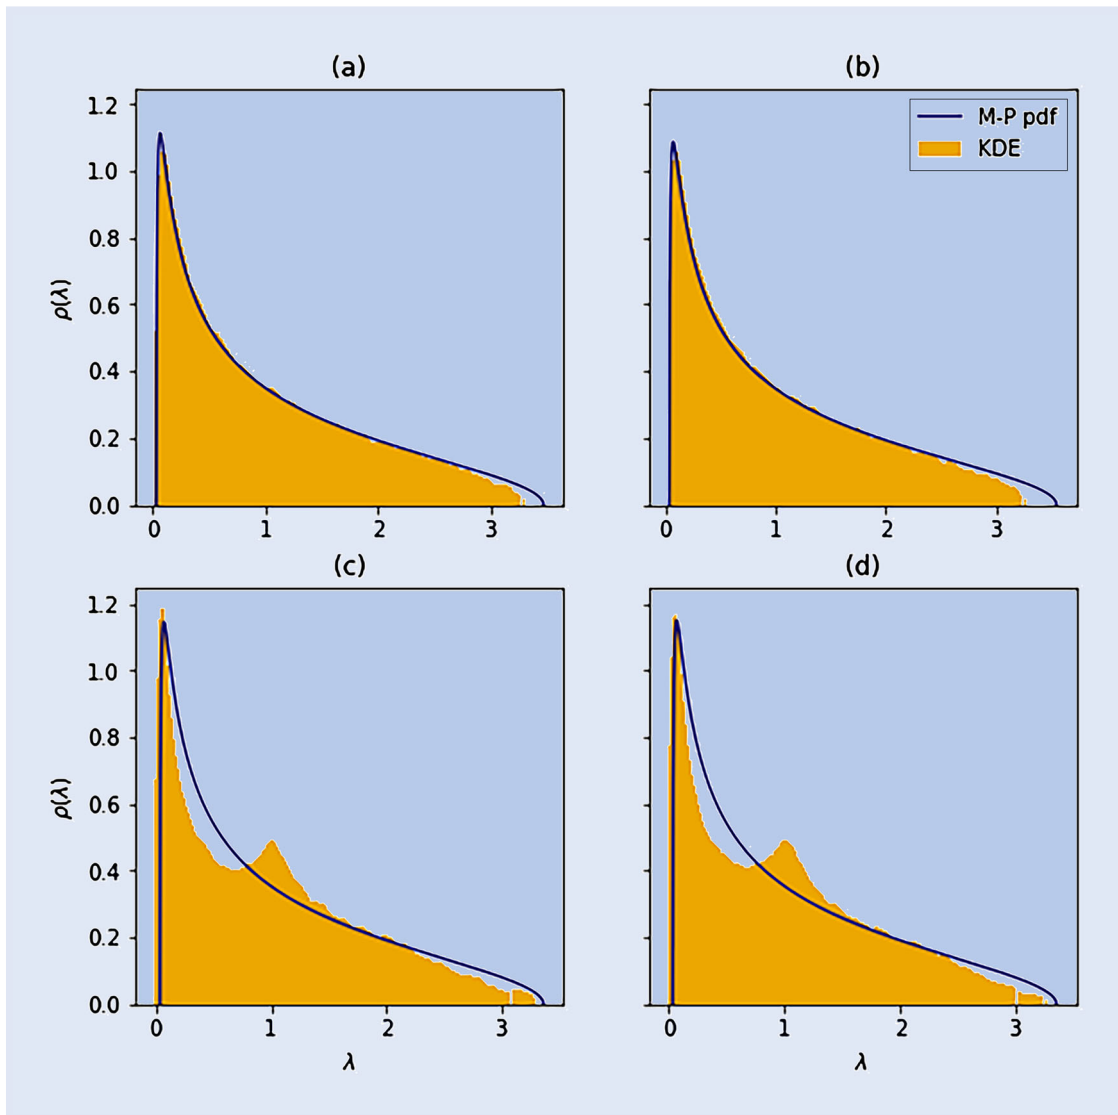


Figure 2. Marchenko–Pastur eigenvalue probability density function and Gaussian KDE fit to simulated data for a selection of known stationary processes and $q = 1.5$. (a) Gaussian ($\mu = 0, \sigma^2 = 1$), (b) Uniform ($\mu = 0, \sigma^2 = 1$), (c) Cauchy ($x_0 = 0, \gamma = 1$), (d) t -dist ($\nu = 1$).

correlation matrix of this column-standardised matrix is then determined and eigen-decomposition performed. Where relevant, the Marchenko–Pastur threshold for separating eigenvalues associated with noise from those associated with the signal is utilised and the kernel density fit is applied to those eigenvalues (the bulk component) that lie within the Marchenko–Pastur boundaries for noise. The fit is optimised by tuning the Marchenko–Pastur process variance parameter in order to minimise the distance between analytic pdf and empirical KDE. In the case of random data with no signal such as in Figures 2 and 3, the optimisation occurs for $\sigma_{opt}^2 = 1$.

This is as it should be since we are dealing with purely random (no signal) data generated from a stationary zero-mean parent process of unit variance (in the case of Gaussian and Uniform generated data). Figure 2 shows the result of having generated data ($N = 500, T = 750$, hence $q = 1.5$) from each of four distinct probability distributions: a standard normal distribution; a uniform distribution of zero mean and unit variance, a Cauchy distribution centred on zero with $\lambda = 1$ and a t -distribution with one degree of freedom. The latter two are included as examples of distributions with fatter tails

than Gaussian processes. The only difference between Figures 2 and 3 is that, in Figure 3, q has been increased to $q = 5$. The main point to note is that in Figure 2, i.e. for $q = 1.5$, the process is shown to be much more agnostic to the source data-generating process. As we move to Figure 3, and hence to $q = 5$, it becomes clear that the process is not appropriate for Cauchy and t -distribution sourced data. This drives the choice of q when we come to work with market data and it has the additional convenience that small q makes the calibration of our stability management system easier. It should be noted that values of $q \approx 1$ will produce erroneous behaviour but $q = 1.5$ produces robust results whilst allowing us to avail of practical design advantages.

Finally, to illustrate what we mean by stationarity, consider the following idealisation. Suppose the underlying data-generating process is Gaussian $N(\mu, \sigma^2)$. When we refer to a period of stationarity, we mean that this process remains the underlying data generating process throughout that period of time, with fixed μ and fixed σ^2 . The outworking of this scenario is depicted in Figure 13(a). There is clear fluctuation (sampling effect) but the fluctuation is confined within

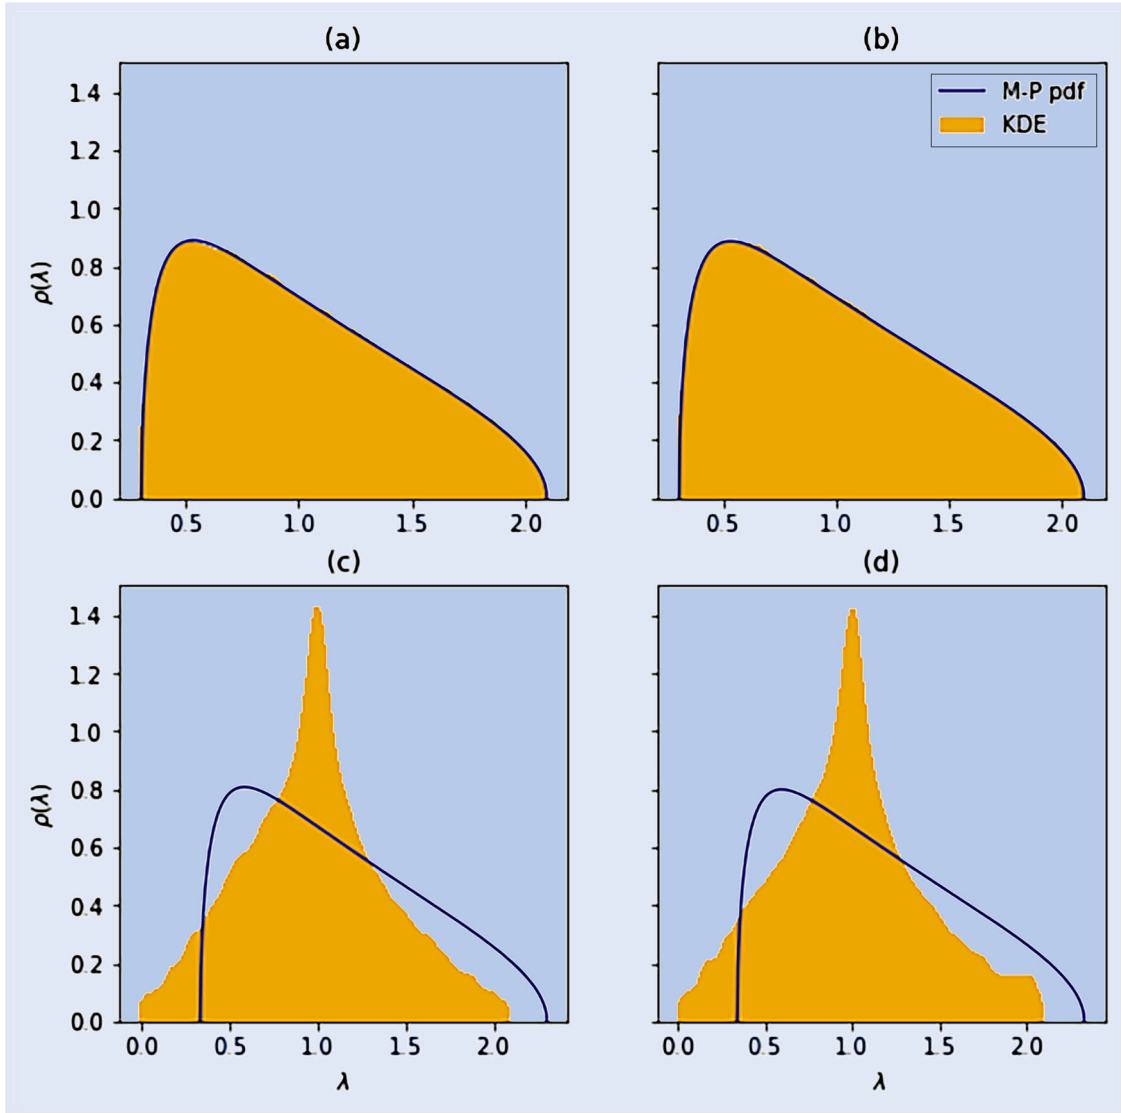


Figure 3. As Figure 2 but with $q = 5$.

a consistent vertical band, indicative of what we interpret as stationarity and with a bandwidth typical of this particular situation ($N(0, 1)$ and $q = 1.5$). Similar effects are shown in the other frames of Figure 13 for other known distributions. This is the link between our stability measure and stationarity. Stationarity, or a break in it, is manifesting through the magnitude of the fluctuation in $\nu(t)$. Figure 13(a–d) are examples of behaviour for defined statistical processes, stationary over the entire time period. Contrast this with Figure 12(a–d) for empirical financial data. A more full explanation of the features of Figures 12 and 13 will be provided in subsequent sections.

6. Random data (with signal)

In the following section, we apply this procedure to real market data and extend the treatment to provide a mechanism for detecting potential non-stationarity. In order to bridge the conceptual gap between considering data generated by a known stationary random process, such as we have just considered,

and data that is generated by stock markets, we consider something of a contrived go-between. Specifically, we generate a correlation matrix for data containing both noise and signal. Signal is injected into a percentage of the channels; we can tune the number of channels containing the signal and the strength of the signal. The injected signal will be deliberately weak, in-keeping with the actual market data we consider subsequently. The purely random component comes from a standard Gaussian source. Column-wise standardisation takes place as before to ensure zero mean and unit variance for each column.

Before presenting results associated with this section, it is instructive to elaborate on what we mean by injecting a signal into channels. To generate the signal carrying covariance matrix, we first generate a factor matrix F of dimensions $N \times f$ where f is the number of channels (we may think of channels as factors) into which we wish to inject the signal and N is the number of assets ($f < N$). The columns of this matrix are populated by random draws from a standard Gaussian, $N(0, 1)$. From this, we generate a covariance matrix FF' which, by virtue of $f < N$, will be rank-deficient. Next, we generate a diagonal matrix of dimensions $N \times N$ where

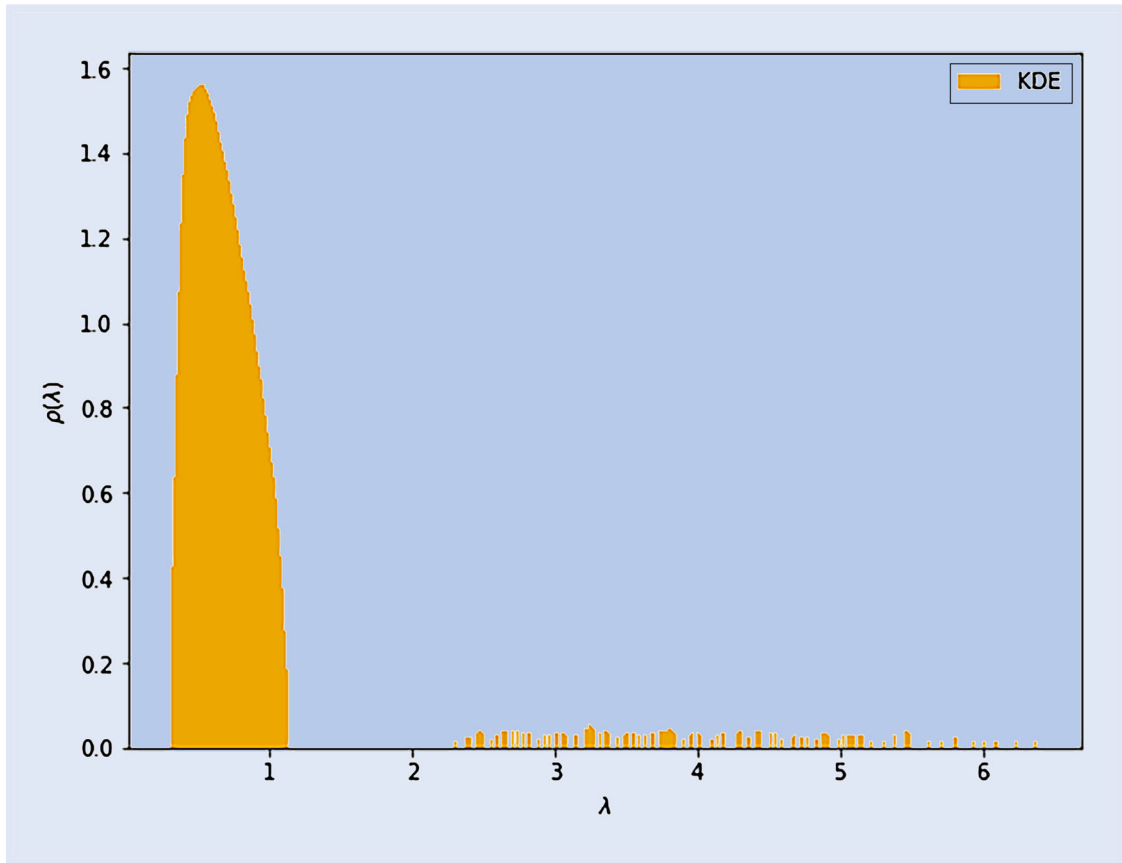


Figure 4. Empirical correlation matrix eigenvalue distribution for a standard Gaussian data-generating process with signal injected into 10% of channels.




Figure 5. Marchenko–Pastur eigenvalue probability density function and Gaussian KDE fit to eigenvalues located within the boundaries for noise $\lambda_- < \lambda < \lambda_+$, for the process described in Figure 4.




Figure 6. As Figure 4 except that signal is injected into 20% of channels.

the main diagonal is populated with random draws from a Uniform distribution, $U(0, 1)$. Summing these two matrices produces a full-rank covariance matrix. This is our signal matrix S . Finally, we add S to a random covariance matrix Q (the latter is generated in the same way as in the previous section) as follows

$$\alpha Q + (1 - \alpha)S \quad (7)$$

producing a noise-with-signal-embedded covariance matrix which is passed to our stability monitoring procedure. The strength of the signal is controlled through α and the number of channels into which the signal has been injected is determined by f . To simulate financial data, we use a very weak signal, setting $\alpha = 0.995$.

Figure 4 shows the result of applying the KDE to the discrete eigenvalue distribution of the associated correlation matrix in the case where the signal has been injected into 10% of the channels (f is 10% of N). The distribution has the characteristic feature of a bulk of eigenvalues separated from a series of discrete eigenvalues representing the signal injected as factors at the outset of the process.

The Marchenko–Pastur probability density function describes the distribution of eigenvalues associated with noise and so we will only consider eigenvalues in the main bulk component. Strictly speaking, we consider only those eigenvalues λ falling below the threshold separating noise and signal,

$$\lambda < \lambda_+ = \sigma^2 (1 + \sqrt{q})^2 \quad (8)$$

We perform the optimisation on this sea of eigenvalues. In other words, we find ν the optimum value of the Marchenko–Pastur process variance parameter,

$$\sigma_{\text{opt}}^2 \equiv \nu \quad (9)$$

The resulting Stability Measure (12) maximises the similarity between the analytic pdf and empirical KDE for the given value of q . The distance metric we utilise in this study is simply the sum of squared differences. Figure 5 visualises the result, showing the fit between the analytic pdf and the empirical KDE. This produces a value $\nu = 0.675$. We interpret this as a measure of signal to noise in the original data. The lower this value, the higher the ratio of signal-to-noise in the data. ν is bounded above by 1 since we are performing an optimisation relative to pure random noise, generated by a zero-mean unit variance process. An optimisation resulting in a value of $\nu = 1$ would be tantamount to observing data which contained no signal and was simply noise. These are the cases we considered in Figures 2 and 3.

Figure 6 depicts the eigenvalue distribution for a case analogous to that in Figure 4 with the distinction that noise has been injected into 20% of the channels. In this case, the increased levels of signal produced a lower optimum parameter value of $\nu = 0.504$. Consistent with higher levels of signal-to-noise.

We do not attempt to infer much meaning from the absolute value of ν but later we shall see that monitoring the change in this parameter over time provides valuable insight into the potential for detecting a stationarity break in the data generation process. This has implications for any prediction




Figure 7. Stability measure ν from 2006 to 2021 for a range of market indices with $q = 1.5$. (a) S&P 400, (b) S&P 500, (c) S&P 600, (d) Russell 1000.

due to be made on the assumption of stationarity in the data-generating process. In particular, this has connotations for portfolio rebalancing, most notably for Minimum Variance portfolio rebalancing since a correlation matrix, which is the core component in our analysis, is the only thing required for the allocation of weights in a Minimum Variance portfolio.

Clearly, any collection of stocks does not become a portfolio until weights are allocated yet our stability monitoring utilises the correlation matrix and its eigen-decomposition without alluding to weights. However, a collection of index stocks such as those of the S&P 500 are all large-cap stocks and the correlations between these stocks are carrying information about what may be happening in the market of specific relevance to this genre of stocks. The correlation matrix formed from S&P 500 stock returns is carrying the stability flag for large-cap stocks and therefore to some extent, for the stability of any portfolio constructed from a weighted combination of these stocks. An efficient frontier can be constructed by varying the weight allocations and there are, of course, some famous portfolios such as the maximum Sharpe Ratio portfolio and the Minimum Variance portfolio.

Our analysis is less concerned with the various features of this efficient frontier and more concerned with making predictions on how this frontier might experience perturbation over time, particularly the Minimum Variance point of the efficient frontier. Moreover, it is focused on identifying whether or not predictions are being made on the basis of data from both (before and after) sides of a break in stationarity; data that were generated, at least in part, by a process that is no longer in place. Our analysis will have maximum efficacy in the Minimum Variance region of the efficient frontier because this portfolio is determined by alluding only to correlations. Hence the stability of this point on the efficient frontier will be most effectively managed through effective monitoring of the stability of the correlation matrix which is precisely what our analysis does.

Furthermore, we would not utilise the S&P 500 correlation matrix to make stability predictions for an S&P 400 Minimum Variance portfolio. There will inevitably be information common across both correlation matrices since there will be changes in the market that will affect both categories of stock (large cap and small cap), but the correlations between stocks




Figure 8. Distribution of values of ν for corresponding frames in Figure 7.

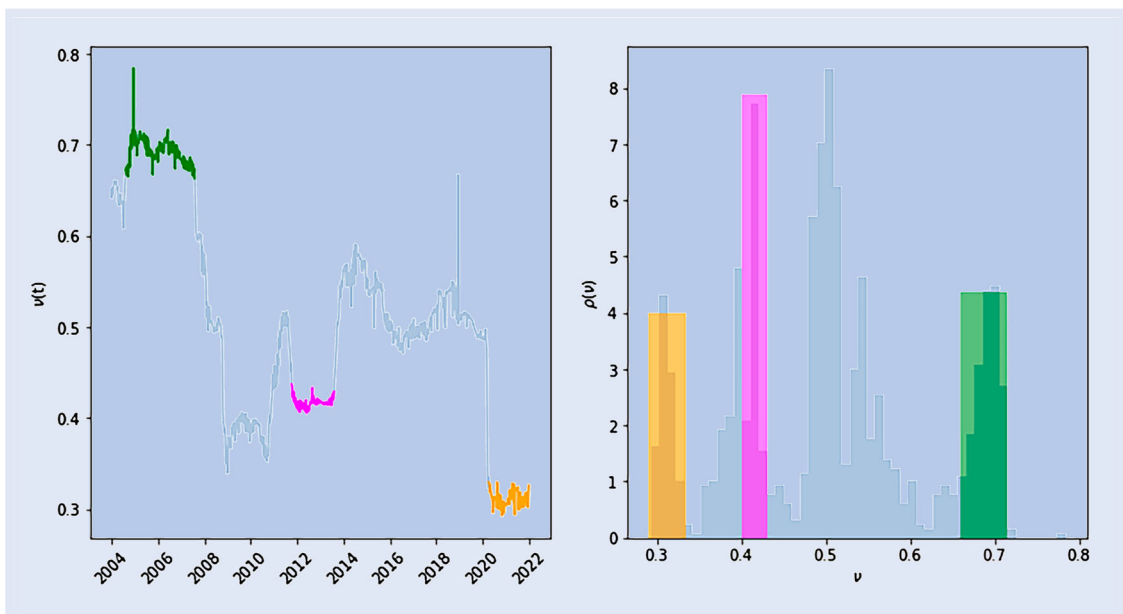


Figure 9. Comparison of frames (a) from Figures 7 and 8 corresponding to the S&P 400 index.




Figure 10. Distribution of values of weekly percentage change in ν for zones identified in Figure 9. (a) Green zone, (b) magenta zone, (c) orange zone, (d) all other zones.

of a particular genre will carry the most pertinent information on index portfolios made up of stocks exclusively from that genre.

A practical feature of portfolio management for index investment is the need to rebalance, to account for stocks that may have entered or left the index since the previous rebalancing. In our analysis, we adopt a short stability monitoring window, one week, which is not a practical frequency for the purposes of rebalancing. However, our results show that we do not need to take account of stocks leaving and entering the index over a given period.

Our results hold for the correlations which exist and evolve in only that core subset of stocks which are present over the rebalancing period, indeed over much longer periods than are typically associated with rebalancing. To be concrete, in the extended period (2002–2021) over which we collected daily returns data, our results hold for the core set of stocks which were ever-present throughout this period. In the case of the S&P 500 that equates to 389 stocks. This collection represents

the large-cap staple, carrying all vital correlation-derived stability information for the genre of large-cap stocks. This makes weekly (or more frequent) stability monitoring eminently practical and readily amenable to automation. Indeed, whilst not reported upon in this study, we have demonstrated that our results hold even for a subset of this subset but since automation means a collection for all these stocks is straightforward there is no real merit in shortening it further. However, it should be noted that since we fix $q = 1.5$ in our analysis, shortening the stock list from 500 to 389, perhaps further to 300, means that the observation window can be shortened from 750 to 450 days.

To perform perfect calibration of the monitoring process we would ideally have access to a starting time-window where stationarity is known to be true. Clearly, it is easier to ensure this for a shorter window. However, offsetting that advantage is the need to ensure that the time-window into which we look back is sufficiently long to provide sample data truly representative of the population process given that each column of




Figure 11. Distribution of values of weekly percentage change in for known stationary processes and $q = 1.5$. (a) Gaussian ($\mu = 0, \sigma^2 = 1$), (b) Uniform ($\mu = 0, \sigma^2 = 1$), (c) Cauchy ($x_0 = 0, \gamma = 1$), (d) t -dist ($\nu = 1$).

the empirical data matrix is standardised as part of a zero-mean unit variance approach. Note, this window is an initial calibration window which then becomes a rolling window. Monitoring frequency is the step size that the rolling window takes each time it moves forward. In this study, we use a step size of one day as we move through the daily returns data.

Of course, the situation can change if we have access to hourly data where rows in the data matrix would represent trading hours rather than trading days. However, it is clear that in the case of noisier hourly data a much longer window (in relative terms) would be required to provide an empirical sample size large enough to be representative of true population dynamics.

At a rate of six trading hours per trading day, it would have to be more than six times longer before it would result in a longer window in absolute time. Of course, it should be borne in mind that as we increase the length of this observation window, given a fixed number of stocks, we are altering the q -value of the Marchenko–Pastur probability density function. Recall, at the beginning of this paper, we demonstrated how

larger values of q are not a good choice for some kinds of random processes, particularly fatter-tailed distributions, which typically underpin stock returns generation.

7. Market data

We backtest this approach on market data to ascertain if it uncovers known features with regard to how the market unfolded over recent years and if these features are meaningfully evidenced by our optimisation. And, of course, to see if our approach lends new insight into how we might consider the concept of stationarity or stability as it pertains to specific investment vehicles such as index portfolios.

The results show our approach accurately bears out major market changes over the past two decades. Major events such as the global financial crisis and the Coronavirus pandemic impacted all the index portfolios we consider, with portfolio-specific nuances.




Figure 12. Evolution of $\Delta\nu(t)$ between 2006 and 2021 for (a) S&P 400, (b) S&P 500, (c) S&P 600, (d) Russell 1000.

Outside these critical events the various indices we consider evolved in different ways and thus stationarity (stability) which exists for one of these index portfolios at a particular point in time and for a particular period of time may not simultaneously exist for one or more of the others. This underpins the central finding of this study that the stability of a Minimum Variance index portfolio can be ascertained through the empirical correlation matrix of core stocks making up that particular index. The stability analysis applies only to the particular index portfolio in question. We cannot infer the stability of one index portfolio from the stability of another.

Through the concept of portfolio stability, we are able to enhance portfolio risk management in a straightforward way. This approach has the greatest applicability to the management of Minimum Variance index portfolios. In forming a Minimum Variance portfolio, the estimation process is limited to estimating the population correlation matrix of the constituent components, whereas non-Minimum Variance portfolios have additional estimation requirements.

Hence, when we assess the stability of the empirical correlation matrix, we are assessing the stability of the associated

Minimum Variance portfolio. At the same time, if we are able to distinguish a portfolio-specific period of stability from a period of instability then we know that any statistical analysis (hypothesis testing, regressions, statistical inference) using data across this break in stationarity is potentially erroneous. Being able to detect a break in stationarity provides the opportunity to take mitigation against such a break negatively impacting the efficacy of predictive models.

Figure 7 shows the evolution of ν (the signal-to-noise ratio of the empirical correlation matrix) over time for the S&P 400, 500, 600 and the Russell 1000 indices over the period from the start of 2006 to the end of 2021. The figure has some interesting features. There are clear similarities and clear differences across indices. As indicated above the critical events of 2008 and 2020 are readily apparent across all indices evidenced by large almost instant changes in ν .

Before and after these critical events ν experiences periods of relative stability when there appears to be stationarity in what we may think of as a portfolio-specific data-generating process. This is different for each of the indices, they do not all experience stationarity at the same time. Our claim for




Figure 13. Evolution of $\Delta v(t)$ between 2006 and 2021 (simulated) for (a) Gaussian ($\mu = 0, \sigma^2 = 1$), (b) Uniform ($\mu = 0, \sigma^2 = 1$), (c) Cauchy ($x_0 = 0, \gamma = 1$), (d) t -dist ($\nu = 1$).

stationarity is empirically supported, and that will be conveyed in some of the visualisations to come. We find the behaviour of v during periods of stationarity is consistent with comparable features of known stationary random processes.

Figure 8 depicts the distribution of v for each of the indices under consideration over the same time period used in Figure 7. The image is characterised by discrete poles around which we have a spread of values. Each pole marks the centre value of a period of stationarity, and hence Minimum Variance portfolio stability, whereas transitions between poles mark breaks in stationarity.

Figure 9 illustrates this in greater detail by focusing on the S&P 400 index and selecting three regions of stationarity. Colour coding is used to link the periods from left to right frames. It is instructive to consider how v changes during each of the periods identified in this image.

To that end, Figure 10 depicts the distribution of weekly percentage change in v . For each of the regions identified as periods of stationarity, the modal weekly change is overwhelmingly zero. Also, in each of these regions, the range of values is contained within $\pm 4\%$ (the green zone has a single

anomalous daily value which can be identified as the spike in Figure 9 causing the scale on the x -axis to broaden). The fourth frame contains the distribution of values in all other regions combined. The mean weekly percentage change for all other regions is still zero but less overwhelmingly so. Indeed, removing the spike at zero for each frame results in values for the mean and variance of the distributions shown which are twice as large for all other regions as for any of the coloured regions whose first and second moments are very similar.

In Figure 11 we display the results of having carried out a structurally identical simulation using data generated from known stationary distributions; specifically, a standard Gaussian, a standard Uniform distribution, and a Cauchy and t -distribution, with parameters as shown, to account for fatter-tailed phenomena. These data were generated by populating an original data matrix of 400 columns and 5000 rows. We then used the same size rolling window as had been used for S&P 400 data [$N = 400$ columns and $T = 600$ rows so that $q = 1.5$]. This rolling window traversed 4400 steps to simulate transiting through 5000 trading days as was the case for

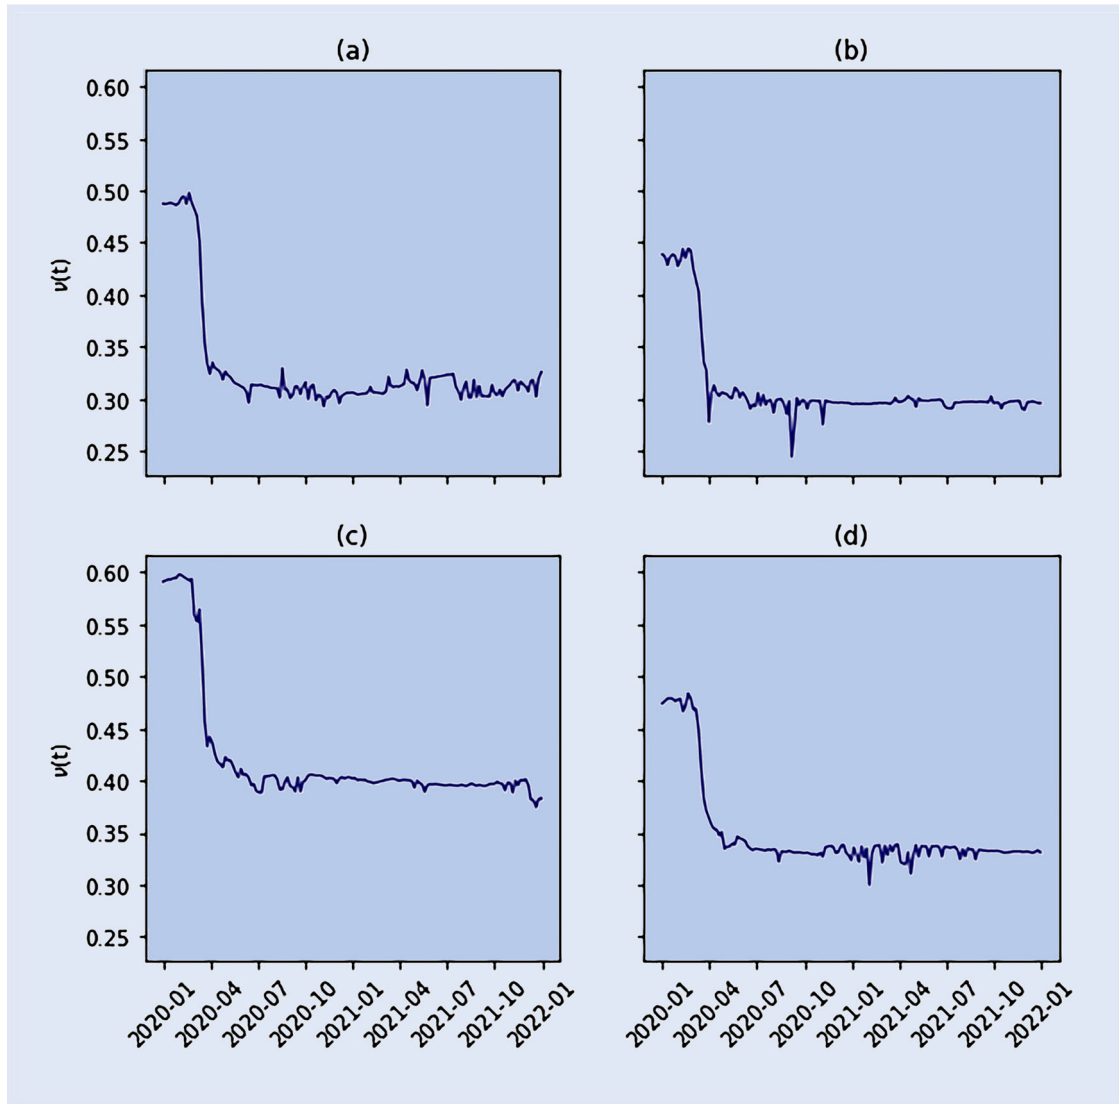


Figure 14. $v(t)$ post-pandemic for $q = 1.5$ (a) S&P 400, (b) S&P 500, (c) S&P 600, (d) Russell 1000.

our market data. Note the width of the distribution in each of the frames in Figure 11. Results for Gaussian and Uniform generated stationary data are contained within $\pm 2\%$ of the centre which is much narrower than for the market distribution in the periods of stationarity. However, for instances where the original data were generated by fatter-tailed stationary distributions, the outcome closely matches actual S&P 400 market results.

To provide an overview to facilitate a comparison between actual market data analysis and simulated data from known stationary processes we include Figures 12 and 13. Each frame depicts the weekly percentage change in v . Figure 12 depicts the backtest outcome for four market indices whereas Figure 13 depicts simulation using known stationary processes. The most obvious featural distinction between the figures is the pulsing behaviour as we move through time for the market data in Figure 12. This pulsing effect, which is completely absent for the perpetually stationary processes of Figure 13, is conveying periods of stationarity and breaks in stationarity over time.

8. Discussion

We highlight the usefulness of the Stability Measure in portfolio rebalancing by reference to an example. The most recent clear break in stationarity occurred as a result of the global Coronavirus pandemic, the effect of which is clearly evidenced in Figure 14 which depicts the evolution of v through the calendar years 2020–2021.

The pandemic produced an almost instant halt to a global economic and conventional financial activity and hence we have something approaching a discontinuity manifesting itself in the same way and at the same time in all four index graphs. Thereafter, however, we see rapid equilibration and the formation of a new period of stationarity across all the indices. As illustrated earlier, with a focus on the S&P 400, a period of stationarity is typified by a distribution of weekly percentage changes in the values of v centred on zero and having a spread of typically $\pm 5\%$. Figure 15 supports this for all four indices throughout these calendar years. This provides a sound basis for analysis commencing in calendar year 2022 in terms of

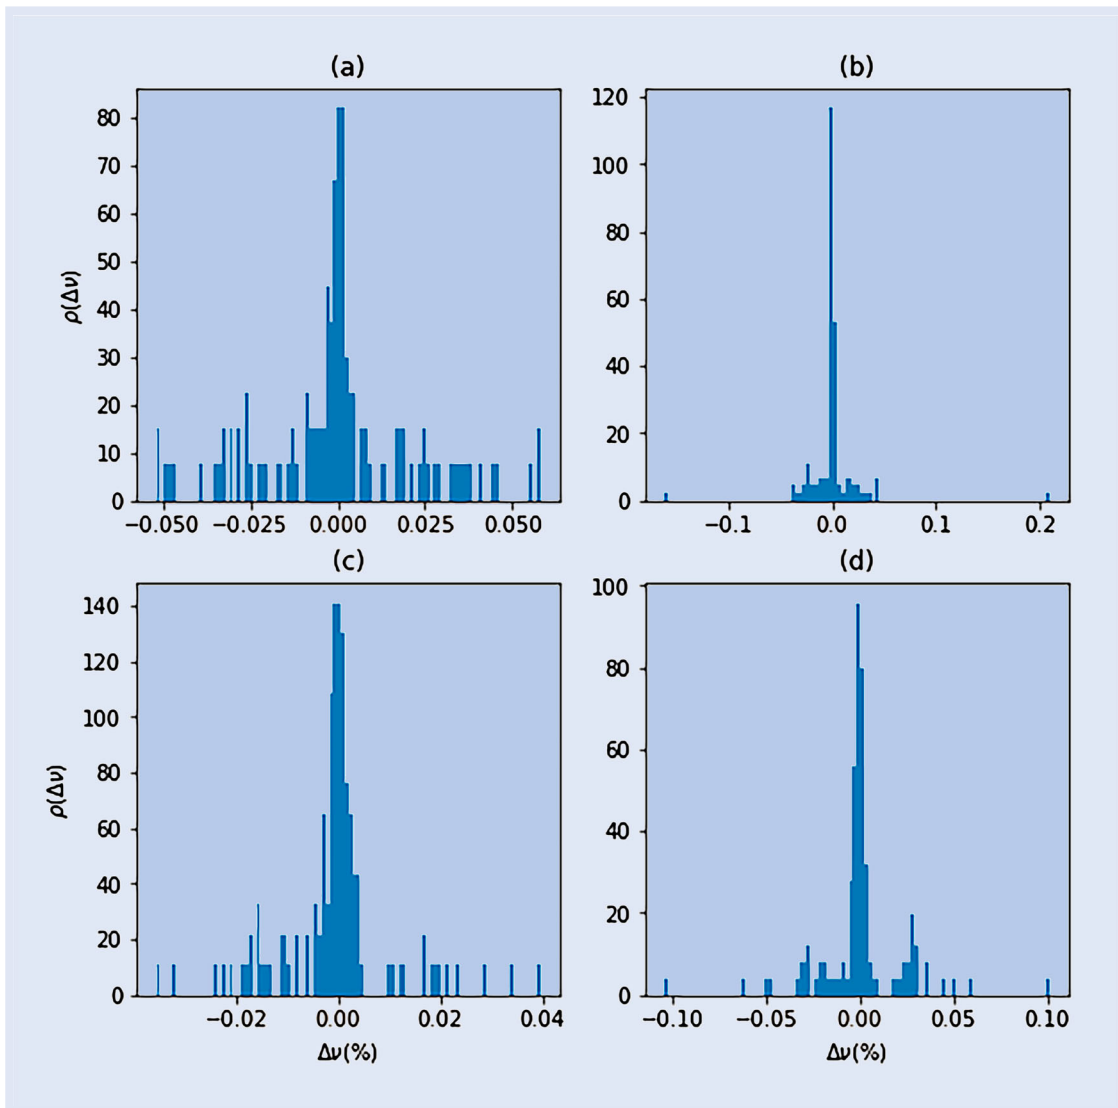


Figure 15. Distribution of values of post-pandemic weekly percentage change in ν for $q = 1.5$, (a) S&P 400, (b) S&P 500, (c) S&P 600, (d) Russell 1000.

monitoring stationarity and hence Minimum Variance index portfolio stability on a rolling weekly basis. This can, in turn, inform decisions around rebalancing which is typically done on a much less frequent basis. Weekly values of ν can be determined for the index portfolio of interest and the empirical percentage change distribution is incorporated as a basis for assessing ongoing stability.

9. Conclusion

We present a mechanism through which the stability of the Minimum Variance portfolio on the efficient frontier can be measured. We call this the Stability Measure. It relies on Marchenko–Pastur theory; specifically, on the distribution of those eigenvalues associated with noise as distinct from those associated with the signal. The measure arises as the product of an optimisation to maximise the similarity between the discrete distribution of the noise-related eigenvalues of an empirical correlation matrix and a benchmark analytic Marchenko–Pastur probability density function.

We apply the Stability Measure to a number of Minimum Variance portfolios derived from key US equity indices. Our central finding is that the stability of a Minimum Variance index portfolio can be ascertained through the empirical correlation matrix of core stocks making up that particular index. In this context, portfolio stability may be thought of as analogous to stationarity in a portfolio-specific data-generating process. The fact that stability monitoring is effective for a core subset of stocks means it is practical to implement it at much higher frequencies than would ordinarily be associated with rebalancing. By identifying a change in portfolio stability, or breaks in stationarity of a portfolio-specific data generating process, our stability monitoring approach can act as a signal for the need to modify the rebalancing schedule.

Disclosure statement

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

ORCID

William Smyth  <http://orcid.org/0000-0002-4502-3818>

Daniel Broby  <http://orcid.org/0000-0001-5482-0766>

References

- Anderson, T.W., Asymptotic theory for principal component analysis. *Ann. Math. Stat.*, 1963, **34**(1), 122–148.
- Bruno, S., Chincarini, L.B. and Ohara, F., Portfolio construction and crowding. *J. Empir. Finance*, 2018, **47**, 190–206.
- Bun, J., Bouchaud, J.P. and Potters, M., Cleaning large correlation matrices: Tools from random matrix theory. *Phys. Rep.*, 2017, **666**, 1–109.
- Chopra, V.K. and Ziemba, W.T., The effect of errors in means, variances, and covariances on optimal portfolio choice. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, edited by W. Y. Ziemba, pp. 365–373, 2013 (World Scientific Publishing).
- DeMiguel, V. and Nogales, F.J., Portfolio selection with robust estimation. *Oper. Res.*, 2009, **57**(3), 560–577.
- El Karoui, N., High-dimensionality effects in the Markowitz problem and other quadratic programs with linear constraints: Risk underestimation. *Ann. Stat.*, 2010, **38**(6), 3487–3566.
- Gennotte, G., Optimal portfolio choice under incomplete information. *J. Finance*, 1986, **41**(3), 733–746.
- Haugen, R.A. and Baker, N.L., The efficient market inefficiency of capitalization-weighted stock portfolios. *J. Portf. Manag.*, 1991, **17**(3), 35–40.
- Klein, R.W. and Bawa, V.S., The effect of estimation risk on optimal portfolio choice. *J. Financ. Econ.*, 1976, **3**(3), 215–231.
- Laloux, L., Cizeau, P., Bouchard, J.P. and Potters, M., Random matrix theory and financial correlations. *Int. J. Theor. Appl. Finance*, 2000, **3**(3), 391–397.
- Laloux, L., Cizeau, P., Bouchaud, J.P. and Potters, M., Noise dressing of financial correlation matrices. *Phys. Rev. Lett.*, 1999, **83**, 1467–1470.
- Ledoit, O. and Wolf, M., Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *J. Empir. Finance*, 2003, **10**(5), 603–621.
- Lopez de Prado, M., *Machine Learning for Asset Managers*, 2020 (Cambridge University Press: Cambridge).
- Marchenko, V.A. and Pastur, L.A., Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sbornik*, 1967, **1**(4), 457–483.
- Merton, R.C., On estimating the expected return on the market: An exploratory investigation. *J. Financ. Econ.*, 1980, **8**(4), 323–361.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L.A.N. and Stanley, H.E., Universal and non-universal properties of cross-correlations in financial time series. *Phys. Rev. Lett.*, 1999, **83**, 1471–1474.
- Rao, C.R., Minimum variance quadratic unbiased estimation of variance components. *J. Multivar. Anal.*, 1971, **1**(4), 445–456.
- Sharifi, S., Crane, M., Shamaie, A. and Ruskin, H., Random matrix theory for portfolio optimization: A stability approach. *Physica A*, 2004, **335**, 629–643.
- Shepard, P.G., Second order risk. arXiv preprint arXiv:0908.2455, 2009.
- Urama, T., Ezepue, P. and Nnanwa, C., Analysis of cross-correlations in emerging markets using random matrix theory. *J. Math. Finance*, 2017, **7**, 291–307.
- Van der Vaart, A.W., *Asymptotic Statistics*, Vol. 3, 2000 (Cambridge University Press: Cambridge).