Covariance-Filtered Historical Simulation using Simultaneous Diagonalization

NICOLA F. ZAUGG[a,*], NORBERTO BROGGINI[a], FRANCOIS DOISNEAU[a]

[a] *Capital Markets Technology, swissQuant Group AG*

Abstract

The classical filtered historical simulation (FHS) offers a robust, non-parametric way to estimate portfolio risk measures and is more and more frequently used by clearing houses to set margin requirements. The simulation filters the variance of the past returns to the current level, providing samples representative of the present market conditions. The FHS however cannot filter changes in the correlation structure of the returns, which typically leads to an underestimation of risk in volatile times, as correlation often increases during these periods. Orthogonal-FHS offers a solution by applying the classical FHS framework to an orthogonal transformation of the returns. We describe a generic framework for the Orthogonal-FHS and show under which conditions it offers robust risk measures. Using a simultaneous diagonalization algorithm we implement an Orthogonal-VaR and show its effectiveness on both simulated and empirical data.

Keywords: Simultaneous Diagonalization, Filtered Historical Simulation, Margin Requirements, Value-at-Risk, Non-Parametric Risk Measures

1. Introduction

Calculating robust portfolio risk measures is an essential task for a financial institution to limit market risk exposure. By extracting the distributional properties of the future return of the portfolio given the information available, a maximal expected loss of the portfolio given a confidence interval can be estimated and hedged accordingly. This process is essential for central counterparty clearing houses (CCPs), which require accurate short-term deep tail modeling to determine margin requirements for their clients' portfolios.

A popular method to extract such distributional properties is based on the historical simulation of the portfolio over a certain lookback period, which defines a set of possible future return scenarios. This method defines a non-parametric, data-driven approach to estimating the risk measure. Since the historical scenarios do not always accurately represent current market conditions, an approach is to transform (or *filter*) the scenarios to reflect a more representative state. This is called a *Filtered Historical Simulation* (FHS) and has been show to improve significantly the accuracy of the risk measures [12, 4, 3].

In the classical FHS, the historical returns are filtered by adjusting the variance of the return to the estimated variance of the return in the current state. This filtering technique, however, does not filter changes in the correlation structure of the assets. As a result, the resulting risk measure estimations can be off when there is a significant temporal change in the correlation structure of the assets contained in the historical simulation. Generally, market risk will be underestimated as correlation increases, leading to significant potential losses for CCPs, and overestimated as correlation drops off, leading to unnecessarily high margins.

To resolve this issue, Du and Nesmith [9] suggest the use of *principal component analysis* (PCA) to transform the correlated returns to a set of uncorrelated (or *orthogonal*) factors and

*Corresponding author.

Email addresses: zaugg@swissquant.com (NICOLA F. ZAUGG), broggini@swissquant.com (NORBERTO BROGGINI), doisneau@swissquant.com (FRANCOIS DOISNEAU)

apply an FHS-type filtering on those factors. The resulting risk measures show significant improvement compared to the classical FHS estimated risk measure when computed on empirical data, where a significant shift in the correlation structure was observed.

While the empirical result offers improvement, the temporal dependence in the data violates the main assumptions of the principal component analysis, which assumes that the observations are independent or at least serially uncorrelated [13]. As a result, the PCA decomposition can only extract unconditional orthogonal components, which are not conditionally uncorrelated, and the correlation filtering is not guaranteed to work as expected.

In this paper, we provide a theoretical framework of the simulation and show under what conditions a successful filtering of historical simulations is achieved. We propose using *simultaneous diagonalization* to extract a set of conditional and unconditional uncorrelated factors which are used to improve the filtering and the quality of the risk measures. The framework has a relevant application for margin calculations for CCPs, which require accurate and computationally tractable tail risk modeling to set initial margins as effectively as possible. While we require the method to be computationally efficient, we do not aim at any compression techniques to reduce the amount of risk factors.

1.1. Problem Setting

We consider a portfolio consisting of M assets of which we observe a matrix of historical returns $R = (r_1, r_2, \dots, r_M)^T$, where each row vector contains N historical returns $r_m = (r_{m,1}, r_{m,2}, \dots, r_{m,N})$. The aim is to use the historical returns R to obtain a risk-measure $\mathcal{X} := \mathcal{X}(R)$ (VaR, Expected Shortfall, etc.), under the rationale that the historical returns reflect some information about the distribution of the future returns $r_{m,N+1}$ of the portfolio. The classical filtered historical simulation estimates this distribution based on filtered scenarios, where the filter rescales the returns to match the current variance. It is common practice to estimate the *conditional variance* of each return r_m using an exponentially weighted moving average (EWMA) with parameter $\lambda \in [0, 1]$

$$\sigma^2_{m,n} = (1-\lambda) r^2_{m,n-1} + \lambda \sigma^2_{m,n-1}, \tag{1.1}$$

and an appropriate seed $\sigma^2_{m,0} > 0$. The filtered historical returns $\tilde{R}$ are obtained by rescaling each historical return by the current conditional variance $\sigma^2_{m,N+1}$:

$$\tilde{r}_{m,n} = r_{m,n} \frac{\sigma_{m,N+1}}{\sigma_{m,n}}, \tag{1.2}$$

Once the filtered historical returns are calculated, we aggregate the asset returns with the portfolio weights $w = (w_1, w_2, \dots, w_M)$ to obtain the filtered returns of the portfolio

$$\tilde{r}_p = w \cdot \tilde{R} \in \mathbb{R}^N. \tag{1.3}$$

The risk measure $\mathcal{X}$ is now obtained using the distributional properties of the filtered portfolio returns. The α-VaR is defined as the $(1-\alpha)$ quantile of the distribution, while the α-expected shortfall is defined as the expected value of all observation at-or-below the $(1-\alpha)$ quantile.

The filtering technique above assumes that the returns of the assets at time t are a random vector $r(t) = (r_1(t), r_2(t), \dots, r_M(t))$, where the variance of the returns $\mathbb{V}\mathrm{ar}(r_m(t_n)) = \sigma^2_{m,n}$ has a time-dependent component. In this case, the value $\tilde{r}_{m,n}$ is a sample of the random variable $\tilde{r}_m(t_n) = r_m(t_n)\frac{\sigma_{m,N}}{\sigma_{m,n}}$, whose variance is given by $\mathbb{V}\mathrm{ar}(\tilde{r}_m(t_n)) = \sigma^2_{m,N}$. A major drawback of this filtering technique is that it is applied on a univariate level. We observe that in this case, the covariance of the filtered returns between two assets $1 \leq k, l \leq M$ is given by

$$\mathrm{cov}(\tilde{x}_k(t_n), \tilde{x}_l(t_n)) = \frac{\sigma_{k,N}}{\sigma_{k,n}} \frac{\sigma_{l,N}}{\sigma_{l,n}} \mathrm{cov}(x_k(t_n), x_l(t_n)) = \sigma_{k,N}\sigma_{k,l}\rho_{k,l}(t_n), \tag{1.4}$$

where $\rho_{k,l}(t)$ is the correlation coefficient between $x_k(t)$ and $x_l(t)$. This shows that the classical FHS adjusts the variance of the returns to the current level but leaves the correlation coefficient at $\rho_{k,l}(t)$. Since the aggregated return of a portfolio depends on the variance as well as the correlation between the individual returns, the risk measure estimation can be significantly impacted if the correlation $\rho_{k,l}(t)$ changes over the lookback window, as the filtered scenarios are not adjusted to reflect the latest correlation structure. This is aggravated by the empirical fact that correlation generally increases during volatile periods [16], potentially overestimating the portfolio netting benefits and underestimating the market risk during distressed markets.

There are multiple ways proposed to improve this issue. The simplest way is to consider a classical univariate FHS on the historical returns of the portfolio rather than considering the assets individually. We refer to this method as a *Portfolio FHS*. In this case, the correlation between the assets is automatically included in the variance of the portfolio returns, and the approach offers excellent results. A significant drawback of this approach is that the filtered returns are no longer available for the individual assets but only for the aggregated portfolio. This is not suitable for risk measures that depend on individual asset returns. For instance, when calculating margin requirements for portfolios, CCPs are often required to include the *gross margin* in the calculation[1], which is defined as the sum of the individual risk measures on the instruments. Although the Portfolio FHS is not suitable for practical purposes, it offers a benchmark for other multivariate approaches, as the risk measures of a well-performing approach should be close to the Portfolio FHS.

Another intuitive approach is to extend the classical FHS to a multivariate FHS by filtering the joint random vector r_m on the covariance matrix rather than the variance of each asset [10, 6]. Suppose that Σ_n is the conditional covariance matrix at t_n, which is obtained through a multivariate EWMA estimation. Similarly to the classical FHS, we obtain the multivariate FHS by

$$\tilde{r}_n = \Sigma_N^{1/2} \Sigma_n^{-1/2} r_n \in \mathbb{R}^M, \tag{1.5}$$

where $\Sigma_n^{1/2}$ is given by the Cholesky factorization of the covariance matrix Σ_n. The resulting filtered returns thus not only reflect the latest variance of the returns but also their correlation structure Σ_N. The main drawback of the method is that it requires the calculation of $\Sigma_N^{1/2}$ and $\Sigma_n^{-1/2}$, which means that Σ_n needs to be positive definite and full-ranked for all n. Numerical instabilities imply that the method is not always suitable for practical use.

Since the classical FHS correctly filters the variance, the classical FHS is equivalent to the multivariate FHS if the correlation between the returns is 0 at all times. In other words, if $r_k(t)$ and $r_l(t)$ are independent for any $k \neq l \leq M$, the classical FHS does not need any adjustment. While it is clear that this is an unrealistic assumption, the idea of Du and Nesmith [9] is to obtain a transformation $W \in \mathbb{R}^{M,M}$, such that $Y = WR$ is a matrix of M uncorrelated column-vectors. In this case, applying a classical FHS on Y yields a vector $\tilde{Y}$, which can then be transposed back using W^{-1} to obtain $\tilde{R} = W^{-1}\tilde{Y}$. In their work, they define the transformation W using a principle component analysis, which is why we refer to the approach as *PCA-FHS*. In the PCA-FHS, the correlation between the factors r_m is included in the variance of the uncorrelated factors y_m. The filtering of Y to obtain $\tilde{Y}$ thus includes the filtering of the correlation to the current level, which greatly improves the results on both historical and artificial data [9].

Although the PCA-FHS approach shows improved results compared to the classical FHS, the principal component analysis ignores any temporal dependence between the returns. This stems from the fact that the PCA assumes all returns to stem from the same distribution. The resulting factors Y obtained through the PCA are thus unconditionally uncorrelated (across the entire time series), but not necessarily *conditionally uncorrelated*[17]. This means that although the covariance matrix of R will be diagonal, the *conditional covariance matrix* obtained through the EWMA is most likely not diagonal, and the PCA-FHS will still underestimate risk when correlation increases sharply.

1.2. Contributions

In this paper we analyse a simple stochastic framework for the filtered historical simulation. We will show under what conditions *orthogonal*-type FHS are valid models to filter the correlation structure in FHS and prove that a valid orthogonal simulation is achieved when the conditional covariance matrix of the returns is diagonal for each time step. Under these conditions, the orthogonal FHS is equivalent to the benchmark Portfolio FHS. Using an algorithm of simultaneous diagonalization to extract a transformation W, we then introduce a novel Orthogonal-FHS framework called *SD-FHS*, such that the classical FHS is valid when applied to the orthogonal vectors of $Y = WR$. Based on empirical and simulated data we show that the model provides solid risk measures and has a potential for improved margin requirements compared to the classical FHS and the PCA-FHS.

2. Model Description

2.1. Filtering Correlation

While asset returns are generally assumed to be temporally independent, it is well known that the time series exhibit a phenomenon called *volatility clustering*[7], where the square of the returns is temporally strongly correlated. For this reason it is common to utilize an ARCH-type model for financial time series, as first indicated by Engle [11]. In this paper we introduce a multivariate EWMA-type[1] model for the asset returns, which is a specific case of an autoregressive model. Suppose that a portfolio consists of M assets with relative portfolio weights $w = (w_1, w_2, \dots, w_M)$ and let $r(t) = (r_1(t), r_2(t), \dots, r_M(t))$ denote the returns of the assets at time t on a fixed grid $t \in [0, 1, \dots)$. We describe the (centered) returns of the asset as a multivariate process

$$\begin{aligned} r(t) &\sim N(0, \Sigma(t)) \\ \Sigma(t) &= (1-\lambda) r(t-1) r^T(t-1) + \lambda \Sigma(t-1) \\ \Sigma(0) &= \Sigma_0, \end{aligned} \tag{2.1}$$

where $\Sigma(t)$ is the covariance matrix of the returns starting at $\Sigma_0 \in \mathbb{R}^{M,M}$ and $\lambda \in [0,1]$ determining the reactivity of the variance of the process. The individual returns of the assets and the portfolio weights form the return of the portfolio $r_p(t)$ at time t as an aggregation of $r(t)$. We denote the aggregation function as $g(.)$:

$$r_p(t) = g(r(t)) := w \cdot r(t) = \sum_{m=1}^{M} w_m r_m(t). \tag{2.2}$$

Since the portfolio returns are calculated as a sum of the individual returns, the variance of $r_p(t)$ depends on both the variances of $r_m(t)$ as well as the correlation between the assets, as we indicated in the introduction. The portfolio filtered historical simulation is then a sampling procedure on the aggregated returns $r_p(t)$, which are univariate.

Definition 2.1 (Portfolio FHS). *Let $N \in \mathbb{N}$ be a lookback window with lookback times $t_1, t_2, \dots t_N$. A Portfolio FHS is the sequence of random variables $(\tilde{r}_p(t_1), \tilde{r}_p(t_2), \dots, \tilde{r}_p(t_N))$, where*

$$\tilde{r}_p(t_n) := r_p(t_n) \frac{\sigma_p(t_N)}{\sigma_p(t_n)}, \tag{2.3}$$

and $\sigma_p(t)$ is the volatility of $r_p(t)$.

As described earlier, the portfolio FHS acts as a benchmark for other FHS-type simulations rather than a viable simulation since the simulations of the individual assets are required for certain calculations. As the variance of $r_p(t)$ contains all the information on the latest correlation structure, the portfolio FHS correctly filters the correlation structure of the returns. The goal is to achieve a similar sampling framework such that the distribution of the samples matches the portfolio FHS, while the individual asset samples of of the portfolio constituents are still available. The classical FHS cannot achieve this due to the lack of correlation filtering. However, if we first transform the return to a new set of factors which are uncorrelated, the correlation structure is then packed into the variance of these factors, and a classical FHS will then suffice to obtain proper filtering. This is called an *orthogonal FHS*. One requirement for the orthogonal FHS is that the transformation exists.

[1]Note that the choice of autoregressive model is not specific. The theory expands to all types of processes.

Definition 2.2. *Let $r(t)$ be defined as in Equation (2.4). We call $r(t)$ orthogonal if there exists a system of M components $y(t)$ and a matrix $W \in \mathbb{R}^{M,M}$, such that*

$$\begin{aligned} r(t) &= Wy(t) \\ y(t) &\sim N(0, \Lambda(t)) \\ \Lambda(t) &= (1-\lambda)\,diag(y_1^2(t-1), y_2^2(t-1), \dots, y_M^2(t-1)) + \lambda\Lambda(t-1) \\ \Lambda(0) &= \Lambda_0, \end{aligned} \tag{2.4}$$

and

$$\Lambda(t) = \begin{pmatrix} \sigma_1^2(t) & 0 & 0 & \dots & 0 \\ 0 & \sigma_2^2(t) & 0 & \dots & 0 \\ & & \dots & & \\ 0 & \dots & 0 & \dots & \sigma_M^2(t) \end{pmatrix}$$

is diagonal. In this case, the processes $y(t)$ are independent EWMA-type processes. The processes $y(t)$ are called the orthogonal components of $r(t)$.

The existence of independent factors driving asset returns is a common theme in the modeling of asset returns and stochastic modeling in general. The setup offers a sound explanation for changes in the correlation structure over time, as a strong increase or decrease in variance of one of the factors changes the relative importance of the factor. If the independent factors are found, all the correlation of the assets is thus contained in the variance of the factors, which means that a classical FHS method to the components successfully filters the correlation. We now prove that the resulting FHS is equivalent to the Portfolio FHS

Theorem 2.1. *Let $N \in \mathbb{N}$ be a lookback window with lookback dates $t_1, t_2, \dots t_N$ and let $r(t)$ be an orthogonal system of returns with matrix W and orthogonal components $y(t)$. The sequence of random variables $(g(\tilde{r}(t_1)), g(\tilde{r}(t_2)), \dots, g(\tilde{r}(t_N))$ with*

$$\tilde{r}_m(t_n) = W \cdot y_m(t_N)\frac{\sigma_m(t_N)}{\sigma_m(t_n)}, \tag{2.5}$$

is equal in distribution to the Portfolio FHS. We call it the O(rthogonal)-FHS.

Proof. We derive the variance of the portfolio returns $r_p(t)$ for any time t conditional on all the information until $t-1$:

$$\mathbb{V}\mathrm{ar}(r_p(t)) = \mathbb{V}\mathrm{ar}(g(Wy(t))) = \mathbb{V}\mathrm{ar}(wWy(t)) = (wW)\Lambda^2(t)(wW)^T$$

The quantity $\tilde{r}_p(t) = \sqrt{\frac{\mathbb{V}\mathrm{ar}(r_p(T))}{\mathbb{V}\mathrm{ar}(r_p(t))}}r_p(t)$ is thus a sum of independent standard normals with variance

$$\frac{\mathbb{V}\mathrm{ar}(r_p(T))}{\mathbb{V}\mathrm{ar}(r_p(t))}\mathbb{V}\mathrm{ar}(r_p(t)) = (wW)\Lambda^2(T)(wW)^T$$

On the other hand, the variance of $g(\tilde{r}(t))$ is given by

$$\begin{aligned} \mathbb{V}\mathrm{ar}(g(\tilde{r}(t))) &= \mathbb{V}\mathrm{ar}\left(wW\Lambda(T)\Lambda(t)^{-1}y(t)\right) \\ &= wW\Lambda(T)\Lambda(t)^{-1}\Lambda^2(t)(wW\Lambda(T)\Lambda(t)^{-1})^T \\ &= wW\Lambda^2(T)(wW)^T \end{aligned}$$

The quantities $g(\tilde{r}(t))$ and $\tilde{r}_p(t)$ are thus both normally distributed with mean 0 and variance $wW\Lambda^2(T)(wW)^T$, which means they are equal in probability.

Since the conditional probabilities for all t and in particular thus $g(\tilde{r}(1)) \overset{d}{=} \tilde{r}_p(1)$, it follows from the total law of probability that $g(\tilde{r}(t)) \overset{d}{=} \tilde{r}_p(t)$ for all t. ■

The theorem implies that if the orthogonal components $y(t)$ exist and can be identified, the classical FHS procedure will yield the same portfolio risk measures as the index FHS.

2.2. Simultaneous Diagonalization

In this section we introduce an algorithm to extract the orthogonal components $y(t)$. The algorithm is a general technique from linear algebra called *simultaneous diagonalization (SD)*. The resulting FHS framework will therefore be called SD-FHS. Simultaneous diagonalization is an algorithm which aims to diagonalize as a set of matrices at the same time. Suppose that $A_i, i \leq I$ is a set of square matrices of size M. The goal is to obtain an (orthogonal) matrix C, such that

$$B_i = CA_iC^T, \quad i \leq I \tag{2.6}$$

is as diagonal as possible (i.e. the off-diagonal elements are as small as possible). We apply the method to obtain the orthogonal components under the following rationale. Suppose that the matrices $\Sigma(t) = \mathbb{V}\mathrm{ar}(r(t))$ are the conditional covariance matrices of the asset returns. Since the covariance matrix of the orthogonal components $y(t)$ is diagonal, we have that

$$\Lambda(t) = \mathbb{V}\mathrm{ar}(y(t)) = \mathbb{V}\mathrm{ar}(W^{-1}r(t)) = W^{-1}\Sigma(t){W^{-1}}^T, \tag{2.7}$$

is a diagonal matrix for all t. Hence, applying SD to the observable matrices $\Sigma(t)$ yields a transformation C, which is exactly the desired transformation W^{-1}. This means that the samples for the orthogonal components $y_n, n \leq N$ are found as

$$Y = C \cdot R \tag{2.8}$$

For this paper we use the *Jacobian Angles SD* implementation, which is due to Cardoso and Souloumiac [5][2]. There are other viable approaches [14, 15, 2] which offer similar results. The iterative algorithm rotates the elementary parts of the matrix piece-wise to achieve a minimization of the squared sums of off-diagonal elements (OSS), defined as

$$OSS := \sum_{n=1}^{N} \mathbf{off}\left(C\ \Sigma(t_n)\ C^T\right), \tag{2.9}$$

where

$$\mathbf{off}\,(A) := \sum_{1\leq i\neq j\leq M} |a_{i,j}|^2. \tag{2.10}$$

We refer to the resulting FHS simulation as *SD-FHS*, a particular type of O-FHS.

3. Numerical Examples

3.1. Simulation Data - Correlation Switch

To show the effectiveness of the SD-FHS, we conduct numerical examples to backtest a Value-at-Risk model based on the simulations and compare the results to the Portfolio-FHS risk-measure. The backtesting procedure consists of running daily VaR calculations and estimating the VaR level based on the simulations. The VaR levels are then compared to the observed next-day returns, of which we expect around $1-\alpha$ to "breach" the VaR level of α. In the first example we consider a similar experiment as in [9], where we consider a portfolio of M assets that are completely independent for a period, after which the correlation between the samples jumps (close) to 1. In the first period, the independent period, the returns are independently normally distributed with constant variances ranging from 20% to 30%. At the end of the period, a switch in correlation happens as the dependent period starts. The returns are now still normally distributed with the same volatility, but the correlation coefficient is set to 1. To avoid numerical instabilities, we add a small amount of noise to the returns so that the correlation matrix is not singular. Table 1 provides an overview of the settings used to construct the backtesting data. After the creation of the initial data of returns R we test the quality of the filtered simulations. We obtain the filtered matrix of returns $\tilde{R}$ for the

[2] A Python implementation of the algorithm can be found on Github thanks to Gabriel Dernbach [8].

SD-FHS, PCA-FHS according to [9], and the classical FHS. Afterwards, we compute the 99% Value-at-Risk measure on $\tilde{R}$, defined as the 1% quantile

$$\mathcal{X}(R) = Q_{0.01}(\tilde{R}) := [\tilde{r}_p]_{\lfloor 0.01 \cdot N \rfloor}, \tag{3.1}$$

where $[\tilde{r}_p]$ are the *sorted* filtered portfolio returns and $\lfloor . \rfloor$ is the floor function. We compare the obtained risk to the benchmark risk measure obtained through the Portfolio FHS. Table 2 shows an overview of the parameters for the backtesting.

Setting	**Value**
Number of assets	5
Minimum Annualized Vol	20%
Maximum Annualized Vol	30%
Length Independent Period	300
Length Dependent Period	300
Total Length	600

Table 1: Overview simulation parameters

Setting	**Value**
VaR Level	99%
One-sided	Left
Ewma Coefficient (λ)	95%
Lookback Period	$\min(500, \text{Avlb})$

Table 2: Parameters of VaR backtest

We show the results of the simulation in Figure 1. The first plot shows the daily portfolio returns as well as the 99% VaR line at this date. Notably, both the PCA-FHS and SD-FHS provide better results than the classical FHS in regards to the breach rates. This is expected as the increased correlation is not captured in the classical simulation. However, the SD-FHS reacts best to the switch in correlation. The second plot shows the average OSS as defined in Equation (2.9) over the time. We see that after the correlation switch, there is a spike in off-diagonal weight in the PCA-FHS, leading to a lower VaR estimation. The SD successfully adapts to the regime switch in correlation.

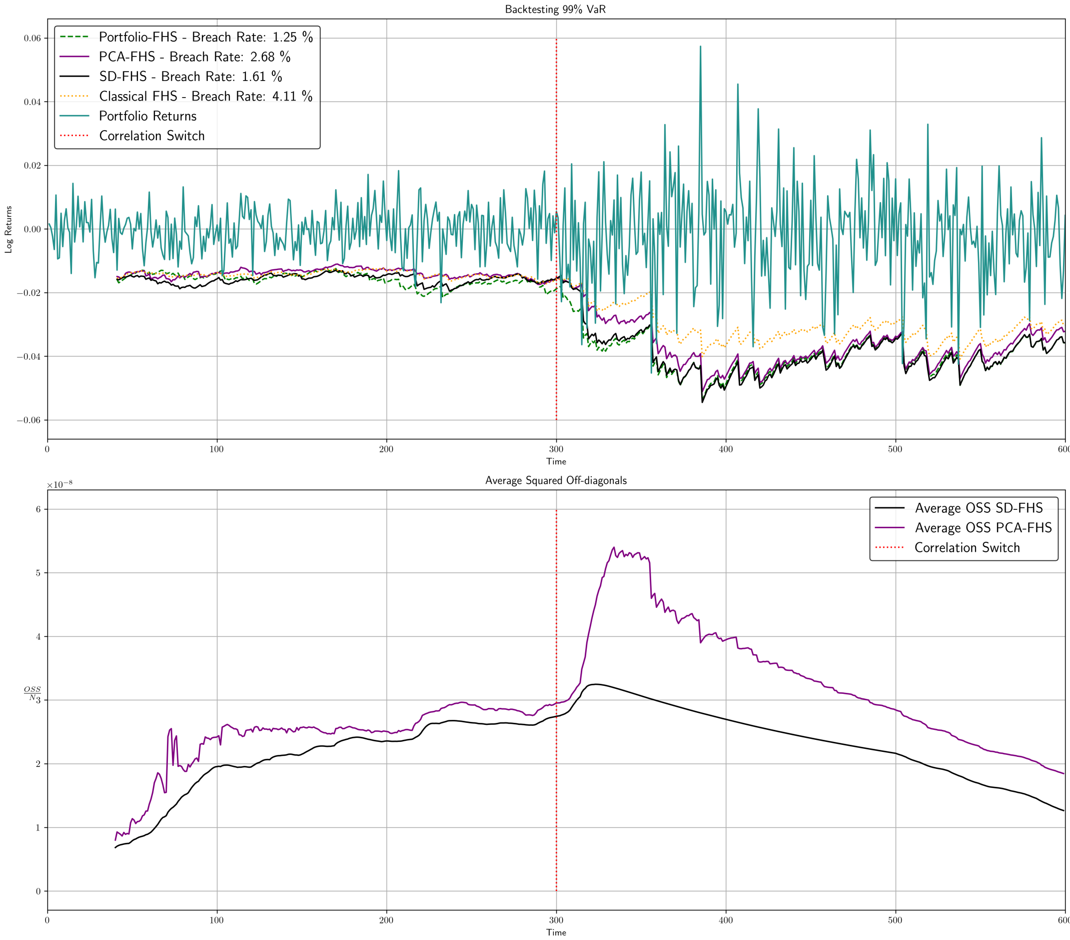

Figure 1: Comparing SD-FHS, PCA-FHS, and the classical FHS. The correlation switch increases the overall portfolio risk, leading to higher VaR levels. The risk measure calculated with SD-FHS tracks the Portfolio-FHS the closest.

We run the same experiment 50 times with different seeds and compute the average breach rates. Section 3.1 shows the statistics of this experiment. The SD FHS is the closest to the Portfolio FHS.

Method	Portfolio FHS	SD FHS	PCA FHS	Classical FHS
Mean	1.396%	1.579%	1.793%	3.261%
Standard Deviation	0.443%	0.414%	0.525%	0.436%

Table 3: Repeated simulation of 50 times yields the average break rate per FHS simulation methods

In the second example of simulated data, we construct the opposite scenario to simulate a sudden drop in correlation. For this experiment, we first simulate strongly correlated assets over 300 days and combine the scenarios with 300 days of uncorrelated assets. The setup of the experiment is thus the same as in Table 1, except that the dependent period precedes the independent period. Figure 2 shows the 99% VaR results of the experiments.

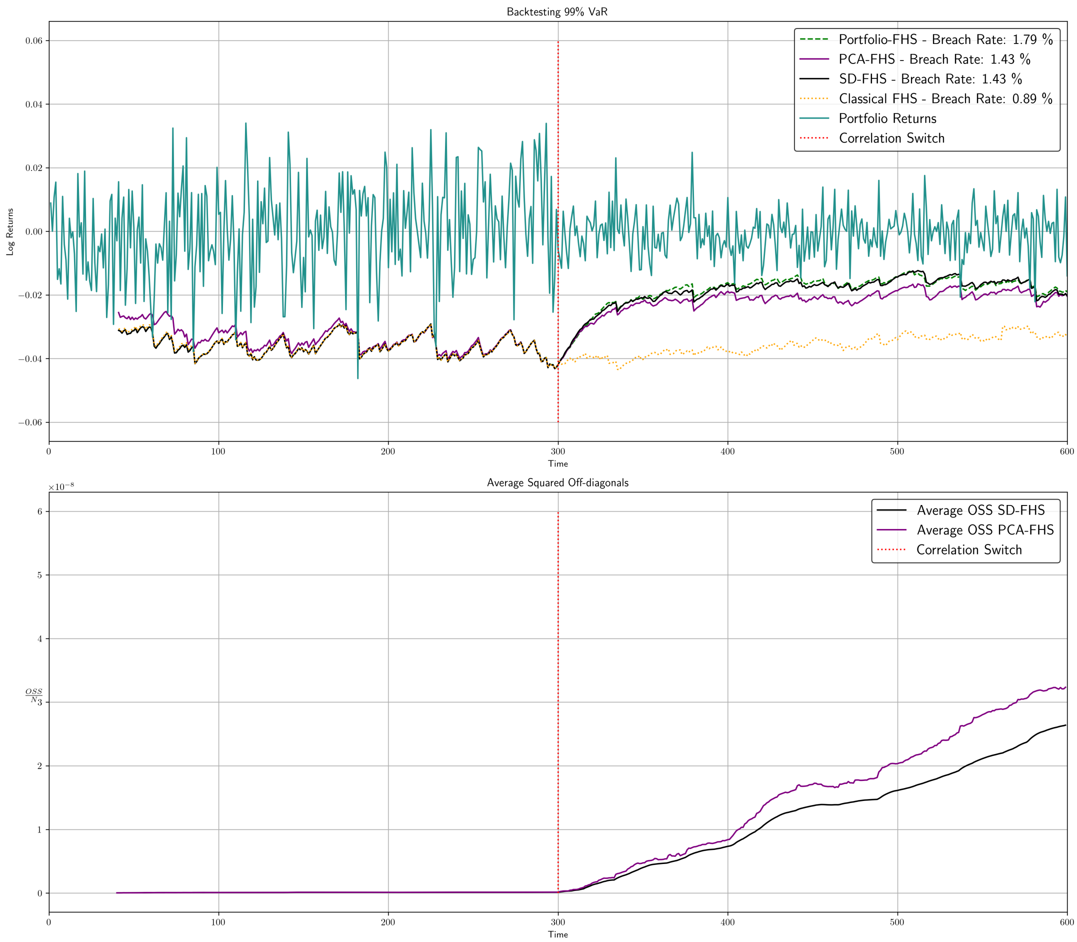

Figure 2: Comparing SD-FHS, PCA-FHS, and the classical FHS. As correlation drops, the portfolio risk decreases significantly. The classical FHS and the PCA-FHS overestimate the residual correlation risk, leading to unnecessarily high margin requirements.

As expected, the PCA-FHS and SD-FHS perform better than the classical FHS due to the correlation filtering as they track the portfolio FHS closer. Additionally, the SD-FHS outperforms the PCA-FHS as it tracks the benchmark VaR a lot closer. We conclude that in case of a correlation drop, the SD-FHS leads to lower risk measures, which means that the CCPs can set the margins lower than with the PCA-FHS.

3.2. Empirical Data - The Covid drawdown

In the second numerical experiment, we consider empirical data from the "Covid Crash" in March of 2020. This backtesting window is a good example of a switch in correlation, as the correlation spiked for a few weeks, returning to a regular level afterward. Figure 4 shows the average correlation observed during the backtesting period using a 50-day rolling window.

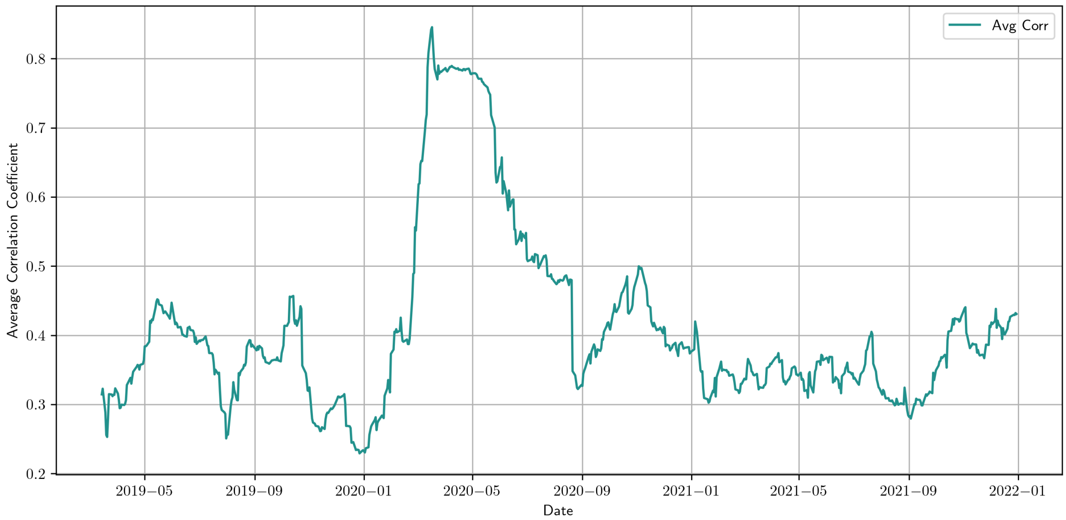

Figure 3: During the Covid crash, the average correlation spikes. Furthermore, we identify two additional minor spikes in the period leading up to 2020.

We consider a portfolio consisting of 10 equally weighted equities, of which all are constituents of the SP500. The constituent list is found in the Appendix A. We again consider a VaR backtesting where the VaR is calculated using the 4 simulation methods as before. Table 4 shows an overview of the data.

Setting	**Value**
Number of assets	10
Period start	2018-12-31
Period end	2021-12-30
Number of days	758
Portfolio weights	Equal

Table 4: Overview simulation parameters

Setting	**Value**
VaR Level	99%
One-sided	Left
Ewma Coefficient (λ)	0.98
Lookback Period	$\min(500, \text{Avlb})$

Table 5: Parameters of VaR backtest

Figure 4 shows the backtesting results over the period. Firstly, we observe that both the PCA-FHS and SD-FHS achieve the same breach rates as the Portfolio-FHS, meaning that both simulations provide accurate VaR levels. From the OSS graph, we see that the SD-FHS provides a more stable conditional covariance matrix, although the algorithm is not able to fully diagonalize during the COVID spike. Nevertheless, the overall OSS is lower than for the PCA, meaning that we can expect more accurate correlation filtering in this case. Although the differences are small, we also notice that the SD-FHS VaR line tracks the Portfolio FHS VaR line better than the PCA-FHS, which exhibits minor VaR spikes during some of the correlation spikes, leading to unnecessarily high margins. The sum of squared differences of the VaR to the benchmark VaR is shown in Section 3.2.

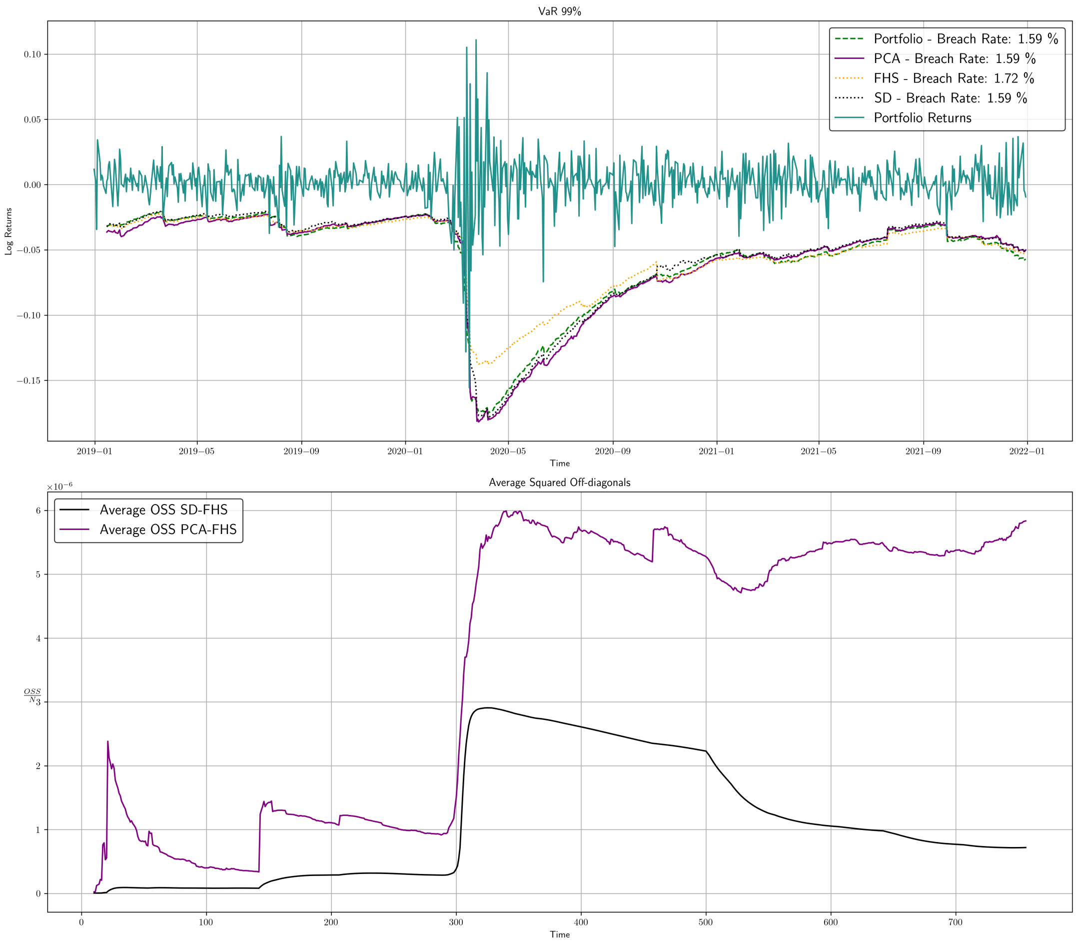

Figure 4: Both PCA-FHS and SD-FHS track the portfolio FHS-VaR better than the classical FHS. The average OSS is more stable for the SD-FHS, which provides better VaR levels during the small correlation spikes leading up to the COVID Crash.

Method	Portfolio FHS	SD FHS	PCA FHS	Classical FHS
VaR Level SSE	0	0.0812	0.11548	0.3845

Table 6: Sum squared difference to Portfolio FHS

4. Conclusion

Filtered Historical Simulation is a widely used technique for estimating risk measures such as Value-at-Risk and Expected Shortfall of asset portfolios. It is commonly employed by central counterparty clearing houses to set initial margin requirements and by risk managers to assess market risk. While FHS effectively adjusts past returns to current volatility levels, it falls short in accounting for sudden shifts in correlation structures, which can lead to underestimating portfolio risk during periods of market turbulence, when correlations tend to increase. Conversely, during stable market conditions, FHS may overestimate risk, resulting in excessively high margins.

To address the limitations of classical FHS, orthogonal FHS-type filtering techniques provide a robust solution by transforming correlated returns into uncorrelated independent factors. This approach inherently adjusts for correlations, offering a more accurate risk assessment. We presented the theoretical underpinnings of this method, demonstrating that extracting orthogonal factors yields risk measures equivalent to those from the benchmark Portfolio FHS. The main challenge lies in accurately extracting these orthogonal factors from time series data.

We propose using a simultaneous matrix diagonalization algorithm, as outlined by Cardoso and Souloumiac [5], to achieve this. The resulting simulation procedure, termed SD-FHS, has been shown to be effective in both simulated and real-world datasets, offering a significant improvement over the classical FHS approach.

References

[1] Regulation (EU) No 153/2013 of the European Parliament and of the Council on OTC derivatives, central counterparties and trade repositories supplementing regulation (EU) No 648/2012 of the European Parliament and of the Council with regard to regulatory technical standards on requirements for central counterparties. EU 153/2013, 2013.

[2] P. Ablin, J.-F. Cardoso, and A. Gramfort. Beyond pham's algorithm for joint diagonalization. *arXiv preprint arXiv:1811.11433*, 2018.

[3] G. Barone-Adesi and K. Giannopoulos. Non parametric var techniques. myths and realities. *Economic Notes*, 30(2):167–181, 2001.

[4] G. Barone-Adesi, K. Giannopoulos, and L. Vosper. Filtering historical simulation. backtest analysis. *Manuscript, March*, 2000.

[5] J.-F. Cardoso and A. Souloumiac. Jacobi angles for simultaneous diagonalization. *SIAM journal on matrix analysis and applications*, 17(1):161–164, 1996.

[6] S. Chrétien, F. Coggins, and Y. Trudel. Performance of monthly multivariate filtered historical simulation value-at-risk. *Journal of Risk Management in Financial Institutions*, 3(3):259–277, 2010.

[7] R. Cont. Volatility clustering in financial markets: empirical facts and agent-based models. In *Long memory in economics*, pages 289–309. Springer, 2007.

[8] G. Dernbach. Approximate joint diagonalization. https://github.com/gabrieldernbach/approximate_joint_diagonalization, 2019.

[9] S. Du and T. D. Nesmith. Portfolio margining using pca latent factors. *Available at SSRN 4745318*, 2023.

[10] D. Duffie and J. Pan. An overview of value at risk. *Journal of derivatives*, 4(3):7–49, 1997.

[11] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007, 1982.

[12] P. Gurrola-Perez and D. Murphy. Filtered historical simulation value-at-risk models and their competitors. *Bank of England. Quarterly Bulletin*, 55(1):103, 2015.

[13] A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.

[14] D. T. Pham. Joint approximate diagonalization of positive definite hermitian matrices. *SIAM Journal on Matrix Analysis and Applications*, 22(4):1136–1152, 2001.

[15] A.-J. Van Der Veen. Joint diagonalization via subspace fitting techniques. In *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, volume 5, pages 2773–2776. IEEE, 2001.

[16] N. F. Zaugg and L. A. Grzelak. Basket options with volatility skew: Calibrating a local volatility model by sample rearrangement. *arXiv preprint arXiv:2407.02901*, 2024.

[17] K. Zhang and L. Chan. Efficient factor garch models and factor-dcc models. *Quantitative Finance*, 9(1): 71–91, 2009.

Appendix A. List of Portfolio Constituents Empirical Section

Ticker
ALGN
ALL
ALLE
AMAT
AMCR
AMD
AME
AMGN
AMP
AMT