Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Portfolio Margining Using PCA Latent Factors

Shengwu Du; Travis D. Nesmith

2025-016

Please cite this paper as:
Du, Shengwu, and Travis D. Nesmith (2025). "Portfolio Margining Using PCA Latent Factors," Finance and Economics Discussion Series 2025-016. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.016.

Portfolio Margining Using PCA Latent Factors

Shengwu Du & Travis D. Nesmith*
Quantitative Risk Analysis
Board of Governors of the Federal Reserve System

February 11, 2025

Abstract

Filtered historical simulation (FHS)—a simple method of calculating Value-at-Risk that reacts quickly to changes in market volatility—is a popular method for calculating margin at central counterparties. However, FHS does not address how correlation can vary through time. Typically, in margin systems, each risk factor is filtered individually so that the computational burden increases linearly as the number of risk factors grows. We propose an alternative method that filters historical returns using latent risk factors derived from principal component analysis. We compare this method's performance with "traditional" FHS for different simulated and constructed portfolios. The proposed method performs much better when there are large changes in correlation. It also performs well when that is not the case, although some care needs to be taken with certain concentrated portfolios. At the same time, the computational requirements can be reduced significantly. Backtesting comparisons are performed using data from 2020 when markets were stressed by the COVID-19 crisis.

KEYWORDS: portfolio risk; Value-at-Risk; margin; CCPs; principal component analysis (PCA); historical simulation; FHS
JEL: G0, G2

1 Introduction

Central counterparties (CCPs) have to calculate margin requirements for any portfolio presented for clearing. As CCP margin requirements are generally defined as meeting a

*20th and C Sts., Washington, DC 20551; E-mail: shengwu.du@frb.gov; travis.d.nesmith@frb.gov
We wish to thank our colleagues in Quantitative Risk Analysis and Financial Stability Assessment for many fruitful discussions. We also greatly appreciate Yang Heppe's excellent research assistance. We also wish to thank Elaine Zhang for her excellent discussion and two reviewers for their helpful comments, both at WFEClear 2024. The views presented are solely our own and do not necessarily represent those of the Federal Reserve Board or its staff. Any remaining errors are our sole responsibility.

percentile at the portfolio level, these calculations resemble estimating portfolio Value-at-Risk (VaR).[1] VaR estimates the size of loss that can be expected at a specified probability level over a certain period.[2] CCPs often add additional charges to their margin, but the core is a market risk measure, like VaR. CCPs have to make these margin calculations frequently—at least daily if not more often—implying that both accurate and efficient risk calculations are needed for the range of possible portfolios each CCP might clear. Furthermore, margin requirement is a, if not the, primary risk management tool for CCPs, many of which themselves are considered to be systemically important to the financial system. Designing robust margin systems for CCPs is therefore both challenging and critically important. An increasing number of CCPs base their margin on forms of filtered historical simulation (FHS). Such methods are efficient to calculate, but can miss important risk characteristics. The challenge is to develop alternatives that maintain the efficiency of FHS methods, while improving accuracy.

To clarify, FHS is a common method for calculating VaR; it is a refinement of historical simulation. Historical simulation calculates VaR nonparametrically as the percentile of the data sample. It assumes that the historical data represents a simulation, where each observation is an independent draw from a fixed underlying distribution. This method has the advantage of simplicity, but the implicit assumption that the whole data sample is produced by a fixed distribution is very strong. An immediate implication is that the volatility of the series should be constant in the sample. However, the volatility of financial returns is nearly always time-varying.

FHS addresses time-varying volatility directly by scaling historical volatility to resemble the volatility of the most recent market data; it therefore estimates conditional VaR. Intuitively, it rescales the historical data to make it more like a simulation conducted under currently observed volatility. Gurrola-Perez and Murphy (2015) show that filtering greatly improves historical simulation. Many CCPs use FHS to calculate margin, usually with conditional volatility estimated semi-parametrically by an exponentially-weighted moving average (EWMA).[3]

Nevertheless, this traditional FHS VaR model has a severe limitation. It rescales the variance in the sample, but does not change the observed correlations. This result can be easily seen. Let X_t and Y_t be two variables at time t, with variances denoted $\sigma_{X,t}$ and $\sigma_{Y,t}$. For simplicity, let us assume both have a mean of zero. Then, if the current time period is T, in FHS the variables are normalized so that $\hat{X}_t = \frac{\sigma_{X,T}}{\sigma_{X,t}} X_t$ and similarly for Y_t; but their

[1]The international standards are found in CPMI-IOSCO (2012). Some CCPs exceed the standards and calculate expected shortfall rather than VaR.

[2]The literature on VaR is vast; two textbook treatments of many are Jorion (2006) and Alexander (2009).

[3]Estimating EWMA requires setting only a single decay parameter.

Figure 1: Time-varying correlation of S&P 500 stocks using EWMA estimator

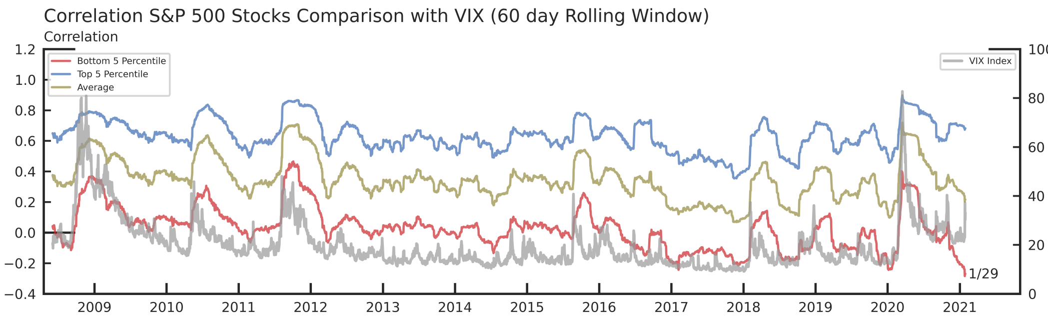

correlation is not changed as,

$$\rho_{\hat{X}_t,\hat{Y}_t} = \frac{\mathbb{E}[\hat{X}_t\hat{Y}_t]}{\sqrt{\mathbb{E}[(\hat{X}_t)^2]}\sqrt{\mathbb{E}[(\hat{Y}_t)^2]}} = \frac{\mathbb{E}[\frac{\sigma_{X,T}}{\sigma_{X,t}}X_t\frac{\sigma_{Y,T}}{\sigma_{Y,t}}Y_t]}{\sqrt{\mathbb{E}[(\frac{\sigma_{X,T}}{\sigma_{X,t}}X_t)^2]}\sqrt{\mathbb{E}[(\frac{\sigma_{Y,T}}{\sigma_{Y,t}}Y_t)^2]}} = \frac{\frac{\sigma_{X,T}}{\sigma_{X,t}}\frac{\sigma_{Y,T}}{\sigma_{Y,t}}\mathbb{E}[X_tY_t]}{\frac{\sigma_{X,T}}{\sigma_{X,t}}\sqrt{\mathbb{E}[(X_t)^2]}\frac{\sigma_{Y,T}}{\sigma_{Y,t}}\sqrt{\mathbb{E}[(Y_t)^2]}}$$

$$= \rho_{X_t,Y_t}.$$

This result is a special case of population and sample correlation being invariant to independent linear scaling.

The implicit assumption that volatility could be time-varying, but correlation constant, is both non-intuitive and likely counterfactual. Observed correlations during periods of higher volatility are certainly higher (Solnik et al., 1996). However, the increase can be the result of selection bias (Boyer et al., 1999, Longin and Solnik, 2001, Forbes and Rigobon, 2002). Nevertheless, time-varying asset correlation is nearly as well-established a stylized fact as time-varying volatility, including in these same papers; see also Groenen and Franses (2000), Baur (2006), Bramante and Gabbi (2007) and Li et al. (2024) among many others. To illustrate, Figure 1 shows that the correlations between components of the S&P 500 exhibit pronounced variation and possibly regime switches. To zoom in on a large shock, when VIX—a common measure of expected volatility in U.S. equity markets—spikes in March 2020 due to COVID-19, the correlation among individual stocks significantly increased. This phenomena is so common that it has a shorthand: "all correlations go to one." The problem is that in the presence of time-varying correlation FHS VaR likely underestimates tail risk, especially for portfolios exposed to correlation risk.

Partly to address this issue, Aramonte et al. (2013) proposed a VaR methodology based on a Dynamic Factor Model (DFM). The method can handle both time-varying volatilities and correlations for a large set of financial variables. DFM represents each individual risk factor by a linear combination of latent risk factors and the residual. The latent factors are derived from principal component analysis (PCA).[4] They model each latent risk factor as a multivariate generalized autoregressive conditional heteroskedasticity (GARCH) process.

[4] The literature on PCA is vast. Shlens (2014) provides a tutorial; Jolliffe and Cadima (2016) and Jolliffe (2022) provide more recent overviews.

There is, however, a second issue with FHS that is specifically a challenge for margin systems: the potential number of risk factors. The high dimensionality of margin systems and the requirement that calculations work for any cleared portfolio means that there is a high premium placed on computational efficiency. FHS can be applied directly to a portfolio, treating it as one risk factor. Alternatively, each position in a portfolio can be defined as a risk factor; each factor is then filtered separately and before being combined to produce the portfolio VaR. For CCP margining, the first approach would entail recalculating the filter every time each cleared portfolio changed every day. The second method involves a univariate calculation for every cleared contract every day. With a large number of changing portfolios and a large space of contracts, either approach can impose a high computational burden. In practice, CCPs have generally followed the second approach, suggesting the computational burden of the second approach is preferable.

To address shortcomings in CCPs' use of FHS, the VaR model needs to address time-varying correlation without losing the other properties that make FHS popular, in particular its ability to be applied to a variety of portfolios, while maintaining reasonable computational efficiency. Although Aramonte et al. (2013) intended for their method to be computationally efficient, the DFM approach requires Monte Carlo simulation of the returns of the latent risk factors based on the estimated GARCH process. Such simulations may be too computationally burdensome for margin systems.

This paper proposes an alternative that tries to address both time-variation in volatility and correlation, while remaining computationally feasible at scale. The method estimates portfolio VaR by applying FHS to the major latent risk factors derived using PCA, rather than to the volatilities of individual risk factors. After filtering and scaling the historical returns of latent risk factors, we then build the historical return of each individual risk factors using the PCA loading matrix. Portfolio P&L history is calculated using those returns of individual risk factors. Portfolio VaR and other tail risk measure then can be estimated using the P&L history. Implicitly, we are assuming that the principal components are relatively stable through time. This assumption is arguably much more reasonable than assuming that correlations are stable. The method is denoted PCA FHS. It is somewhat similar to orthogonal GARCH or EWMA (Alexander, 2008, 2009, McNeil et al., 2015).

To reduce computational burden, only a few principal components are used. Consequently, the residual risk of each individual risk factor is not rescaled. By using only the first few principal components, the method focuses on the primary systematic drivers of risk. Besides reducing computational burden, using the first few principal components means only those principal components are assumed to be stable, representing market structure, rather than assuming that all principal components are stable when smaller ones are much noisier. Nevertheless, for certain portfolios with concentrated idiosyncratic risk, different principal components may need to be calculated to account for the specific risk exposure.

Using PCA to reduce the complexity associated with a high number of risk factors can be done successfully if a limited number of principal components account for the most of the variance in the joint distribution. It therefore would not be surprising for our method to work in applications where factor modeling is the dominant paradigm, for example places where the yield curve is the dominant risk factor and a three- factor model captures most of the risk. Such circumstances include rates, bonds, and futures. See Alexander (2008, Chapter II.2) for examples and discussion of applying PCA in interest rate sensitive portfolios. To make it harder, we explore the performance of this method in equity markets, where factor models are used, but the factor structure is less strongly established and consistent.

Our application is motivated by such work as Laloux et al. (2000), Pafka and Kondor (2002) and Plerou et al. (2002), among others, that have applied Random Matrix Theory (RMT) to study large dimensional financial time series systems, such as stock markets. Those researches have found that the eigenvalues of the correlation matrix of stock returns are consistent with those calculated using random returns, with the exception of a few large eigenvalues, implying a large degree of randomness in the measured cross-correlation coefficients. In addition, they find that these "deviating eigenvectors" are stable in time. They analyze the components of the deviating eigenvectors and find that the largest eigenvalue corresponds to an influence common to all stocks. Moreover, their findings indicated that these large eigenvalues, which do not conform to random returns, had eigenvectors that were more stable over time. These findings support using the major latent factors derived from the PCA method to model the systematic risk in equity markets at a large scale. Nevertheless, this application is arguably more challenging; if the method performs well here, we expect its performance would be confirmed for margining assets that are generally recognized to have strong stable factor structure.

To evaluate the PCA FHS method, we perform two sets of tests. In the first, we simulate data that have a variety of sharp changes in volatility or correlation. We then conduct comparisons of PCA FHS coverage at 99 th percentile against traditional FHS. We include other comparisons, particularly against the DFM method of Aramonte et al. (2013). For a volatility break that does not change correlation, PCA FHS performs as well as traditional FHS, and both perform acceptably. This result may be surprising, as this simulation exactly follows the null assumption underlying of FHS. The situation changes dramatically when correlation changes are evaluated. PCA FHS continues to perform well, even as the performance of traditional FHS sharply degrades, so much so that the model is strongly rejected. Interestingly, for simulations with changing correlations, the performance of PCA FHS consistently approaches results produced by the full parametric estimation of the DFM method, but with much lower computational costs.

Subsequently, we test the method with constructed equity portfolios over 2019 to 2021, so that the volatile Covid-19 period is included. For a diversified portfolio, the test clearly accepts the PCA FHS method, while traditional FHS is on the borderline. The results are very similar for a long/short equity portfolio. For a portfolio concentrated in idiosyncratic positions, which is chosen explicitly to violate the PCA FHS method's focus on systematic risk, the initial application of PCA FHS is rejected while FHS is accepted. Interestingly, the performance of both methods is similar in response to the initial COVID shock, with one large exceedance, but FHS keeps margin elevated and avoids subsequent small breaches that impact the PCA FHS method's coverage. Several alternatives to improve the model performance in this situation are examined.

To more thoroughly examine potential limitations of the proposed method, we explore some particularly challenging simulated portfolios including sparse portfolios, which are more likely to deviate from the broader market behavior and some simulated long/short portfolios, which have sharply different correlations than observe in the market. The performance of the PCA FHS method can degrade. These results reflect the challenge of producing a margin system that works for any portfolio, and illustrates how CCPs need to closely monitor margin performance, particularly for smaller and idiosyncratic portfolios.

In the end, we find that the PCA FHS provides an alternative VaR method for calculating margin, improving performance in the face of correlation changes and breaks while retaining many of the practical features that makes traditional FHS appealing.

The remainder of the paper is organized as follows. Section 2 describes a general framework to study the impact of correlation change on portfolio VaR estimation. Section 3 details the estimation of the PCA FHS VaR. Section 4 provides several simulations to study PCA FHS VaR performance. Section 5 provides the empirical analysis on stock portfolio. Performance of the different methods of calculating VaR and the associated statistical tests are compared for both the simulations and the empirical analysis. Section 6 has some more simulations looking at the impact of sparse portfolios and portfolio direction on the method's performance. The last section contains concluding remarks and thoughts for future research, followed by an appendix.

2 Correlation in FHS VaR

FHS-VaR does not capture changing correlations between time-varying volatilities; only unconditional correlation among the filtered variables is captured (Pritsker, 2006). Sun and Zhang (2021) provides a theoretic framework to study how correlation impacts in FHS VaR estimation. They examined two approaches to estimate portfolio VaR using a FHS technique. The first approach applies FHS to individual risk factors, and the second approach estimates the FHS VaR based on the portfolio's P&L history. They found that applying FHS to individual risk factors could underestimate portfolio VaR, when there is large change in correlation, as this approach only considers the sample average correlation.

To illustrate the problem, let us assume a single risk factor r_i follows a GARCH(1, 1) process:

$$\begin{aligned} r_{i,t} &= \sigma_{i,t} e_{i,t} \\ e_{i,t} &\sim N(0,1) \\ \sigma_{i,t}^2 &= \omega_i + \alpha_i r_{i,t-1}^2 + \beta \sigma_{i,t-1}^2 \end{aligned} \tag{1}$$

with parameters $\omega, \alpha,$ and β.

A portfolio's P&L with position w_i on risk factor r_i is

$$PR_t = \sum_{i \le N} w_i r_{i,t}. \tag{2}$$

To run FHS on single risk factor, set

$$\bar{r}_{i,t} = \frac{\sigma_{i,T}\sigma_{i,t}}{\sigma_{i,t}} e_{i,t} = \sigma_{i,T} e_{i,t} \tag{3}$$

so that $\bar{r}_{i,t}$ is the filtered return at time $t - T$.

As an approximation, we assume portfolio daily returns follow a normal distribution, so the VaR at time T can be estimated by applying a constant multiplier to the standard deviation of portfolio return:

$$\text{VaR}^2_{T,p} = \text{Factor} \cdot E[\overline{PR}^2_t] = \text{Factor} \cdot E\left[\left(\sum_{i \leq N} w_i \sigma_{i,T} e_{i,t}\right)^2\right]. \tag{4}$$

If we assume that all the risk factors have same volatility σ at time T, and ρ is the average correlation between risk factors for the time period from 1 to T, we can rewrite the portfolio VaR at T:

$$\text{VaR}_{T,p} = \text{Factor} \cdot \sigma_T E\left[\left(\sum_{i \leq N} w_i e_{i,t}\right)\right] \approx \text{Factor} \cdot \sigma_T \cdot \rho. \tag{5}$$

Equation (5) shows that by applying the FHS technique on single risk factors, the estimated portfolio VaR is a function of risk factors' volatility at time T and sample average correlation for the time period from 1 to T. Time-varying correlation is not considered.

Time-varying correlation has been extensively studied in financial literature. A well-known example is Engle (2002), which developed a Dynamic Conditional Correlations model to address the time varying nature of correlation between financial asset returns. Aramonte et al. (2013) built on this work in developing their DFM method.

To illustrate the impact of time-varying correlation on the portfolio VaR estimation, assume that each individual risk factor in a portfolio can be modeled by a one-factor model:

$$r_{i,t} = \rho_{i,t} I_t + \sqrt{1 - \rho^2_{i,t}}\, e_{i,t} \tag{6}$$

where I_t is the systematic common risk factor, $\rho_{i,t}$ is the time-varying coefficient for individual risk factor i, and $e_{i,t}$ is the idiosyncratic residual risk. Then a portfolio's P&L with position w_i on risk factor i can be written as

$$PR_t = \sum_{i<N} w_i r_{i,t} = \sum_{i<N} (w_i \rho_{i,t} I_t + w_i \sqrt{1 - \rho^2_{i,t}}\, e_{i,t}). \tag{7}$$

By assuming zero correlation between systematic risk I_t and idiosyncratic risk $e_{i,t}$ as well as between idiosyncratic risk $e_{i,t}$ and $e_{j,t}$, we have

$$PR^2_t = \left(\sum_{i<N} w_i \rho_{i,t} I_t\right)^2 + \left(\sum_{i<N} w_i \sqrt{1 - \rho^2_{i,t}}\, e_{i,t}\right)^2 = I^2_t \cdot \sum_{i,j<N} w_i w_j \rho_{i,t} \rho_{j,t} + \sum_{i<N} \left[w_i (1 - \rho_{i,t})\, e_{i,t}\right]^2. \tag{8}$$

We further assume that all individual risk factors have the same average correlation ρ_t and the number of risk factors, N, is large, then we can drop the residual risk term in (8) to produce

$$PR_t^2 \approx \rho_t I_t^2 = \rho_t \sigma_{I,t}^2 \tag{9}$$

where $\sigma_{I,t}$ is the volatility of systematic risk.

Let us assume that portfolio returns follow a normal GARCH(1, 1) process. Portfolio VaR then can be written as

$$\begin{aligned} \text{VaR}_{p,T}^2 = \text{Factor} \cdot \sigma_{p,T+1}^2 &= \text{Factor} \cdot (\omega_p + \alpha_p PR_T^2 + \beta\sigma_{p,T}^2) \\ &= \text{Factor} \cdot (\omega_p + \alpha_p \rho_T \sigma_{I,T}^2 + \beta\sigma_{p,T}^2) \end{aligned} \tag{10}$$

If correlation changes without changing volatility, the change of portfolio VaR will be:

$$\text{VaR}_{p,T}^2 - \text{VaR}_{p,T-1}^2 = \text{Factor} \cdot \alpha_p(\rho_T - \rho_{T-1})\,\sigma_{I,T}^2 \tag{11}$$

showing how changes in correlation can impact VaR—linearly in this case—even in simplified examples. Missing the impact of correlation changes seems like a significant hole in the FHS-VaR model. In the next section, we propose a method to fill this hole.

3 FHS VaR Model Using Latent Factors from PCA

As an alternative to traditional FHS, we propose filtering based on principal components. Effectively, we rotate the data to construct uncorrelated risk factors—the principal components—and then historically filter those latent risk factors. The volatility of each selected latent factor is normalized and then rescaled to be consistent with the corresponding latent factor in the current data. In contrast to traditional FHS, when the historical returns of individual risk factors are rebuilt both the *volatility* and *correlation* has been normalized to be more consistent with current volatility and correlation.

To calculate principal components, denote historical data with T observations on N correlated asset or risk factor returns by a $T \times N$ matrix X. PCA will produce up to N uncorrelated returns, called the principal components of X, each component P being a simple linear combination of the original returns as:

$$P_{k,t} = \sum_{i<N} X_{i,t} W_{i,k}^T \tag{12}$$

where W is the matrix of eigenvectors of $X'X/T$. The weights in these linear combinations are determined by the eigenvectors of the correlation matrix of X, and the eigenvalues of this matrix are the variances of the principal components. The principal components are ordered according to the size of eigenvalue so that the first principal component, the one corresponding to the largest eigenvalue (i.e., the one with the largest variance) explains most of the variation; in a highly correlated system the largest eigenvalue will be much larger than the rest and only the first few eigenvalues will be significantly different from zero. Thus, in such systems, only a few principal components are required to represent the original variables to a high degree of accuracy.

Since W is an orthogonal matrix, we can rewrite (12) as $X = PW$, that is

$$x_{i,t} = \sum_{k<N} W_{i,k}^T P_{k,t}. \tag{13}$$

We use the first r principal components, which are the key risk factors for the system; it is important to choose only a few of these, as the empirical results show (Plerou et al. (2002, pp. 9–13) and (Laloux et al., 2000, pp. 14, 15)). Then we can rewrite (13) as

$$x_{i,t} = w_{i,1}^T P_{1,t} + w_{i,2}^T P_{2,t} + \cdots + w_{i,r}^T P_{r,t} + \varepsilon_{i,t} \tag{14}$$

where $\varepsilon_{i,t}$ is the error term picking up the approximation from using only the first r of the N principal components.

Alexander (2001, 2008) proposed an O-GARCH/O-EWMA model to build a principal component covariance matrix for a large N-dimensional multivariate process, where the name reflects that the principals components are orthogonal. Following Alexander's framework, we can estimate the volatility for those principal components using either a GARCH or EWMA model. To reduce estimation requirements, we use the EWMA model:

$$\sigma_{P_k,T+1}^2 = (1-\lambda)P_{k,T}^2 + \lambda\sigma_{P_k,T}^2. \tag{15}$$

Applying FHS to principal component returns:

$$\bar{P}_{k,t} = \frac{\sigma_{P_k,T+1}\sigma_{P_k,t}}{\sigma_{P_k,t}} e_{P_k,t} = \sigma_{P_k,T+1} e_{P_k,t}. \tag{16}$$

The single risk factor return based on FHS using the first r principal components can be written as:

$$\bar{x}_{i,t} = w_{i,1}^T \sigma_{P_1,T+1} e_{P_1,t} + w_{i,2}^T \sigma_{P_2,T+1} e_{P_2,t} + \cdots + w_{i,r}^T \sigma_{P_r,T+1} e_{P_r,t} + \varepsilon_{i,t}. \tag{17}$$

For a portfolio with large number of risk factors, we can ignore the residual risk ε_i. Then a diversified portfolio's P&L with position u_i on risk factor i can be written as

$$PR_t = \sum_{i\le N} u_i \bar{x}_{i,t} \approx \sum_{i\le N, j\le r} u_i w_{i,j}^T \sigma_{P_j,T+1} e_{P_j,t}. \tag{18}$$

Portfolio VaR is then estimated from the P&L time series data. To capture the most recent correlation impact, we propose to use the EWMA correlation or covariance matrix to derive the PC factors' loading matrix. Factors' volatility $\sigma_{P_j,T+1}$ can be estimated using either EWMA or GARCH(1, 1). Given PC factor return is a weighted average return of individual risk factor, $\sigma_{P_j,T+1}$ will be determined by the volatility and correlation of individual risk factors.

To test the effectiveness of this method, we next compare its performance to other methods in a variety of simulations.

4 Simulation Studies

We run four different simulations to examine the model performance of PCA FHS VaR. First, we study the scenario when correlation is constant while volatility has a sharp regime switch. This simulation tests how PCA FHS VaR responds to a pure volatility change; this scenario corresponds precisely to the assumptions of traditional FHS. Second, we examine the case when volatility is constant, but correlation experiences a significant regime change. Third, we study the model's performance when correlation has multiple regime switches. Lastly, we test the FHS VaR model in a scenario when both volatility and correlation regimes change together. This scenario is more consistent with the structure of financial time series data.

For each simulation, we estimate the 99% VaR for the portfolio using three different techniques: FHS, PCA FHS, and DFM. The FHS method estimates VaR based on the traditional FHS technique by filtering individual risk factors. PCA FHS is the approach described in previous section. The DFM method is based on the method proposed by Aramonte et al. (2013), also discussed previously. For the first two simulations, we also consider the index approach, which applies the standard FHS technique to the portfolio P&L history data calculated from the percentage change of portfolio value. For simplicity but also to make it more challenging for the PCA methods to be effective, we use only the top three principal components to estimate PCA FHS and DFM VaR for our simulation studies. The next section will show that the model performance could change using different numbers of principal component, especially for non-diversified portfolios.

The performance of the different methods will be evaluated by looking at the realized coverage ratio, the number of observed breaches, and two standard statistical tests: the Kupiec test (Kupiec, 1995) and the Conditional Coverage Independence (CCI) test (Christoffersen, 1998).[5] The Kupiec test is the standard test of whether the number of breaches is consistent with the targeted coverage quantile. It is two- sided and can reject for too many or too few breaches. The CCI test evaluates whether breaches are clustered rather than independently distributed as would be expected if the model is accurate. Each simulation has 300 observations. In the first two, the regime change occurs after 50 observations so that there is roughly a year worth of daily observations after the change. This duration is designed to reflect how the Kupiec test is often implemented in practice. Nevertheless, as we expect breaches to occur more frequently following regime changes, the CCI test may more accurately reflect how well the models adjust.

The first scenario is a baseline analysis to compares model performance when there is no correlation change and only a jump in volatility. As said, this simulation exactly matches the underlying assumptions of the FHS method. We simulated two sets of data from student t-distribution with degree of freedom 4 (set A and set B). This distribution is fat-tailed and empirically fits the typical estimated log-return distribution of financial assets. Both sets of data have 100 variables with variety of volatility levels. Data set A and

[5] Although Gurrola-Perez (2018) suggests that validation of FHS VaR models needs to go beyond just backtesting, comparisons in this paper primarily will stick to backtesting results.

Figure 2: simulation study 1–Portfolio VaR vs daily P&L with volatility switching

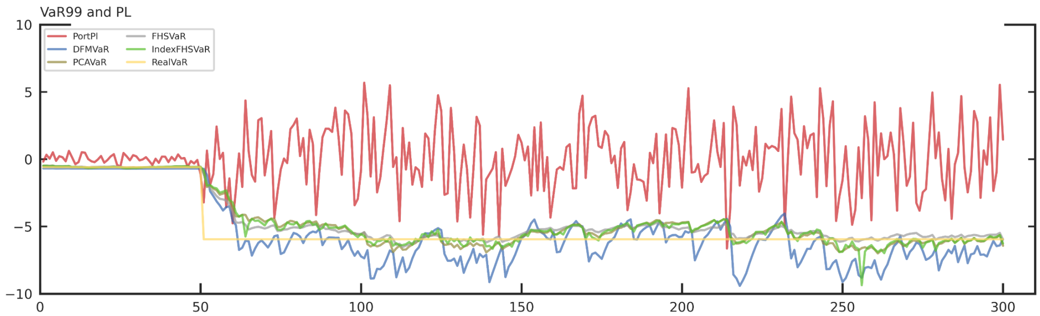

Table 3: Portfolio VaR performance with volatility switching

Model	Breaches	Cov. Ratio	POF Test*	CCI test*
FHS	6	98.0%	2.348	5.570
PCA FHS	6	98.0%	2.348	5.570
DFM	5	98.4%	1.120	5.570
RealVaR	2	99.3%	0.380	0.128
Index VaR	7	97.6%	3.916	6.319

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

set B both have 0.5 average correlation, but volatility in set B is ten times higher than in set A. Then, we build a time series with volatility regime switching by appending set B to set A. We examine an equal weighted portfolio across all 100 risk factors.

Figure 2 presents the time series of P&L moves and margin estimates for each method for the first scenario. Visually, the jump in volatility is easy to see, as movements in the first 50 observations are much more muted than the subsequent observations. Table 3 presents the 99% VaR backtesting results. Although this simulation is designed to mimic the underlying assumptions of the standard FHS method, PCA FHS produces very similar results; both methods have 6 breaches in 300 simulated dates with 98% coverage ratios. This result demonstrates that the PCA FHS adapts to changing volatility in line with the standard method. The PCA results are achieved even though the top 3 principal components explain only 50% of overall variance. The DFM method performs slightly better than both FHS and PCA FHS in terms of the coverage ratio. All the three models are accepted by the Kupiec test and the CCI test. The DFM method performs slightly better on the Kupiec test due to its lower number of breaches, but performs the same as the two FHS methods on the CCI test. Surprisingly, the index VaR has more breaches and fails both the Kupiec and CCI tests. The result could highlight another limitation of traditional FHS that its VaR estimation relies on the forecasted volatility that is used to scale the historical data. If the forecasted volatility is overly sensitive to the method used, the model performance could be comprised, especially when the volatility of the data sample varies frequently.

Figure 4: Simulation study 2–Portfolio VaR vs daily P&L with correlation switching

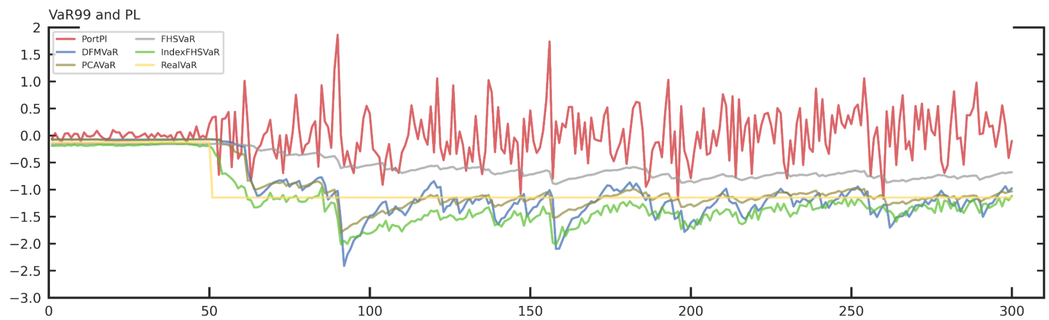

Second, we examine a scenario with correlation switching. Again, we simulated two sets of data from student t-distribution with degree of freedom 4 (set A and set B). Both sets of data have 100 variables with variety of volatility levels, and the volatilities are the same in both sets of data. Set A has low correlation between the variables, while set B has high correlation. Then we build a time series with correlation regime switching by appending set B to set A. Again, we examined an equal weighted portfolio across all 100 risk factors and estimate the 99% VaR for the portfolio using the same four different techniques: FHS, PCA FHS, DFM, and an Index method; as discussed previously. All the estimation methods use 500 days of history up to time T to calculate daily VaR. Real VaR at time T is the parametric VaR based on the pre-specified distribution used to simulate the data. We also compare the average computational time for the traditional FHS, PCA FHS, and DFM methods.

Figure 4 presents the VaR and daily P&L in percentage for 300 time periods where the correlation regime switches in the time $t = 50$. Since there is no change in volatility between two regimes, correlation switching is the driver of changes in portfolio VaR. Nevertheless, there is a visual jump in the apparent volatility of the portfolio that looks very similar to the volatility driven jump in Figure 2. This similarity illustrates that identifying the driver of changes in portfolio volatility is not straightforward, and how restricting attention solely to changes in volatility—as traditional PCA implicitly does—can be misleading. Table 5 on the following page shows the number of breaches, one year coverage ratio, and Kupiec test, as well as Conditional Coverage Independence (CCI) Test. Average run times for traditional FHS, PCA FHS, and DFM methods are also presented in the table.

As expected, FHS VaR seriously underestimates the tail risk when the correlation switches from low to high. There are over 23 VaR breaches during the studied period when correlation switched. The coverage ratio is only 90.8% for the FHS VaR model. Both Kupiec test and CCI test reject the model. Traditional FHS VaR, as we discussed in the previous section, effectively uses the average correlation of the data set. The model responses to the correlation change gradually and can only capture the correlation change when the majority of sample data are rolling into the high correlation regime. Consequently, it underperforms for a period after the correlation jumps.

Table 5: Portfolio VaR performance with correlation switching from low to high

Model	Breaches	Cov. Ratio*	POF Test*	CCI test*	Run Time
FHS	23	90.8%	62.810	57.640	0.15
PCA FHS	6	97.6%	3.550	3.822	0.04
DFM	4	98.4%	0.760	4.871	20.00
RealVaR	2	99.2%	0.108	0.144	
Index VaR	1	99.6%	1.176	1.196	

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Both PCA FHS and DFM methods significantly improve VaR estimation performance. PCA FHS has 6 VaR breaches during the studied period while DFM only has 4. The coverage ratio for PCA FHS is 97.6% and DFM achieves 98.4%, both close to the 99% coverage ratio objective. Kupiec POS test and CCI test accept both PCA FHS and DFM VaR models. Interestingly, variance explained by the top 3 principal components varies from only 33% prior to the correlation change to 94% afterwards. It is worth pointing out that the simulation is effectively a regime change—always a challenge for VaR calculations–so the ability of these two models to adapt quickly is impressive.

The DFM method does perform slightly better here. But, there is a cost. PCA FHS, as a nonparametric method, has the relative advantage of computation efficiency. The average run time for PCA FHS in this simulation is only 0.04 seconds compared to 20.0 seconds for the DFM model.[6] Among traditional FHS, PCA FHS, and DFM, PCA FHS is the most efficient approach.

Out of the four approaches, IndexVaR has the best model performance with coverage ratio over 99% and only one VaR breach. This result is not surprising because IndexVaR is applying FHS to portfolio P&L history, and factors in both systematic and idiosyncratic risk. Nevertheless, the second simulation study shows that FHS VaR model can capture the tail risk when there is a significant change in the correlation.

To further examine performance, our third simulation studies model performance when there are multiple major correlation changes. The correlation changes from low to high, a couple of times. We compare the model performance for three VaR model: traditional FHS, PCA FHS, and DFM. Figure 6 on the next page presents the 99% VaRs against portfolio daily P&L and Table 7 on the following page shows the number of breaches, one year coverage ratio, POF test, and CCI test. Again, traditional FHS under-performs in this scenario, as it does not adapt to the correlation regime switches; its coverage ratio is less than 99%, and both the POF and CCI tests reject the model. The PCA FHS and DFM models perform well in this scenario. Their coverage ratio is over the 99% requirement, and both statistical tests accept those two models.

[6] Run times are calculated using a desktop personal computer without any attempt to optimize performance. The run times provide a metric for comparing relative efficiency, but an industry application would certainly be more efficient.

Figure 6: Simulation study 3–Portfolio VaR vs daily P&L with multiple correlation regimes

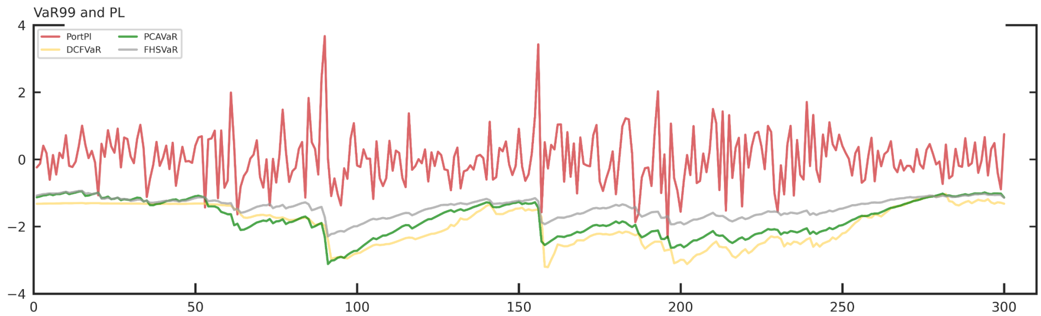

Table 7: Portfolio VaR performance with multiple correlation regimes

Model	Breaches	Cov. Ratio	POF Test*	CCI test*
FHS	7	97.6%	3.916	6.077
PCA FHS	3	99.0%	0.000	0.060
DFM	2	99.3%	0.380	0.408

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Our fourth simulation studies the model performance when both volatility and correlation change. We make two changes to the simulated data sets. First, when correlation is switching from low to high, we also increase the volatility level for all the risk factors. Second, we make the data switching back to low correlation and low volatility regime so that there are two regime switches in the time series data. This allows us to compare FHS VaR with PCA FHS VaR results when both correlation and volatility return to normal after a spike.

Figure 8 on the next page presents the 99% VaRs against portfolio daily P&L. First regime switch is happened at data point $t = 50$ and second regime switch at date point $t = 150$. There is a significant jump in portfolio's daily PL volatility when both correlation and volatility increase. Traditional FHS VaR fails to respond to regime change and underestimates the portfolio tail risk, especially in the beginning of the period when the regime changed. In stark contrast, PCA FHS captures the change and is very responsive to the risk profile dynamic, driven by both correlation and volatility changes. When both correlation and volatility switch back from high to low, DFM VaR responds quickly to the change in the portfolio risk profile. The VaR estimated using DFM method quickly decreases and reverts to the level of low volatility and correlation regime period. PCA FHS and traditional FHS VaR adjust to the regime change more smoothly and stay elevated for an extended period.

Table 9 on the following page shows the coverage ratio, number of breaches, Kupiec POF test, and CCI test. FHS has 8 VaR breaches to fail the 99% coverage ratio test. The

Figure 8: Simulation study 4–Portfolio VaR vs daily P&L with correlation & volatility changes

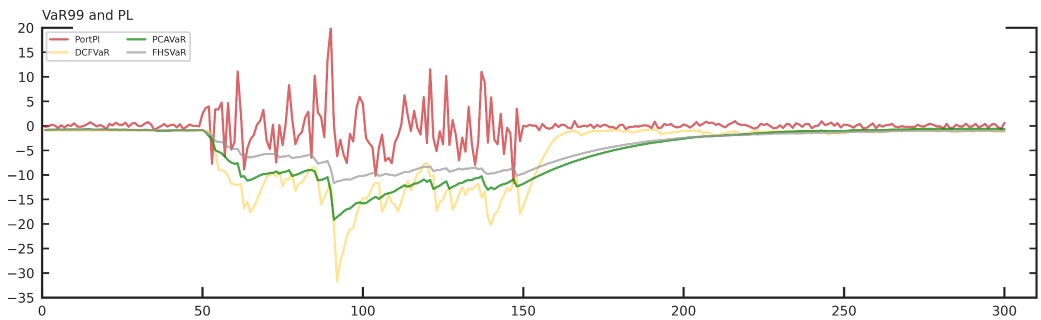

Table 9: Portfolio VaR performance with both volatility and correlation changes

Model	Breaches	Cov. Ratio	POF Test*	CCI test*
FHS	8	97.3%	5.778	6.218
PCA FHS	2	99.3%	0.381	0.408
DFM	2	99.3%	0.381	0.408

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Kupiec tests and CCI test both reject the model for 99% VaR. PCA FHS, on the other hand, only has 2 breaches and meet the 99% coverage ratio test. It also passed both Kupiec POF and CCI statistic test. Again,in this simulation, the variance explained by the principal components varies from 33% in low correlation regimes to 94% in the high correlation regime. The more computationally intensive DFM method had identical results.

There are different methods for calculating principal components. To test the potential sensitivity to different estimation methods, we rerun the PCA FHS model on this fourth set of simulated data with principal components derived from four methods: the first method uses the EWMA estimate covariance matrix; the second method uses the sample return data set; the third method uses covariance matrix filtered using RMT technique of Plerou et al. (2002); and the last method is based on robustly-estimation of the covariance matrix using the technique proposed by Maronna and Zamar (2002). For each PCA FHS estimation, we use the top three largest PCs as latent risk factors.

Table 10 on the next page shows the performance of 99% VaRs and 95% VaRs for four different PCA FHS methods. We did not find meaningful differences in model performance across the four methods. All of them have the same number of breaches and pass the tests. This lack of sensitivity to the PCA method used in VaRs estimation highlights that eigenvectors corresponding to large eigenvalues are consistently identified and are relatively stable over time, as Plerou et al. (2002) has shown. Importantly, given its efficient estimation, using EWMA to estimate the covariance matrix seems to produce results that are consistent with the other methods.

Table 10: Portfolio VaR performance with different PCA FHS methods

Model	99% VaR			95% VaR		
	Cov. Ratio	Breaches	POF test*/ CCI test*	Cov. Ratio	Breaches	POF test*/ CCI test*
PCA FHS1	99.3%	2	0.381	96.0%	12	0.021
			0.408			1.855
PCA FHS2	99.3%	2	0.381	96.0%	12	0.021
			0.408			1.855
PCA FHS3	99.3%	2	0.381	96.0%	12	0.021
			0.408			1.855
PCA FHS4	99.3%	2	0.381	96.0%	12	0.021
			0.408			1.855

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

5 Empirical Study

In this section, we study whether the PCA FHS VaR model can address the correlation risk for portfolios with different risk profiles. We run empirical studies for three different portfolios: a long only diversified stock portfolio across the S&P 500 index, a 130-30 (long 130 short 30) balanced portfolio, as well as a portfolio concentrated on names with high idiosyncratic risk. Since we exclude the non-systematic risk in our PCA FHS model setting, studying a portfolio with concentrated idiosyncratic risk provides a stern test of the method.

In all three cases, we use daily returns on the S&P 500 stocks from the Center for Research in Security Prices, LLC (CRSP) and use only stocks that have non-missing returns on all trading days from 2017 to 2022. The total number of stocks in the universe is 385. Sample portfolios have $100 000 market values and portfolio weights remain constant throughout the sample period. We run daily 99% VaR calculations using traditional FHS and PCA FHS methods, then compare their performances during the COVID-19 crisis along four dimensions: the number of VaR breaches, the coverage ratio, Kupiec POF test, and CCI test. We study the PCA FHS VaR results using the four different methods for deriving principal components, that were discussed in the simulation study section.

Diversified Portfolio

Plerou et al. (2002) has shown that the top latent factor derived from S&P 500 stock returns can consistently represent the systematic risk in equity market. Given a well diversified portfolio is mainly exposed to systematic market risk, we expect that our PCA FHS VaR with only the top two latent factors will perform well for this portfolio. We could of course include more, but testing the effectiveness of the method with only two makes the test more severe. Figure 11 on the following page presents the percentage of variance

Figure 11: Percentage of variance explained by top 2 PCs for S&P 500

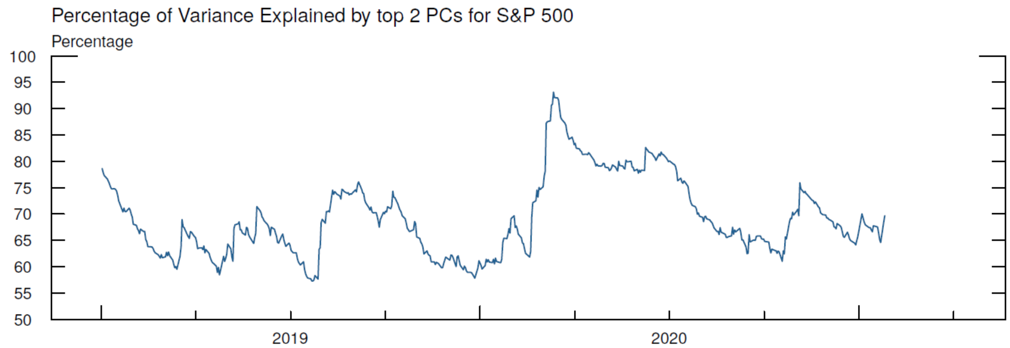

explained by the top two principal components of S&P 500 stock returns. The top two principal components always explain more than 55% of variance during the period. The variance explained is often a good bit higher and, when market volatility escalated during the COVID-19 crisis, the percentage exceeded 90%.

Table 12 compares the performance between FHS VaR and PCA FHS VaR during COVID-19 crisis period. For 99% VaR, the FHS model has 6 breaches over the one-year period covering the COVID-19 financial market turmoil. its VaR coverage ratio is below the 99% requirement. Although both Kupiec and CCI tests suggests that the model performance is acceptable, the test results are very close to the critical value to reject the model. All four PCA FHS methods achieve a better VaR performance, even limiting the method to

Table 12: FHS and PCA FHS VaR performance for diversified portfolio

	99% VaR			95% VaR		
Model	Cov. Ratio	Breaches	POF test*/ CCI test*	Cov. Ratio	Breaches	POF test*/ CCI test*
FHS	98.0%	6	3.555	94.7%	16	0.951
			5.957			2.530
PCA FHS1	99.0%	3	0.095	95.0%	15	0.496
			0.164			1.651
PCA FHS2	98.7%	4	0.769	95.0%	15	0.496
			0.887			1.651
PCA FHS3	98.7%	4	0.769	95.0%	15	0.496
			0.887			1.651
PCA FHS4	98.7%	4	0.769	95.0%	15	0.496
			0.887			1.651

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Figure 13: Daily P&L vs VaR for diversified portfolio

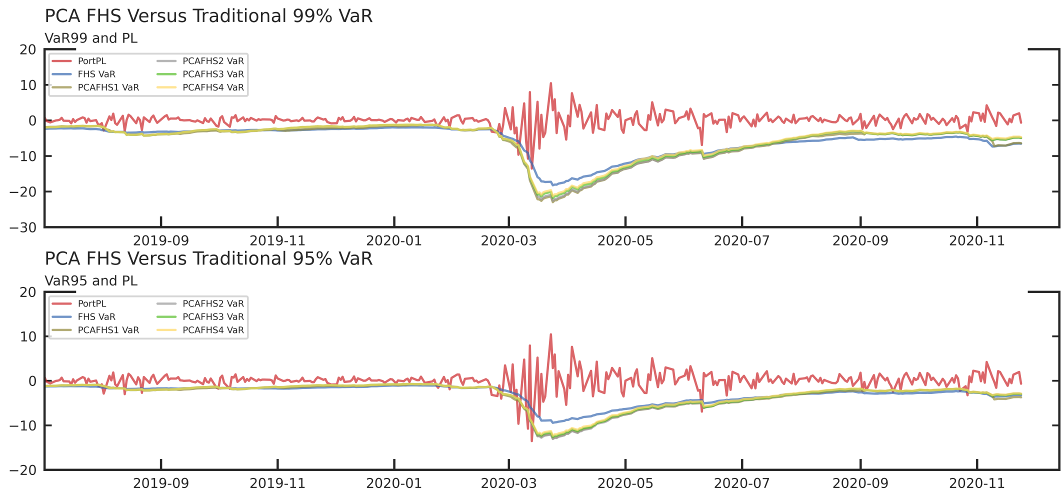

only use the first two latent factors. EWMA based PCA FHS VaR performs best with only three breaches. It also has the coverage ratio of 99% level. Both the Kupiec test and the CCI test accept the model with 95% confidence. VaR from the other three PCA FHS methods have very similar performances. All of them pass the Kupiec test and CCI test with 4 VaR breaches over the testing period. For 95% VaR, the results show that the performance between PCA FHS and traditional FHS are very similar, as FHS has only one more breach. Both models are accepted by both the Kupiec and CCI tests.

Figure 13 presents the history of portfolio's daily P&L and daily VaR results. When market volatility elevated during the COVID-19 period, correlation among stocks also increased significantly. The superior performance of the PCA FHS method through this period of market stress demonstrates the benefit of capturing correlation changes in VaR estimation, especially at higher quantiles.

To test whether adding more latent factors will significantly improve VaR estimation, we rerun our PCA FHS model with additional PCA latent risk factors. The average VaR during the 2020 COVID-19 period estimated using the four PCA FHS methods with different numbers of PCs is presented in Table 14 on the next page. Again, all four different PCA FHS methods produce very similar VaR results. Adding more PCs to VaR estimation does not change the VaR results meaningfully for this diversified portfolio. This result is expected given a well-diversified portfolio is mainly exposed to systematic market risk, which is well captured by the top principal component.

The last column of Table 14 presents the average run times for the EWMA approach. The increase in time as the number of principal components increases reflects that there is a trade-off between the effectiveness and the efficiency for the PCA methods. Implementation of a PCA FHS margin system would have to consider where to set the trade-off for a particular market type of portfolios. For the diversified equity portfolio, a small number of factors is sufficient to calculate VaR effectively, so the calculation can be chosen to be

Table 14: Average 99% VaR with different number of PCs for diversified portfolio

PC Factors	EWMA	Sample Avg.	RMT adj.	Robust	Run Time
1PC	−3113.38	−3047.98	−3092.39	−3089.37	2.0
2PC	−3121.12	−3048.35	−3099.65	−3091.32	2.0
5PC	−3058.82	−2960.70	−3005.94	−2999.13	3.0
10PC	−3086.62	−2975.45	−3030.45	−3025.80	5.0
30PC	−3144.34	−2995.30	−3051.37	−3044.12	12.0

* The run time is average running time in seconds to estimate VaR once.

extremely efficient.

130-30 portfolio

This section evaluates the PCA FHS model performance for another popular portfolio: the 130-30 strategy portfolio. This strategy is often called a long/short equity strategy, referring to an investing methodology popularly used by institutional investors. A 130-30 designation implies using a ratio of 130% of starting capital allocated to long positions made possible by taking in 30% of the starting capital from shorting stocks. This strategy tends to have larger exposure to systematic risk than a traditional diversified long-only portfolio.

Figure 15 on the following page shows the history of the daily VaR vs portfolio P&L. Table 16 on the next page compares VaR results between traditional FHS and PCA FHS methods during the COVID-19 crisis period. All our four PCA FHS methods performs well in measuring the tail risk for this type of portfolio. For 99% VaR, PCA FHS methods only have three breaches over the 300-day period, they all meet the 99% coverage ratio test and are accepted by the Kupiec test as well as CCI test. For 95% VaR, PCA FHS has 15 VaR breaches, two breaches less than FHS. The Kupiec test and CCI test favor the PCA FHS method too. Traditional FHS VaR has more breaches and falls below the coverage ratio requirement. The Kupiec test and CCI test results are close to the critical value to reject the model. This result is expected given this portfolio is also diversified and mainly exposed to systematic risk.

To further test our PCA FHS VaR model, we perform the same analysis for different long/short combination portfolios. Table 17 on page 21 presents the coverage ratio result using EMWA based PCA FHS and traditional FHS. For hedge ratio up to 70%, PCA FHS VaR with only the top two principal components performed better than traditional FHS. Idiosyncratic risk exposure for those portfolios is relatively small compared to systematic risk exposure. For the higher hedge ratio portfolio, idiosyncratic risk starts to outweigh general market risk due to the extensive offsetting due to the short positions. Interestingly, starting at a 170-70 strategy portfolio, the traditional FHS method starts out performing the PCA FHS method and even overestimates risk at nearly fully offset levels. These results suggest a need to include residual risk in the VaR estimation in such cases; either more

Figure 15: Daily P&L vs VaR for 130-30 portfolio (PCA FHS with top 2 PCs)

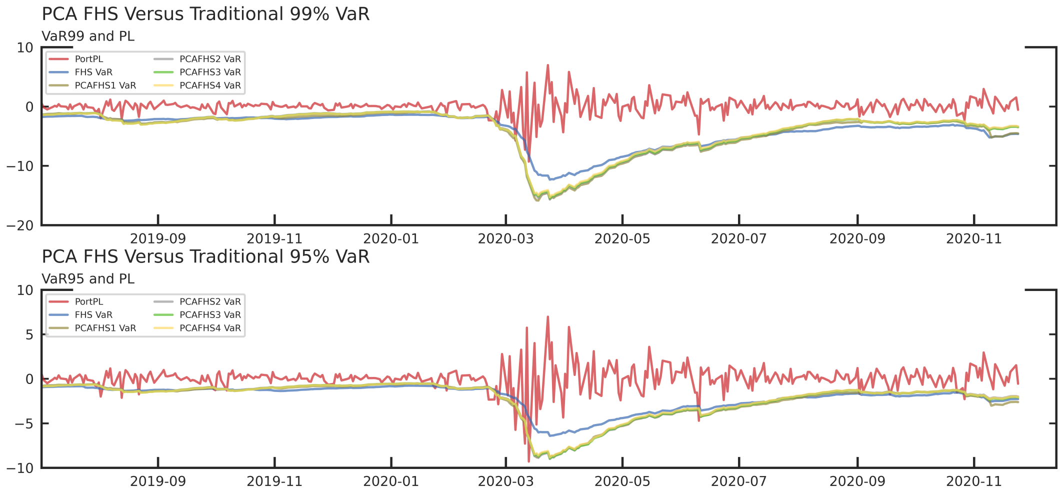

latent factors are needed to achieve better performance or latent factors need to be tailored for the specific residual risk remaining in the portfolio by calculating principal components on the actual positions. The trade offs will be examined more closely in the next section which considers concentrated portfolios.

Table 16: Average 99% PCA FHS VaR for 130-30 portfolio

	99% VaR			95% VaR		
Model	Cov. Ratio	Breaches	POF test*/ CCI test*	Cov. Ratio	Breaches	POF test*/ CCI test*
FHS	98.0%	6	3.555	94.3%	17	1.540
			5.785			1.605
PCA FHS1	99.0%	3	0.069	95.0%	14	0.096
			0.126			1.737
PCA FHS2	99.0%	3	0.095	95.0%	15	0.496
			0.126			1.567
PCA FHS3	99.0%	3	0.095	95.0%	15	0.496
			0.126			1.567
PCA FHS4	99.0%	3	0.095	95.0%	15	0.496
			0.126			1.567

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Table 17: VaR coverage ratio for long/short portfolios

Long/Short hedge ratio	99% VaR coverage		95% VaR coverage	
	PCA	FHS	PCA	FHS
100/0	99.00%	98.40%	95.60%	94.00%
100/10	99.00%	98.40%	95.60%	94.00%
100/20	99.00%	98.40%	96.40%	94.00%
100/30	99.00%	98.00%	96.00%	94.00%
100/40	99.00%	98.40%	95.60%	94.00%
100/50	99.00%	97.60%	97.20%	94.00%
100/60	98.80%	98.40%	95.20%	94.40%
100/70	98.00%	99.60%	94.76%	94.94%
100/80	98.00%	99.60%	93.62%	95.88%
100/90	91.83%	99.91%	84.20%	96.06%
100/100	64.12%	100.00%	59.84%	97.26%

Concentrated Portfolio

This section examines the performance of PCA FHS VaR model for a portfolio with concentrated positions on stocks with large market volatility. We test two different approaches to derive the PCA latent risk factors. The first approach estimates PCA latent risk factors using the daily return of 385 stocks in our data sample. The second approach only uses the stock return in the concentrated portfolio. Given that a concentrated portfolio is mainly exposed to idiosyncratic risk, we expect that PCA FHS method using only the top principal components for the overall market could underestimate the tail risk for this portfolio. More factors might be needed to capture the idiosyncratic risk.

Figure 18 on the next page presents portfolio P&L against traditional FHS VaR and PCA FHS VaR. Table 19 on the following page shows the average PCA FHS VaR results during 2020 COVID-19 period with the top two PCs. As expected, PCA FHS produces consistently smaller estimates of VaR than FHS VaR here. The top two principal components cannot capture the idiosyncratic risk in the portfolio and consequently PCA FHS underestimates the tail risk. For 99% VaR, EWMA PCA FHS has 9 VaR breaches over the 300-day period while other PCA FHS methods have 13 breaches. All PCA FHS models fail to meet 99% the coverage ratio test and are rejected by the Kupiec test and CCI test. Similar results are observed for 95% VaR. On the other hand, the traditional FHS model performs well on this concentrated portfolio. At the 99% level, traditional FHS method only has two VaR breaches over the 300-day period. It meets the 99% coverage ratio test requirement. At 95%, FHS VaR has only 12 breaches with coverage ratio over 95%. The Kupiec test and CCI test accept the traditional FHS model for both 99% and 95% VaR. This result indicates

Figure 18: Daily P&L vs VaR for concentrated portfolio (PCA FHS with top 2 PCs)

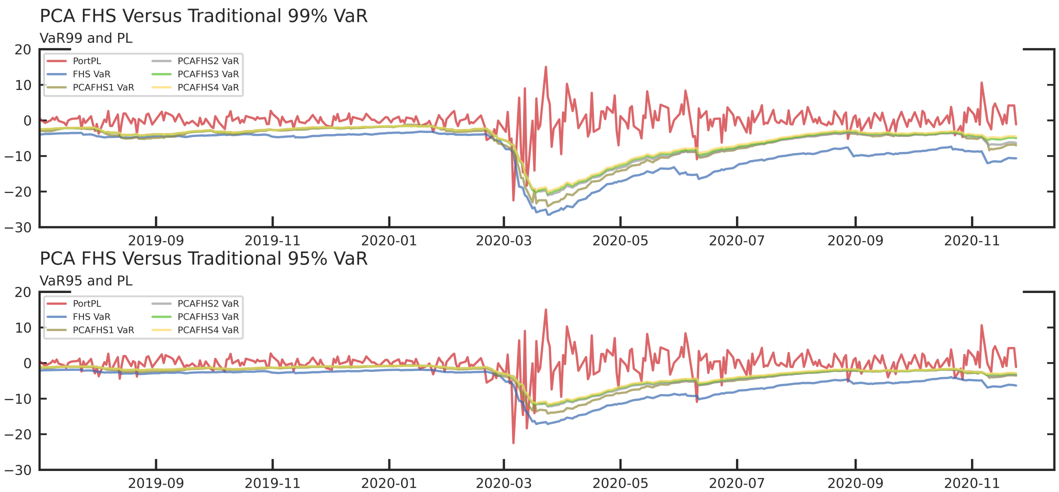

that traditional FHS VaR model can capture idiosyncratic risk that may be more driven by volatility rather than correlation.

This poor performance shows that the top two principal components are insufficient to accurately estimate VaR for the concentrated portfolio. It can be addressed by adding more PCs into the model estimation. Figure 20 on the next page shows portfolio P&L against traditional FHS VaR as well as PCA FHS VaR based on 30 PCs. VaR estimations from PCA FHS and traditional FHS are close to each other. however, traditional FHS model

Table 19: FHS and PCA FHS VaR performance for concentrated portfolio (with Top 2 PCs)

	99% VaR			95% VaR		
Model	Cov. Ratio	Breaches	POF test*/ CCI test*	Cov. Ratio	Breaches	POF test*/ CCI test*
FHS	99.3%	2	0.108	96.0%	12	0.021
			0.188			0.417
PCA FHS1	96.7%	9	9.661	90.0%	25	9.189
			10.731			17.071
PCA FHS2	95.0%	13	21.395	87.0%	29	15.462
			21.571			25.309
PCA FHS3	95.0%	13	21.395	87.7%	29	15.462
			21.571			25.309
PCA FHS4	95.0%	13	21.395	87.0%	33	22.979
			21.571			28.945

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Figure 20: Daily P&L vs VaR for concentrated portfolio (PCA FHS with top 30 PCs)

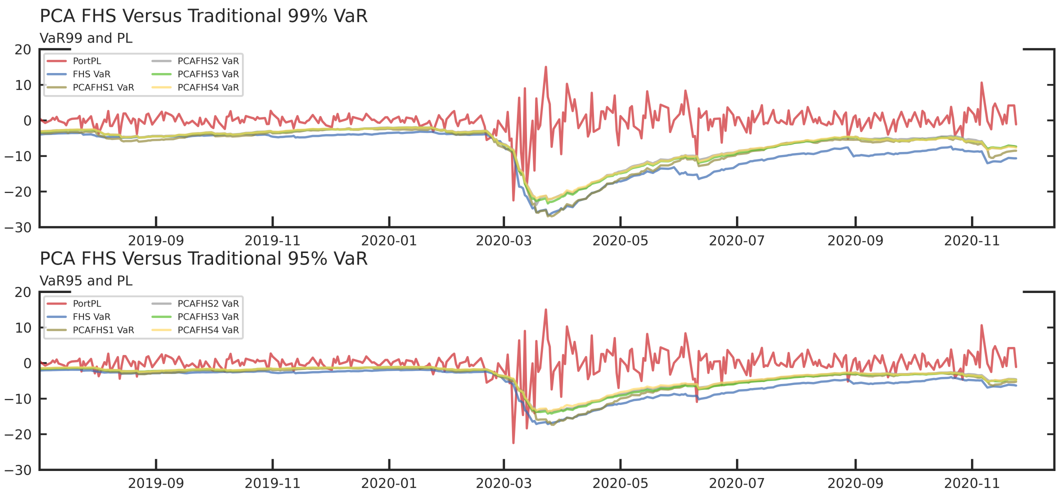

still has better model performance with fewer breaches. As seen in Table 21, calculating the PCA FHS model with 30 PCs significantly reduces the number of breaches at the 99% VaR level. Clearly, the PCA FHS model needs more latent factors than two for this concentrated portfolio. To see how performance improves with more latent factors, Table 22 on the next page presents the average VaR numbers using PCA FHS methods with a increasing number of PCs. While adding more principal components can improve PCA FHS VaR performance, it also requires more computation and reduces the method's efficiency benefit. As seen in

Table 21: FHS and PCA FHS VaR performance for concentrated portfolio (with 30 PCs)

	99% VaR			95% VaR		
Model	Cov. Ratio	Breaches	POF test*/ CCI test*	Cov. Ratio	Breaches	POF test*/ CCI test*
FHS	99.3%	2	0.108	96.3%	11	0.197
			0.187			0.415
PCA FHS1	99.0%	3	0.069	93.0%	21	4.577
			0.126			9.795
PCA FHS2	98.0%	6	3.555	91.7%	25	10.327
			3.571			17.945
PCA FHS3	98.0%	6	3.555	91.7%	25	10.327
			3.571			17.945
PCA FHS4	98.0%	6	3.555	92.4%	23	6.527
			3.571			13.294

* The $\varepsilon = .05$ quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Table 22: Average 99% PCA FHS VaR for concentrated portfolio

PC Factors	EWMA	Sample Avg.	RMT Adj.	Robust	Run Time*
1PC	−3194.50	−2661.46	−2698.81	−2695.97	2.0
2PC	−3284.58	−2878.23	−2795.35	−2814.58	2.0
5PC	−3749.87	−3560.97	−3584.53	−3524.78	3.0
10PC	−3935.98	−3720.82	−3837.27	−3760.70	5.0
30PC	−4236.35	−3844.83	−3880.20	−3809.51	12.0

* The run time is average running time in seconds to estimate VaR once.

the last column, moving from two to thirty factors increases the computation time by an order of six. This result suggests that for a portfolio with large idiosyncratic risk, it may be necessary to more directly model the risk.

An alternative approach to enhance the model performance for PCA FHS is to conduct PCA on stock returns in this concentrated portfolio only. Notably, the principal components we have used to estimate VaR for this concentrated portfolio are calculated using all of S&P 500 stock returns, although this portfolio is concentrated only on 20 stocks. Tailoring the method, we can estimate PCA FHS VaR using the top principal components derived from the return data for stocks in this portfolio only. If those top principal components capture a large portion of the volatility for this portfolio, our PCA FHS model should perform well for this concentrated portfolio. Table 23 on the following page presents the results of both upside and downside 99% VaR estimated using the first 2 principal components derived with this new approach. We only test the PCA FHS method based on sample correlation matrix given our estimated VaR numbers are not sensitive to the method used to conduct PCA.

Table 23 on the next page shows that PCA FHS model achieves the same performance as the traditional FHS method when estimating tail risk for a concentrated portfolio. The two approaches have the same number of VaR breaches as well as the coverage ratio. The Kupiec test and CCI test results suggest both models are acceptable. Figure 24 on the following page displays the daily history of 99% VaR vs portfolio P&L in percentage of portfolio value. In general, when market volatility is muted, VaR estimated by PCA FHS using portfolio specific latent factors is similar to the estimate from traditional FHS. However, when market volatility spikes as in March 2020, the jump in VaR is noticeably higher with PCA FHS compared to traditional FHS. Since the PCA FHS model implicitly includes the sharp changes in correlations observed in this period, this result is perhaps not surprising.

6 Limitations

The results to date have shown how the PCA FHS model can be an effective margin method in the face of changing correlations. However, the results also suggest that the method's effectiveness depends on how exposed a portfolio is to more systematic risk factors rather

Table 23: FHS and PCA FHS VaR for concentrated portfolio (PCA from selected risk factors)

Model	99% VaR Upside			99% VaR Downside		
	Cov. Ratio	Breaches	POF test*/ CCI test*	Cov. Ratio	Breaches	POF test*/ CCI test*
FHS	99.0%	3	0.095	99.3%	2	0.108
			0.731			0.731
PCA FHS	99.0%	3	0.095	99.3%	2	0.108
			0.185			0.185

* The ε = .05 quantiles for the POF test and for the CCI test are 3 841 and 5 991 respectively.

Figure 24: Daily P&L vs VaR for concentrated portfolio (PCA FHS with top 2 PCs using selected risk factors)

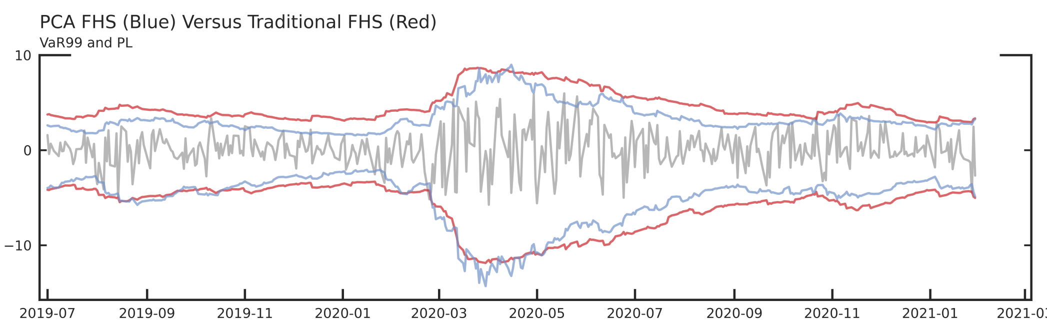

than idiosyncratic risks. In one sense, such results highlight the challenge faced by margin systems: effectively estimating risk for any feasible portfolio is a difficult problem. In this section, we further study the limitations of the PCA FHS model for two different portfolio characteristics: sparsity and offsets.

Sparsity

The first factor we examine is sparsity. If a portfolio has relatively few stocks it may be more exposed to idiosyncratic risk. To look at the impact, we simulate 100 risk factors from a multivariate student t-distribution with degree of freedom 4 and an average correlation of 0.5. We are not changing the parameters so that the overall VaR is constant. We then take samples of the risk factors, so the resulting risk will depend on the number of risk factors. We start with all the risk factors and then take smaller and smaller samples. Smaller samples, of course, can introduce more sampling volatility due to less diversification. The results are shown in Table 25 on the next page. The table shows the portfolio volatility, the actual parametric VaR, and VaR estimated by PCA FHS in two ways. The first estimates principal components from all the risk factors, while the second uses only the risk factors in the portfolio. In each case, three principal components are used.

Table 25: 99% PCA FHS VaR for shrinking portfolios

Number of factors	Volatility	Parametric VaR	PCA FHS VaR All factors	PCA FHS VaR Portfolio factors
100	2.59	6.12	6.13	6.13
90	2.60	6.07	6.08	6.02
80	2.66	6.20	6.20	6.17
70	2.62	6.11	6.12	6.30
60	2.62	6.11	6.11	6.34
50	2.63	6.21	6.17	6.70
40	2.74	6.37	6.30	6.88
30	2.77	6.48	6.35	6.72
20	3.02	7.04	6.90	7.15
10	3.02	7.05	6.45	7.00

It is informative to compare the results here to those for long/short portfolios we studied in the previous section. When the short portfolio reached 70%—so that the residual exposure was around 30%—the PCA FHS performance started to degrade compared to traditional FHS. Similarly here, when the portfolio contains 30% of the risk factors, the VaR estimate starts falling 2% or more below the actual parametric value when calculated using all risk factors. In this case, switching to PCA FHS estimation using only selected risk factors is a better approach to capture the tail risk of the portfolio.

Taken together, these results suggest the performance of PCA FHS for portfolio with more unsystematic exposures, whether due to portfolio concentration or extensive hedging, would need to be monitored. Such a specific portfolio might bring less systematic risk, but could still result in large losses to a CCP especially when the exposure is sizable and concentrated.

Direction & Offsets

Building on the long/short results, here we conduct a similar simulation to the one used to look at sparsity, but we allow constructing portfolios with both long and short positions and allow the portfolios to even be primarily short in direction. So for example, a 40/60 portfolio here would include 40 long simulated risk factors and 60 short, all equally weighted. Here we only present results using the top 3 principal components calculated using all the risk factors.

The results are in Table 26 on the following page. Not surprisingly, as the hedging ratio grows, the performance of the PCA FHS estimate worsens. Specifically, the PCA FHS VaR estimates when 40 to 50% of the positions are offset or hedged fall well short of the parametric estimates. This result is intuitive as hedged portfolios are mostly exposed

Table 26: 99% PCA FHS VaR for long and short portfolios

Long factors	Short factors	Volatility	Parametric VaR	PCA FHS VaR
90	10	2.58	4.82	4.84
80	20	2.70	3.81	3.83
70	30	2.63	2.36	2.28
60	40	2.62	1.29	1.18
50	50	2.74	0.72	0.23
40	60	2.87	1.28	1.10
30	70	3.05	2.35	2.27
20	80	3.52	3.80	3.81
10	90	3.57	4.82	4.85

to idiosyncratic risk, and the top principal components do not model it as well. Adding more principal components can improve model performance in this case. For example, for the 50/50 portfolio, using 40 principal components can remarkably improve performance: the PCA FHS VaR estimate increases to 0.73 almost identical to the parametric estimate. However, this solution reduces the efficiency of this model as the run time increases to 16 seconds from 2 seconds.

7 Conclusion

Traditional FHS VaR, which filters volatility for each individual risk factor, does not explicitly address the time-varying correlation observed in financial returns. When there is major correlation regime switching among individual risk factors, the traditional FHS model cannot respond to the correlation change; it will either underestimate or overestimate portfolio tail risk, depending on portfolio correlation exposure. CCPs who use traditional FHS VaR to estimate margin requirements could therefore under-collateralize cleared portfolios. This paper proposes a new method to estimate portfolio VaR by applying filtering to the historical return of PCA latent risk factors. Both simulation and empirical studies support that this PCA FHS method can effectively capture correlation dynamics and address risks driven by the correlation changes. Backtesting analysis performed using data from the 2020 COVID-19 financial crisis shows that for portfolios with limited idiosyncratic risk exposure, PCA FHS VaR outperforms FHS VaR.

The PCA FHS method also provides an efficient and time-saving alternative for estimating portfolio's VaR. The required running time for the PCA FHS method is much less than traditional the FHS method for large portfolios. CCPs could potentially use the PCA FHS method as an efficient margin method on its own.

There are limitations to the method. Portfolios that are more exposed to idiosyncratic risk—either because they are highly concentrated or because they hedge the systematic risk—may not be sufficiently margined by a naive application of the proposed method. Of course, traditional FHS or other margin methods may also struggle to perform across all possible portfolio types. But, more work to map the limitations of the proposed PCA FHS method likely are needed before it is used a primary margin method. Some of the limitations may be addressed by including more principal components than we have used in our analysis. Doing so increases the effectiveness at the cost of decreasing efficiency. An implementation for a particular market would need to explore the trade off in order to satisfy these competing objectives. Nevertheless, the efficiency and demonstrated ability of the PCA FHS method to respond to correlation changes would seem to argue for its adoption at least as a backstop or a comparison for other margin methods, particularly the widely adopted traditional FHS, which ignores correlation dynamics.

Furthermore, although the simulation studies and empirical exercises in this paper were limited to equity portfolios, we suspect that the method would perform well for other asset classes. In particular, rates are commonly modeled with just a few principal components, so the method would seem natural there. Also Alexander (2001, 2008) applied the O-EWMA model to commodities, suggesting the approach may extend to futures markets. In addition, CME Group (2023, pp. 7) indicates that the new SPAN2 model for futures implements some form of correlation scaling as well. Because of the potential nonlinear exposures, the method likely would outperform straight FHS for equity options, although neither method might be sufficient. Further studies are needed, however, to examine PCA FHS model performance for other asset classes and portfolios containing multiple asset classes.

References

Alexander, Carol. 2001, *Mastering Risk: Volume 2 - Applications*. Financial Times/Prentice Hall.

Alexander, Carol. 2008, *Practical Financial Econometrics*, *Market Risk Analysis*, volume II. Wiley, URL http://www.wiley.com/WileyCDA/WileyTitle/productCd-0470998016.html.

Alexander, Carol. 2009, *Market Risk Analysis: Value-at-Risk Models*, *Market Risk Analysis*, volume IV. Wiley, URL http://www.wiley.com/WileyCDA/WileyTitle/productCd-0470998016.html.

Aramonte, Sirio, Marius del Giudice Rodriguez, and Jason Wu. 2013, "Dynamic factor Value-at-Risk for large heteroskedastic portfolios." *J Bank Financ*, 37(11): 4299–4309. doi:10.1016/j.jbankfin.2013.07.038.

Baur, Dirk. 2006, "Multivariate Market Association and Its Extremes." *J. Int. Finan. Markets*, 16(4): 355–369. doi:10.1016/j.intfin.2005.05.006.

Boyer, Brian H., Michael S. Gibson, and Mico Loretan. 1999, "Pitfalls in Tests for Changes in Correlations." Technical report, Federal Reserve Board, URL https://www.federalreserve.gov/econres/ifdp/pitfalls-in-tests-for-changes-in-correlations.htm.

Bramante, Riccardo, and Giampaolo Gabbi. 2007, "Correlation Breakdowns in Asset Management." In *Advances in Risk Management*, ed. Greg N. Gregoriou, chapter 12, 226–240, London: Palgrave Macmillan UK. doi:10.1057/9780230625846_12.

Christoffersen, Peter F. 1998, "Evaluating Interval Forecasts." *International Economic Review*, 39(4): 841–862, URL http://www.jstor.org/stable/2527341.

CME Group, "SPAN 2 Margin Framework Overview Presentation." URL https://www.cmegroup.com/clearing/files/cme-span-2-margin-framework.pdf.

CPMI-IOSCO. 2012, *Principles for Financial Market Infrastructures*. URL http://www.bis.org/cpmi/publ/d101.htm.

Dobrev, Dobrislav, Travis D. Nesmith, and Dong Hwan Oh. 2017, "Accurate Evaluation of Expected Shortfall for Linear Portfolios with Elliptically Distributed Risk Factors." *J. Risk Financ. Manage.*, 10(1): 1–14. doi:10.3390/jrfm10010005.

Engle, Robert. 2002, "Dynamic Conditional Correlation." *J. Bus. Econom. Statist.*, 20(3): 339–350. doi:10.1198/073500102288618487.

Forbes, Kristin J., and Roberto Rigobon. 2002, "No Contagion, Only Interdependence: Measuring Stock Market Comovements." *J. Financ.*, 57(5): 2223–2261, URL http://www.jstor.org/stable/3094510.

Groenen, Patrick J. F., and Philip Hans Franses. 2000, "Visualizing Time-Varying Correlations across Stock Markets." *J. Empir. Financ.*, 7(2): 155–172. doi:10.1016/S0927-5398(00)00009-8.

Gurrola-Perez, Pedro. 2018, "The Validation of Filtered Historical Value-at-Risk Models." *J. Risk Model Valid.*, 12(1): 85–112. doi:10.21314/JRMV.2018.185.

Gurrola-Perez, Pedro, and David Murphy. 2015, "Filtered Historical Simulation Value-at-Risk Models and Their Competitors." Working Paper 525, Bank of England, URL http://www.bankofengland.co.uk/research/Documents/workingpapers/2015/wp525.pdf.

Jolliffe, Ian. 2022, "A 50-year Personal Journey through Time with Principal Component Analysis." *J. Multivariate Anal.*, 188: 104820. doi:10.1016/j.jmva.2021.104820.

Jolliffe, Ian T., and Jorge Cadima. 2016, "Principal Component Analysis: A Review and Recent Developments." *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.*, 374(2065): 20150202. doi:10.1098/rsta.2015.0202.

Jorion, P. 2006, *Value at Risk: The New Benchmark for Managing Financial Risk: The New Benchmark for Managing Financial Risk*. Mcgraw-Hill, 3rd edition, URL https://books.google.com/books?id=dbVLjgEACAAJ.

Kupiec, Paul H. 1995, "Techniques for Verifying the Accuracy of Risk Measurement Models." *J. Derivatives*, 3(2): 73. doi:10.3905/jod.1995.407942.

Laloux, Laurent, Pierre Cizeau, Marc Potters, and Jean-Philippe Bouchaud. 2000, "Random Matrix Theory and Financial Correlations." *Int. J. Theoretical Appl. Finance*, 3(3): 391–397. doi:10.1142/S0219024900000255.

Li, David, Fernando Cerezetti, and Roy M. Cheruvelil. 2024, "Correlation Breakdowns, Spread Positions and Central Counterparty Margin Models." *J. Financ. Market Infrastructures*, 11(3): 19–40. doi:10.21314/JFMI.2024.006.

Longin, François, and Bruno Solnik. 2001, "Extreme Correlation of International Equity Markets." *J. Financ.*, 56(2): 649–676, URL http://www.jstor.org/stable/222577.

Maronna, Ricardo A., and Ruben H. Zamar. 2002, "Robust Estimates of Location and Dispersion for High-Dimensional Datasets." *Technometrics*, 44(4): 307–317. doi:10.1198/004017002188618509.

McNeil, Alexander J., Rüdiger Frey, and Paul Embrechts. 2015, *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton Series in Finance, Princeton, New Jersey: Princeton University Press, 2nd edition, URL http://qrmtutorial.org/book.

Pafka, Szilárd, and Imre Kondor. 2002, "Noisy Covariance Matrices and Portfolio Optimization." *Eur. Phys. J. B*, 27(2): 277–280. doi:10.1140/epjb/e20020153.

Plerou, Vasiliki, Parameswaran Gopikrishnan, Bernd Rosenow, Luís A. Nunes Amaral, Thomas Guhr, and H. Eugene Stanley. 2002, "Random Matrix Approach to Cross Correlations in Financial Data." *Phys Rev E*, 65: 066126. doi:10.1103/PhysRevE.65.066126.

Pritsker, Matthew. 2006, "The Hidden Dangers of Historical Simulation." *J. Bank. Financ.*, 30(2): 561–582. doi:10.1016/j.jbankfin.2005.04.013.

Shlens, J. 2014, "A Tutorial On Principal Component Analysis." *ArXiv e-prints*, URL http://adsabs.harvard.edu/abs/2014arXiv1404.1100S. 1404.1100.

Solnik, Bruno, Cyril Boucrelle, and Yann Le Fur. 1996, "International Market Correlation and Volatility." *Financ. Anal. J.*, 52(5): 17–34. doi:10.2469/faj.v52.n5.2021.

Sun, Heng, and Zhen Zhang. 2021, "Portfolio and Index VARs by Filtered Historical Simulation." *SSRN Working Paper*, URL https://ssrn.com/abstract=3776455.

Appendix: Correlation change and Portfolio VaR

By assuming the underlying asset returns follow a normal distribution,[7] a portfolio VaR with two assets, with a notation emphasizing the dependence on correlation, equals:

$$\text{VaR}^{\rho} = F * \sigma^{\rho}_{\text{port}} = F * \sqrt{{w_1}^2{\sigma_1}^2 + {w_2}^2{\sigma_2}^2 + 2\rho_{12}w_1w_2\sigma_1\sigma_2}$$

where F is a constant and $\forall i \in \{1,2\}$ w_i is the portfolio weight on asset i and, σ_i is the volatility of asset i. The correlation between asset i and asset j is denoted ρ_{ij}. The notation $\sigma^{\rho}_{\text{port}}$ denotes portfolio variance. For both VaR^{ρ} and $\sigma^{\rho}_{\text{port}}$ values for particular correlations will be denoted by either a letter for an unspecified value or a number for a particular value, so VaR^{a} will be the Value-at-Risk with $\rho_{12} = a$ and σ^{0}_{port} will represent the portfolio variance when the two assets are uncorrelated, which equals $\sqrt{{w_1}^2{\sigma_1}^2 + {w_2}^2{\sigma_2}^2}$.

Let ρ^{a}_{ij} and ρ^{b}_{ij} denote different correlations in different periods. Then, the squared relative change in VaR, holding all other variables fixed is

$$\left(\frac{\text{VaR}^{b}}{\text{VaR}^{a}}\right)^2 = \left(\frac{\sigma^{b}_{\text{port}}}{\sigma^{a}_{\text{port}}}\right)^2 = \frac{w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2\rho^{b}_{12}w_1w_2\sigma_1\sigma_2}{w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2\rho^{a}_{12}w_1w_2\sigma_1\sigma_2} = 1 + \frac{2\left(\rho^{b}_{12} - \rho^{a}_{12}\right)w_1w_2\sigma_1\sigma_2}{\left(\sigma^{a}_{\text{port}}\right)^2}.$$

This equation shows that the squared relative change is a linear function of the change in correlation from a to b. Let $\Delta(\rho) = \rho^{b}_{12} - \rho^{a}_{12}$ denote the change. Then, we have

$$\begin{aligned}
\frac{2\left(\rho^{b}_{12} - \rho^{a}_{12}\right)w_1w_2\sigma_1\sigma_2}{\left(\sigma^{a}_{\text{port}}\right)^2} &= \frac{\Delta(\rho)\cdot 2w_1w_2\sigma_1\sigma_2}{\left(\sigma^{a}_{\text{port}}\right)^2} \\
&= \frac{\Delta(\rho)\cdot \left(\sigma^{1}_{\text{port}}\right)^2 - \left(\sigma^{0}_{\text{port}}\right)^2}{\left(\sigma^{a}_{\text{port}}\right)^2} \\
&= \Delta(\rho)\cdot\left(\frac{\left(\text{VaR}^{1}\right)^2 - \left(\text{VaR}^{0}\right)^2}{\left(\text{VaR}^{a}\right)^2}\right).
\end{aligned}$$

This result can be used to show that the change in VaR is equal to the change in correlation times the relative sensitivity of VaR to correlation, as measured by how much VaR would change going from no correlation to perfect correlation, as

$$\begin{aligned}
(\text{VaR}^{b})^2 &= (\text{VaR}^{a})^2 \cdot \left[1 + \Delta(\rho)\cdot\left(\frac{\left(\text{VaR}^{1}\right)^2 - \left(\text{VaR}^{0}\right)^2}{\left(\text{VaR}^{a}\right)^2}\right)\right] \\
&= (\text{VaR}^{a})^2 + \Delta(\rho)\cdot\left(\left(\text{VaR}^{1}\right)^2 - \left(\text{VaR}^{0}\right)^2\right).
\end{aligned}$$

[7]The analysis could be readily extended to elliptical distributions, like Student-t distributions, because they also have an analytical formula that depends similarly on correlation (Dobrev et al., 2017).

Then the percentage change in VaR for a correlation change at time t is

$$\frac{\Delta \text{VaR}}{\text{VaR}^{\rho_t}} \approx \frac{\Delta \rho_t}{2} \cdot \left(\frac{\left(\text{VaR}^1\right)^2 - \left(\text{VaR}^0\right)^2}{(\text{VaR}^{\rho_t})^2} \right).$$

For portfolios with more than two assets, an analogous result holds so that

$$\frac{\Delta \text{VaR}}{\text{VaR}^{\rho_t}} \leq \frac{\max_{i,j}(\Delta \rho_t)}{2} \cdot \left(\frac{\left(\text{VaR}^1\right)^2 - \left(\text{VaR}^0\right)^2}{(\text{VaR}^{\rho_t})^2} \right).$$

where $\max_{i,j}(\Delta \rho_t)$ is the maximum change of correlation across all assets i and j in the portfolio. Under the assumption of normality or more generally ellipticity, it is straightforward to calculate the VaR under perfect and zero correlation even for large portfolios. It is therefore easy to calculate correlation sensitivity for the portfolio and bound the change in VaR for a change in the correlation matrix. Not surprisingly, holding volatility constant, for a given change in correlation, the change in VaR will be larger for portfolios with higher correlation sensitivity.