UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

Covariance Matrix Estimation: A Comparative Analysis

Andrea Di Guida

Supervisor: Marcin Zamojski
Master's thesis in Finance
Spring 2024
Graduate School, School of Business, Economics and Law, University of Gothenburg, Sweden

Abstract

For accurate risk assessment and portfolio optimization in finance, the covariance matrix is crucial. This research evaluates various estimation techniques, comparing their accuracy and limitations. Starting with the sample covariance matrix, the study explores both static and dynamic shrinkage approaches and concludes with a factor copula model to offer a different perspective on estimation. Through simulations and empirical analysis of high-dimensional stock market data, the study demonstrates that dynamic models outperform static ones in accuracy, even though they come at a higher computational cost. These findings underscore the need for financial practitioners to select models based on their specific requirements for accuracy, computational efficiency, and application context. Future research should investigate a broader range of assets and additional techniques to further enhance decision-making in covariance matrix estimation.

Keywords: *Covariance matrix estimation, High-dimensional, Shrinkage, Factor copula, Stock market.*

Acknowledgements

I want to first thank my supervisor, Marcin Zamojski, for his guidance and help throughout the entire thesis writing process. I am also incredibly grateful to the University of Rome Tor Vergata for giving me the wonderful opportunity to participate in this dual degree program with the University of Gothenburg. This whole experience has been amazing, both professionally and personally. Of course, none of this would have been possible without the continuous support of my family, for which I am very grateful. A special thank you goes to my girlfriend, Camilla, for her constant presence and encouragement in both the highs and lows during my entire university journey. Lastly, I want to thank all the amazing friends who joined me on this adventure, making it a truly special and unforgettable time.

Contents

1 Introduction

The covariance matrix of asset returns holds a central position in financial analysis, as emphasized by Markowitz (1952), who underscored its crucial role in portfolio selection alongside expected returns. This study focuses on estimating this covariance matrix in high-dimensional settings. Over the years, various methods have emerged for this purpose. This research aims to examine and compare some of them, particularly concerning their efficacy in capturing stock market dynamics and accuracy feasibility for a large number of assets.

This analysis begins with the well-known sample covariance matrix approach, but after that, I delve into more sophisticated methods to estimate covariance matrices. First, I explore two static linear shrinkage approaches proposed by De Nard (2019) to refine the base model. Then, I move on to the DCC with nonlinear shrinkage presented by Engle, Ledoit, and Wolf (2019) for capturing dynamic dependencies. Finally, I explore the factor copula model from Opschoor et al. (2021) to analyze a different time-varying methodology.

The sample covariance matrix stands as the most prevalent estimator for the covariance matrix, yet it suffers from high estimation errors. These are especially apparent when the number of stocks exceeds the length of the return time series, as demonstrated by Jobson and Korkie (1980). Furthermore, Michaud (1989) argues that employing this approach results in "error maximization". He suggests that the extreme coefficients in the sample covariance matrix are not necessarily inherently extreme, but rather they contain a significant amount of error.

Ledoit and Wolf (2003, 2004a,b) propose a solution termed linear shrinkage estimation, aiming to mitigate these errors by combining the sample covariance matrix with a structured estimator, effectively reducing estimation errors. They use different methods to build the structured estimator, but each has its limitations and involves strong assumptions. The most basic approach, utilizing the identity matrix, assumes zero correlation between assets, an unrealistic assumption in the stock market. Another approach, the single-factor method, assumes that the underlying variable follows the model proposed by Sharpe (1963), which is suitable only for stock returns, limiting the extension of the model to other asset classes.

To address these limitations, De Nard (2019) proposes another linear shrinkage estimator, called constant-variance-covariance (CVC), aiming to shrink both diagonal and off-diagonal coefficients of the covariance matrix, thus shrinking both variances and covariances towards their respective means through a linear approach.

Additionally, De Nard (2019) introduces a generalized version (GCVC) designed to better handle portfolios with assets of diverse characteristics, such as different industries or asset classes. This approach works by clustering the assets into groups and then

applying to each cluster the same linear shrinkage method as for the CVC. He offers various methodologies to create the groups based on either external information or statistical clustering approaches. In this study, I focus solely on the clustering approach relying on external information, specifically the industry classification.

Moving beyond linear shrinkage, Ledoit and Wolf (2012) introduce a nonlinear shrinkage (NL) approach, enhancing their previous work by applying individualized shrinkage intensity to every eigenvalue of the sample covariance matrix. Engle, Ledoit, and Wolf (2019) utilize this technique to enhance the dynamic conditional correlation (DCC) model presented by Engle (2002), which captures the time-varying nature of both the volatility of asset returns and the correlation between them. The DCC-NL model, applying the NL shrinkage approach to the DCC, overcomes issues related to the use of the sample covariance matrix as an estimator of the unconditional correlation matrix utilized in the DCC model.

In addition to the DCC-NL model, I also consider a model proposed by Opschoor et al. (2021). This model leverages a factor copula structure and offers separate estimations of variance and correlation. This is in contrast to the DCC-NL approach, which estimates these parameters jointly. This model aims to model the dependence between assets using a reduced set of factors with time-varying loadings. The model assumes that assets can be grouped based on observable characteristics, such as industry classification, as in the GCVC approach. They propose several structures to model correlation, and among these options, the 2-factor copula model is considered in this study. The model assumes a single common factor shared among all groups with equiloadings and a common factor within assets in each group with group-specific loadings.

This study aims to evaluate and compare the accuracy of these models for estimating both variance and correlation. I analyze the commonalities and differences between the models to identify the most effective approach. To achieve this, a simulation analysis is conducted using the Monte Carlo method. The analysis employs three distinct scenarios with varying asset counts and lengths to assess the models' adaptability across diverse market conditions. The key findings from the simulation reveal interesting details about the performance of each model. Dynamic models provide accurate estimates across all scenarios, but they require an initial adjustment period for changes and are computationally more expensive. Conversely, static models, as expected due to their inherent structure, struggle with dynamic market scenarios. However, they are computationally inexpensive and easier to implement compared to dynamic models.

Following the simulation analysis, I conduct an empirical application to evaluate the real-world performance of these models. The empirical analysis focuses on stocks listed in the S&P500 from January 2nd, 2004, to December 29th, 2023. After making necessary

adjustments to the data based on Standard Industrial Classification (SIC) codes, I discern distinct groups, each representing a sector. These groups are formed by portfolios consisting of stocks from various sub-industries. This empirical application offers valuable insights into the practical applicability of the models, especially during periods of crisis. Considering their inherent construction, static models struggle to capture the evolving dynamics that characterize financial markets. Their estimates are constant, failing to reflect the increased variability and potential shifts in correlations during market stress periods. In contrast, the two dynamic models demonstrate their strength in capturing these changes reflecting their ability to adapt to changing market conditions. However, it is important to note that the estimates of the two dynamic models may differ in value due to their underlying methodologies. This highlights the importance of considering not just the ability to capture dynamics, but also the specific details each model reveals. The true power lies in combining the analysis of empirical results with simulation studies. This allows for a more comprehensive understanding of each model's accuracy under various market scenarios, including crisis periods.

In conclusion, determining which model is better is not straightforward and depends on the specific needs of the users. Each model offers varying levels of accuracy and computational time. Users must identify the optimal balance between these factors to make an informed choice about the best model for their needs.

2 Theoretical Framework

2.1 Sample Covariance Matrix

The covariance matrix, denoted as Σ, plays a pivotal role as a risk measure in the financial analysis, but it is not directly observable. Consequently, it needs to be estimated. Among the various techniques available, the simplest approach is to use the sample covariance matrix estimator $\hat{S}$.

Let $R = [r_1, \ldots, r_N]$ denote the $T \times N$ asset return matrix, where r_i is the $T \times 1$ return vector of each asset, T is the sample length and N represents the number of assets. Hence, the sample covariance matrix estimator $\hat{S}$ can be computed as:

$$\hat{S} := \frac{1}{T-1}\tilde{R}\tilde{R}',$$

where $\tilde{R}$ represents the demeaned asset return matrix. This estimation approach often leads to high estimation errors, particularly when the number of assets N exceeds the length of the return time series T (Jobson and Korkie (1980)). Moreover, Michaud (1989) suggests that using this methodology results in "error maximization", which means that the extreme coefficients in the sample covariance matrix are so extreme due to a significant amount of error. Therefore, various methods have been proposed to minimize these estimation errors.

A widely adopted approach is the linear shrinkage estimator, provided by Ledoit and Wolf (2003, 2004a,b). This estimation method combines the sample covariance matrix S with a highly structured estimator F. The latter can be defined through various techniques. One approach is to consider it as a scalar multiple of the identity matrix, assuming that all pairwise correlations are equal to zero. Alternatively, F can be defined as the covariance matrix generated through the estimation of the single-index model introduced by Sharpe (1963). Ledoit and Wolf (2003) propose a convex linear combination to estimate the covariance matrix:

$$\hat{\Sigma} := \delta F + (1-\delta)S,$$

where δ is the shrinkage constant (or intensity), ranging from 0 to 1. This formulation shrinks the sample covariance matrix S towards the structured estimator F with an intensity determined by δ. According to Ledoit and Wolf (2003), there is only one optimal value for the shrinkage intensity δ^*, obtainable through the minimization of a loss function based on the Frobenius norm. This function quantifies the quadratic distance between the true covariance matrix Σ and the estimated matrix $\hat{\Sigma}$:

$$E[L(\delta)] := E[||\hat{\Sigma} - \Sigma||_F^2] = E[||\delta F + (1-\delta)S - \Sigma||_F^2],$$

where $||\cdot||_F^2$ is the (squared) Frobenius norm: $||A||_F^2 = \langle A, A\rangle = tr(AA') = \sum_{i=1}^{N}\sum_{j=1}^{N} a_{ij}^2$ for a squared matrix A with entries $(a_{ij})_{i,j=1,..N}$.

The optimal shrinkage intensity δ^* needs to be estimated due to the presence of unobservable parameter Σ in its formula. The mathematical proof of this estimation is explained in Subsection 2.2 in more depth. Consequently, the optimal linear shrinkage estimator $\hat{\Sigma}^*$ can be expressed as:

$$\hat{\Sigma}^* := \hat{\delta}^* F + (1 - \hat{\delta}^*)S.$$

The estimation of the optimal shrinkage intensity depends on the shrinkage target F selected. The most straightforward choice is to assume that the structured estimator is a scalar multiple of the identity matrix: $F_I = \eta I$, where I represents the $N \times N$ identity matrix and η is the scalar multiple. Hence, the optimal covariance matrix becomes:

$$\Sigma_I^* := \delta^* \eta I + (1 - \delta^*)S,$$

where

$$\eta := \frac{tr(\Sigma I)}{N},$$
$$\delta^* := \frac{E[||S - \Sigma||_F^2]}{E[||S - \eta I||_F^2]}.$$

The parameters η and δ^* need to be estimated due to the dependence on the unobservable parameter Σ. Therefore, $\hat{\eta} = \frac{tr(SI)}{N} = \frac{1}{N}\sum_{i=1}^{N} s_i^2 = \bar{s}^2$ and $\hat{\delta}^*$ is a feasible estimator of the optimal shrinkage intensity.

The described shrinkage approach of Ledoit and Wolf (2003, 2004a,b) effectively mitigates estimation errors in the covariance matrix compared to the basic estimation method utilizing the sample covariance matrix. However, it shrinks the off-diagonal covariance matrix coefficients towards zero, a process that fails to accurately represent reality given the inherent linear dependencies among the analyzed assets.

2.2 CVC Shrinkage

De Nard (2019) suggests an extension to the shrinkage approach of Ledoit and Wolf (2003), known as the Constant-Variance-Covariance (CVC) shrinkage approach. This estimation model adds a second parameter, ν, to not only shrink the diagonal elements (variances) of the sample covariance matrix towards their mean, but also shrink the off-diagonal elements (covariances) towards their own mean value. The formulation of the

shrinkage target under this approach is given by:

$$\hat{F}_{CVC} = \hat{\eta} I + \hat{\nu} J,$$

where $J = 11' - I$ represents the off-diagonal matrix, 1 is a conformable vector of ones, and both $\hat{\eta}$ and $\hat{\nu}$ are the estimates of the optimal values of the two scalar parameters. Hence, in this approach, the optimal covariance matrix is defined as:

$$\Sigma^*_{CVC} := \delta^* F_{CVC} + (1 - \delta^*) S,$$

where δ^* is the optimal shrinkage intensity, as defined for the linear shrinkage of Ledoit and Wolf (2003, 2004a,b).

The shrinkage target F_{CVC} can be estimated using two different methods. Both approaches involve asymptotic derivations, with the first method using large-dimensional (or Kolmogorov) asymptotics, where both N and T can go up to infinity. Conversely, the second method utilizes standard asymptotics, where only T tends to infinity. The first derivation is an extension of Ledoit and Wolf (2004a), establishing the relationship:

$$\iota^2 + \tau^2 = \zeta^2,$$

where $\iota^2 = ||\Sigma - F_{CVC}||_F^2$, $\tau^2 = E[||S - \Sigma||_F^2]$ and $\zeta^2 = E[||S - F_{CVC}||_F^2$.
Consequently, the shrinkage intensity estimation problem can be formulated as follows:

$$\min_{\delta} \; E[||\Sigma^* - \Sigma||_F^2] \tag{1}$$

$$\text{subject to } \Sigma^* = \delta F_{CVC} + (1 - \delta) S \tag{2}$$

$$F_{CVC} = \eta I + \nu J \tag{3}$$

The solution of the minimization problem ensures that:

$$\Sigma^* = \frac{\tau^2}{\zeta^2}(\eta I + \nu J) + \frac{\iota^2}{\zeta^2} S,$$

$$E[||\Sigma^* - \Sigma||_F^2] = \frac{\tau^2 \iota^2}{\zeta^2}.$$

The optimal values of the shrinkage target parameters are estimated as:

$$\hat{\eta} \equiv \bar{\sigma}^2 = \frac{1}{N} \sum_{i=1}^{N} \sigma_{ii}$$

$$\hat{\nu} \equiv \bar{\sigma}_{ij} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \sigma_{ij},$$

where $\bar{\sigma}^2$ and $\bar{\sigma}_{ij}$ are, respectively, the average population variance and covariance. Since these parameters are not directly observable, they are estimated using the sample covariance matrix. Specifically, $\bar{\sigma}^2$ is substituted with $\bar{s}^2 := \frac{1}{N}\sum_{i=1}^{N} s_{ii}$ and $\bar{\sigma}_{ij}$ with $\bar{s}_{ij} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} s_{ij}$. Finally, the optimal shrinkage intensity can be computed as:

$$\hat{\delta}^* = \frac{\hat{\tau}^2}{\hat{\iota}^2 + \hat{\tau}^2} = \frac{\hat{\tau}^2}{\hat{\zeta}^2} = 1 - \frac{\hat{\iota}^2}{\hat{\zeta}^2},$$

where $\hat{\tau}^2 := \min\{T^{-2} \sum_{t=1}^{T} ||(r_{.t} - \bar{r})' - S||_F^2, \hat{\zeta}^2\}$, $\hat{\zeta}^2 := ||\hat{\eta} I + \hat{\nu} 11' - S||_F^2$ and $\hat{\iota}^2 := \hat{\zeta}^2 - \hat{\tau}^2$ with $r_{.t}$ representing the t-th $N \times 1$ column vector of the return matrix R and $\bar{r}$ being the average $N \times 1$ asset returns vector.

The second approach to derive the optimal shrinkage intensity and the shrinkage target is based on Ledoit and Wolf (2003, 2004b). In this approach, the optimal shrinkage intensity converges asymptotically to zero as T tends to infinity. Ledoit and Wolf (2003) assume the following expression for the shrinking intensity:

$$\delta^* = \frac{\kappa}{T},$$

where $\kappa = \frac{\pi - \rho}{\gamma}$, π is the sum of asymptotic variances of the entries of S, ρ is the sum of asymptotic covariances of the entries of S with the entries of F and γ is the misspecification of the (population) shrinkage target, with all the three parameters scaled by $\sqrt{T}$. Following the estimation of all parameters, it is possible to show an estimator for the optimal shrinkage intensity:

$$\hat{\delta}^* = \min\{\max\{\frac{\hat{\kappa}}{T}, 0\}, 1\}.$$

De Nard (2019) proves that the two derivation approaches of the asymptotically optimal and consistent shrinkage intensity are equivalent, ignoring $\hat{\rho}$, which is very close to zero for F_{CVC}. In the end, the Constant-Variance-Covariance (CVC) shrinkage estimator is expressed as:

$$\hat{\Sigma}^*_{CVC} = \hat{\delta}^*(\bar{s}^2 I + \bar{s}_{ij} J) + (1 - \hat{\delta}^*) S.$$

This approach, according to De Nard (2019) outperforms the traditional shrinkage towards the identity matrix I, as it also considers covariances between asset returns in the shrinkage procedure. An additional advantage is that the shrinkage estimator is always positive definite and, therefore, invertible.

2.3 GCVC Shrinkage

De Nard (2019) suggests an alternative shrinkage approach that relies on dividing the N assets into K groups based on either external information or a clustering algorithm without the need of additional information. In both scenarios, the clustered population shrinkage target F_{GCVC} can be defined as a block matrix of the form:

$$F_{GCVC} := \begin{bmatrix} \Phi_{11} & \Phi_{12} & \dots & \Phi_{1K} \\ \Phi_{21} & \Phi_{22} & \dots & \Phi_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \Phi_{K1} & \Phi_{K2} & \dots & \Phi_{KK} \end{bmatrix},$$

where Φ_{lm} represents the (l,m)-th cluster of F_{GCVC} and K is the number of homogeneous groups. Therefore, it is possible to apply the concepts discussed in the Subsection 2.2 to each cluster Φ_{lm} as follows:

$$\Phi_{lm} = \begin{cases} \eta_k I_{N_k} + \nu_k J_{N_k} & for\ l = m = k \in \{1, ..., K\} \\ \nu_{lm} 1_{N_l} 1'_{N_m} & for\ l \neq m \in \{1, ..., K\} \end{cases},$$

where I_{N_k} and J_{N_k}, respectively, represent the $N_K \times N_K$ identity matrix and the $N_K \times N_K$ off-diagonal matrix, N_K is the number of assets in the group k and 1_{N_k} denotes the $N_K \times 1$ vector of ones.

The parameters of the shrinkage target estimator are: $\hat{\eta}_k \equiv \bar{s}_k^2$ for the average variance of the group (cluster) k, $\hat{\nu}_k \equiv \bar{s}_k$ for the average covariance of the group (cluster) k, and $\hat{\nu}_{lm} \equiv \bar{s}_{lm}$ for the average covariance of the joint cluster between group l and m. The optimal shrinkage intensity can be estimated using the formula employed in the CVC approach:

$$\hat{\delta}^* = \min\{\max\{\frac{\hat{\kappa}}{T}, 0\}, 1\},$$

where $\hat{\kappa} := \frac{\hat{\pi} - \hat{\rho}}{\hat{\gamma}}$. In particular, also for F_{GCVC}, the sum of asymptotic covariances $\hat{\rho}$, can be neglected because its value is very close to zero. Finally, the Generalized CVC shrinkage estimator is:

$$\hat{\Sigma}^*_{GCVC} = \hat{\delta}^* F_{GCVC} + (1 - \hat{\delta}^*) S.$$

The clustering procedure can be based either on external information or statistical methods. The chosen clustering approach should create homogeneous groups with similar variances and covariances for the assets within the same group k, while ensuring that assets belonging to the other $K - 1$ groups exhibit distinct variances and covariances. In this study, I cluster using external information, particularly utilizing an industry classification.

2.4 DCC-NL Shrinkage

Another approach for estimating the covariance matrix, provided by Engle (2002), is the dynamic conditional correlation (DCC) estimator. The DCC model describes the dynamics of the correlation matrix over time. For simplicity, it is assumed that the univariate volatilities follow a GARCH(1,1) process:

$$d_{i,t}^2 = \omega_i + a_i r_{i,t-1}^2 + b_i d_{i,t-1}^2,$$

where (ω_i, a_i, b_i) are the model parameters. Then, it is possible to express a conditional pseudo-correlation matrix as:

$$Q_t = (1 - \alpha - \beta)C + \alpha r_{t-1}^* r_{t-1}^{*'} + \beta Q_{t-1},$$

where $r_t^* = (\frac{r_{1,t}}{d_{1,t}}, \ldots, \frac{r_{N,t}}{d_{N,t}})'$ is the vector of devolatilized residuals, (α, β) are the DCC parameters and C is the unconditional correlation matrix. This matrix needs to be adjusted in order to obtain both the conditional correlation matrix P_t and the conditional covariance matrix H_t as follows:

$$P_t = \text{Diag}(Q_t)^{-\frac{1}{2}} Q_t \text{Diag}(Q_t)^{-\frac{1}{2}}$$

$$H_t = D_t P_t D_t,$$

where D_t is the diagonal matrix of the square root of the conditional volatilities $d_{i,t}$.

When estimating large-dimensional portfolios, the proposed model encounters challenges, particularly regarding the inversion of the conditional covariance matrix H_t, which is essential for computing the log-likelihood. Engle, Ledoit, and Wolf (2019) propose an improved model to overcome this issue, known as DCC-NL shrinkage. Specifically, they use the composite likelihood method proposed in Pakel et al. (2021). This method emphasizes the sum of the log-likelihood of pairs of assets rather than the joint log-likelihood of all assets.

An additional improvement in DCC-NL involves applying the nonlinear (NL) shrinkage technique, as discussed in Ledoit and Wolf (2012), to estimate the unconditional correlation matrix C instead of relying on the sample covariance matrix. Nonlinear shrinkage allows the use of different shrinkage intensities for the different entries of the sample covariance matrix S. Specifically, an alternative expression for the linear shrinkage combination is essential, employing the spectral decomposition of the sample covariance matrix, as detailed by Ledoit and Wolf (2004a). It allows shrinking the sample eigenvalues rather than directly shrinking the entries of the sample covariance matrix. Consequently, this

refinement enables the utilization of distinct shrinkage intensities for individual sample eigenvalues. While the population covariance matrix Σ remains unknown, it is possible to derive the population eigenvalues from the sample eigenvalues. This result is achievable through the numerical inversion of the Marcenko and Pastur (1967) equation, which explains the asymptotic relationship between population and sample eigenvalues.

The process starts with the estimation of the sample covariance matrix $\hat{C}$ of the devolatilized returns matrix R^*:

$$\hat{C} := \frac{1}{T} R^* R^{*'} = \frac{1}{T} \sum_{t=1}^{N} r_t^* r_t^{*'}.$$

Consequently, let $\hat{C}$ be a set of eigenvalues $(\lambda_1, \lambda_2, ..., \lambda_N)$ and their corresponding eigenvectors $(u_1, u_2, ..., u_N)$, yielding:

$$\hat{C} = \sum_{i=1}^{N} \lambda_i u_i u_i'.$$

By inverting the Marcenko and Pastur (1967) equation, it becomes feasible to derive a deterministic equivalent of the sample eigenvalues $Q_{N,T}(t) = (q_{N,T}^1(t), q_{N,T}^2(t), ..., q_{N,T}^N(t))$ from the population eigenvalues $t = (t_1, ..., t_N)$:

$$\tilde{\tau} = \arg\min_{t \in [0,\infty)^N} \frac{1}{N} \sum_{i=1}^{N} [q_{N,T}^i(t) - \lambda_i]^2.$$

Utilizing Theorem 4 from Ledoit and Péché (2011), it is possible to compute the nonlinear shrinkage formula, which is asymptotically optimal under large-dimensional asymptotics. Thus, the nonlinear shrinkage estimator of the covariance matrix can be expressed as:

$$\tilde{C} = \sum_{i=1}^{N} \tilde{\lambda}_i(\tilde{\tau}) u_i u_i',$$

where $\tilde{\lambda}_i(\tilde{\tau})$ represents the shrunk eigenvalues.

2.5 Factor Copula Model

The methods above focus on estimating both volatility and correlation simultaneously. However, recent advancements have seen the rise of factor copula models, which enable the estimation of dependence structures independently from volatility. Notably, Creal and Tsay (2015), Oh and Patton (2017, 2018), and Lucas, Schwaab, and Zhang (2017) have proposed different models to extend the applicability of factor copula to high cross-sectional dimensions. These models introduce dynamic loadings within the single-factor

framework. Opschoor et al. (2021) present a series of multi-factor copula models, wherein observable characteristics drive the dynamic factor loadings. The primary objective of these models is to decompose the conditional joint distribution $G_t(r_t)$ of asset returns vector r_t into N marginals and a conditional copula:

$$r_t|\mathcal{F}_{t-1} \sim G_t(r_t)$$

$$= C_t(G_{1,t}(r_{1,t};\theta_{M,1,t}),\dots,G_{N,t}(r_{N,t};\theta_{M,N,t});\theta_{C,t}),$$

where $C_t(\cdot;\theta_{C,t})$ represents the conditional copula given the information set $\mathcal{F}_{t-1} = \sigma(r_{t-1}, r_{t-2},\dots)$ and the parameter vector $\theta_{C,t}$. Meanwhile, $G_{i,t}(r_{i,t};\theta_{M,i,t})$ with $i = 1,\dots,N$ denotes the conditional marginal distribution of asset i given $\mathcal{F}_{t-1}$ and the time-varying marginal distribution parameter vector $\theta_{M,i,t}$.

The conditional copula can also be interpreted as the conditional distribution $C_t(u_t,\theta_{C,t})$ of the probability integral transforms (PITs) $u_t = (u_{1,t},\dots,u_{N,t})^\top$ of r_t, where $u_{i,t} = G_{i,t}(r_{i,t};\theta_{M,i,t})$ for $i = 1,\dots,N$.

Opschoor et al. (2021) assume that the N asset returns can be clustered into K groups, with the assets in the same group having identical factor loadings. The copula structure proposed by Opschoor et al. (2021) can be summarized as:

$$u_{i,t} = G_{x,i}(x_{i,t};\tilde{\xi}_{i,t},\sigma_{i,t},\psi_C),$$

$$x_{i,t} = \tilde{\xi}_{i,t}^\top z_t + \sigma_{i,t}\epsilon_{i,t},$$

$$z_t \sim G_z(z_t|\psi_C),\ \epsilon_{i,t} \sim G_\epsilon(\epsilon_{i,t}|\psi_C),$$

where $x_{i,t}$ represents a latent random variable, z_t is a vector of common latent factors, $\epsilon_{i,t}$ is an idiosyncratic shock, and $\tilde{\xi}_{i,t}$ is a vector of scaled factor loadings. Specifically, they define:

$$\tilde{\xi}_{i,t} = \frac{\xi_{i,t}}{\sqrt{1+\xi_{i,t}^\top\xi_{i,t}}},$$

with $\xi_{i,t}$ representing an unrestricted vector of factor loadings such that $x_{i,t}$ has zero mean and unit variance. Moreover, $\sigma_{i,t}$ can be defined as the square root of $\sigma_{i,t}^2 = \frac{1}{1+\xi_{i,t}^\top\xi_{i,t}}$. Therefore, the correlation matrix of $x_t = (x_{1,t},\dots,x_{N,t})^\top$ is equal to:

$$P_t = \tilde{L}_t^\top\tilde{L}_t + D_t,$$

$$D_t = \mathrm{diag}(\sigma_{1,t}^2,\dots,\sigma_{N,t}^2).$$

The loadings' matrix $\tilde{L}_t^\top$ can be specified in various ways, leading to a more specific

model. From the models discussed by Opschoor et al. (2021), I consider the 2-factor (2F) copula model, which assumes a single common factor with (common) equiloadings and a common factor with $K+1$ group-specific loadings. For instance, considering $K = 4$ groups with 2 firms in each group, it can be expressed as:

$$\tilde{L}_t^T = \begin{pmatrix} \tilde{\xi}_{1,t} & \tilde{\xi}_{2,1,t} \\ \tilde{\xi}_{1,t} & \tilde{\xi}_{2,2,t} \\ \tilde{\xi}_{1,t} & \tilde{\xi}_{2,3,t} \\ \tilde{\xi}_{1,t} & \tilde{\xi}_{2,4,t} \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

where $\otimes$ is the Kronecker product.

Opschoor et al. (2021) integrate convenient distributional assumptions for the factors z_t and $\epsilon_{i,t}$, alongside a score-driven transition equation for the factor loadings $\xi_{i,t}$. These assumptions facilitate the expression of the copula structure as follows:

$$u_{i,t} = T(x_{i,t}; \nu_C),$$

$$x_{i,t} = \sqrt{\zeta_t}(\tilde{\xi}_{i,t}^\top z_t + \sigma_{i,t}\epsilon_{i,t}),$$

here $T(\cdot, \nu_C)$ represents the cumulative distribution function (CDF) of the univariate Students's t distribution with ν_C degrees of freedom, centered at zero, and with unit scale. Additionally, $z_t \sim N(0, I_k)$, $\epsilon_{i,t} \sim N(0, 1)$, $\zeta_t \sim$Inv-Gamma$(\frac{\nu_C}{2}, \frac{\nu_C}{2})$. In the scenario where $\nu_C \to \infty$, ζ_t is equal to 1, yielding a Gaussian copula setting.

Another important component involves specifying the dynamics of the unique factor loadings. Here, the authors employ an observation-driven approach, wherein factor loadings are contingent on functions of past observables. Particularly, they use score-driven dynamics, as introduced by Creal, Koopman, and Lucas (2013). In this framework, defining the unique time-varying parameters as f_t, the score-driven dynamics leverage the score of the conditional copula density to drive f_t. This approach ensures that the factor loadings evolve in response to the observed data's evolving characteristics, capturing the underlying processes' dynamic nature. The transition equation of f_t is the following:

$$f_{t+1} = \omega + As_t + Bf_t,$$

where $s_t = \frac{\partial \log c(x_t; \tilde{\xi}_t; \nu_C)}{\partial f_t}$ and $c(\cdot, \tilde{\xi}_t, \nu_C)$ is the Student's t copula density.

The last element to specify is related to the marginal distributions. Specifically, I assume that they are distributed following a GARCH(1,1) model to maintain coherence with the other dynamic model (DCC-NL). However, this model assumes the normality of the error terms when modeling the variance, which can have several disadvantages. Firstly,

it leads to an increased susceptibility to outliers, as the model's reliance on normality may not adequately capture the heavy-tailed nature of financial data. This can result in an underestimation of the likelihood and impact of extreme events, such as market crashes or volatility spikes. Furthermore, the model's forecast accuracy may be compromised due to misspecification when the true distribution of conditional variances deviates from Gaussian. In conclusion, analyzing various assumptions regarding volatility can yield differing results, necessitating further research.

3 Simulation Analysis

To assess the performance of the models outlined in Section 2, I conduct a study involving simulated paths for both asset return variances and correlations. Specifically, I examine three distinct paths: static, with structural breaks, and AR(1). For each simulated path, I construct a covariance matrix for each period that reflects the true covariance for the portfolio under consideration at a point in time. In this matrix, the diagonal entries represent the variances $\sigma_{i,t}^2$, while the off-diagonal elements correspond to the covariances $\sigma_{ij,t} = \rho_{ij,t}\sigma_{i,t}\sigma_{j,t}$. Therefore, the covariance matrix can be expressed in the following form:

$$\Sigma_t = \begin{bmatrix} \sigma_{1,t}^2 & \dots & \sigma_{1N,t} \\ \vdots & \ddots & \vdots \\ \sigma_{N1,t} & \dots & \sigma_{N,t}^2 \end{bmatrix}.$$

It is crucial to note that assets are categorized into different groups, with the assumption that assets within the same group share the same correlation path. This assumption is made to maintain consistency with the industry clustering approach of GCVC shrinkage and 2-factor copula models, as well as with the approach used in the empirical framework (Section 4), which considers group clustering based on the industry sector.

Consequently, I generate a time series of returns assuming that they are normally distributed with a mean of zero for all the assets and a covariance matrix equivalent to the true covariance matrix for that period. Subsequently, this time series of observations is utilized as input to estimate the covariance matrices according to the aforementioned models. The analysis is carried out via Monte Carlo (MC) simulation, with 1000 replications, in order to minimize sampling errors. To streamline the initial analysis, the simulation focuses on a portfolio composed of 4, 10, and 20 assets. The above steps are executed for three distinct sample lengths, comprising 1000, 2500, and 5000 observations, to account for potential effects arising from the length of the time series.

After estimating variances and correlations across different paths and sample lengths, I conduct a visual analysis by creating figures that depict the estimated paths along with MC confidence bands at a confidence level of 90% across replications. To gain deeper insights, I employ two performance measures: average Root Mean Squared Error (RMSE) and average Mean Absolute Error (MAE). The former represents the average magnitude of the errors between estimated values and true values, instead the latter measures the average absolute deviations of the estimated values from the true ones. Notably, RMSE is more sensitive to outliers compared to MAE due to its construction, which is based on squaring the errors. Additionally, only for correlations, I calculate the average Mean Absolute Percentage Error (MAPE) to provide a more nuanced understanding. It com-

putes the average percentage difference between estimated and true values, relative to the true ones. These measures are obtained through a Monte Carlo simulation with 1000 replications, where a new simulated path is generated in each iteration. This approach mitigates potential errors arising from a single-path simulation. It is noteworthy that, for the computation of performance measures, the first 10% of observations have been discarded. This exclusion aims to eliminate distortions associated with the initialization of the models, ensuring a more accurate evaluation of performance across the remaining periods. Finally, I report the absolute values of these measures alongside their relative values compared to the CVC shrinkage approach. This comparative analysis facilitates a clearer evaluation of the different models.

3.1 Static Paths

The initial simulated path, termed static, assumes constant values for both variances and correlations during all time periods. Investigating these paths is crucial not only for evaluating the performance of static models like the sample covariance, CVC, and GCVC shrinkage approaches but also for understanding the adaptability of time-varying models. It provides insights into how these models behave in scenarios where the asset value parameters remain constant, offering valuable comparisons for assessing their versatility.

Figure 1 compares the true volatility with the estimated ones for a representative example, including the 90% Monte Carlo (MC) confidence bands.

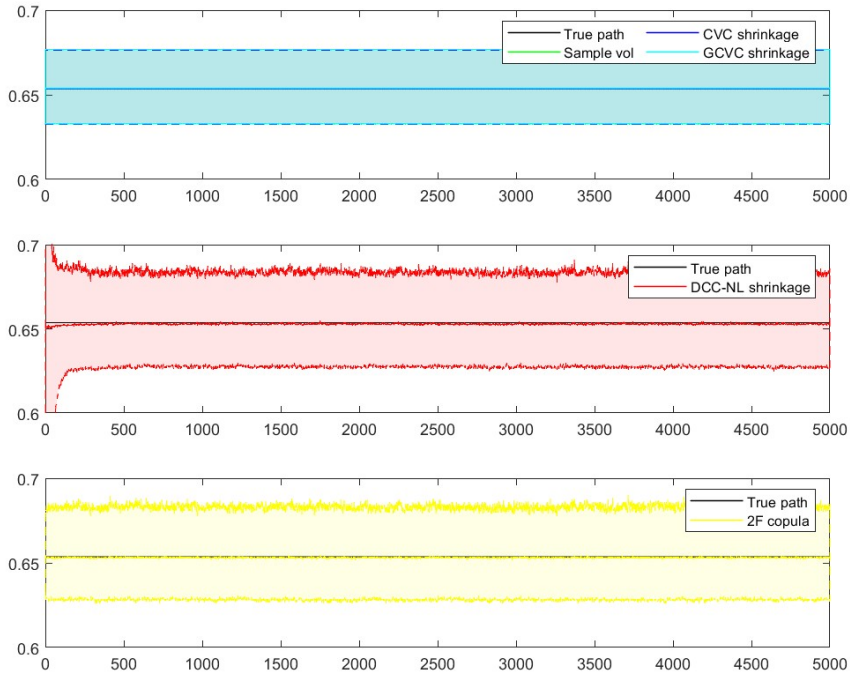

Figure 1: Estimated variances in the static scenario

The analysis reveals the effectiveness of all the estimation approaches for capturing the variance's true Data Generating Process (DGP). They produce highly accurate median estimates and the MC confidence intervals successfully capture the true paths. As expected, the three static models achieved the most precise path estimates, evident in both the median values and MC confidence bands. Interestingly, the two time-varying models also demonstrate good estimation properties. However, their median values deviate slightly from the true path. This versatility suggests their potential adaptability to a wider range of scenarios, even when the underlying process is static.

Figure 2 compares estimated and true correlations for a representative example.

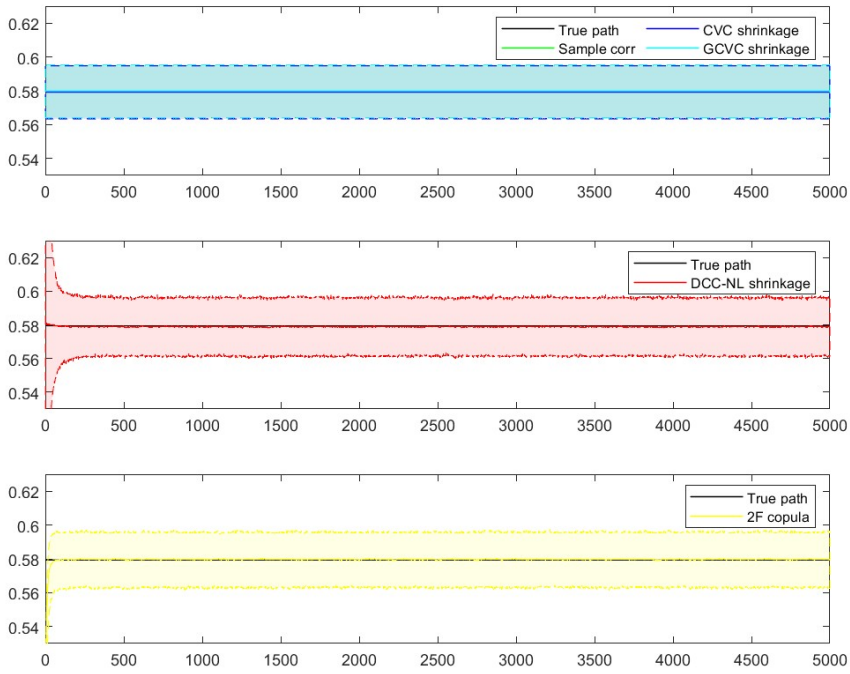

Figure 2: Estimated correlations in the static scenario

Similar to the observed performance with variances, the three static methodologies demonstrate good estimation for correlation paths. The two dynamic models, while exhibiting wider Monte Carlo (MC) confidence bands in the initial periods, still produce good estimates for the static path in terms of median values. This initial widening of confidence bands is likely due to the time required for the dynamic models to adapt to the data.

To quantify the fit, I construct two performance measures for the volatility paths: the average Root Mean Squared Error (RMSE) and average Mean Absolute Error (MAE). The results of these errors are summarized in Tables 1 and 2. They suggest that the static models, including sample variance, CVC, and GCVC shrinkage, exhibit lower errors compared to dynamic models. This finding is in line with my expectations and supported by graphical analysis. However, it is noteworthy that among the static models, CVC

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.0147	0.0097	0.0072	0.9597	0.9827	0.9902
CVC shrinkage	0.0153	0.0099	0.0073	1.0000	1.0000	1.0000
GCVC shrinkage	0.0147	0.0097	0.0072	0.9731	0.9840	0.9909
DCC-NL shrinkage	0.0206	0.0141	0.0101	1.3712	1.4304	1.3981
2F Copula	0.0200	0.0138	0.0099	1.3341	1.3949	1.3699
	N = 10					
Sample covariance	0.0194	0.0146	0.0115	0.9731	0.9899	0.9927
CVC shrinkage	0.0200	0.0147	0.0116	1.0000	1.0000	1.0000
GCVC shrinkage	0.0194	0.0146	0.0115	0.9738	0.9895	0.9927
DCC-NL shrinkage	0.0260	0.0181	0.0140	1.3005	1.2310	1.2047
2F Copula	0.0254	0.0176	0.0138	1.2721	1.1963	1.1916
	N = 20					
Sample covariance	0.0225	0.0181	0.0159	0.9777	0.9841	0.9926
CVC shrinkage	0.0230	0.0184	0.0160	1.0000	1.0000	1.0000
GCVC shrinkage	0.0225	0.0181	0.0159	0.9778	0.9847	0.9930
DCC-NL shrinkage	0.0279	0.0216	0.0181	1.2123	1.1751	1.1303
2F Copula	0.0271	0.0213	0.0180	1.1799	1.1577	1.1227

Table 1: Average RMSE of the estimated variances in the static scenario

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.0147	0.0097	0.0072	0.9597	0.9827	0.9902
CVC shrinkage	0.0153	0.0099	0.0073	1.0000	1.0000	1.0000
GCVC shrinkage	0.0147	0.0097	0.0072	0.9596	0.9840	0.9909
DCC-NL shrinkage	0.0182	0.0122	0.0089	1.1874	1.2418	1.2312
2F Copula	0.0177	0.0120	0.0087	1.1523	1.2125	1.2050
	N = 10					
Sample covariance	0.0194	0.0146	0.0115	0.9731	0.9899	0.9927
CVC shrinkage	0.0200	0.0147	0.0116	1.0000	1.0000	1.0000
GCVC shrinkage	0.0194	0.0146	0.0115	0.9738	0.9895	0.9927
DCC-NL shrinkage	0.0230	0.0165	0.0128	1.1509	1.1169	1.1026
2F Copula	0.0226	0.0161	0.0127	1.1313	1.0905	1.0913
	N = 20					
Sample covariance	0.0225	0.0181	0.0159	0.9777	0.9841	0.9926
CVC shrinkage	0.0230	0.0184	0.0160	1.0000	1.0000	1.0000
GCVC shrinkage	0.0225	0.0181	0.0159	0.9778	0.9847	0.9930
DCC-NL shrinkage	0.0252	0.0199	0.0170	1.0962	1.0811	1.0626
2F Copula	0.0246	0.0196	0.0169	1.0708	1.0655	1.0555

Table 2: Average MAE of the estimated variances in the static scenario

shrinkage consistently demonstrates higher errors across all dimensions. Additionally, the dynamic models show slightly higher error values compared to the static models across all time periods. Notably, increasing the number of observations tends to decrease error levels for all estimated models, as observed in both RMSE and MAE measures. This finding suggests that increasing the sample length in the static scenario results in better estimation accuracy. Additionally, analyzing performance measures across varying asset counts reveals a concerning trend in the static scenario. All models exhibit increasing errors in volatility estimation as dimensionality rises. However, the relative analysis indicates that as the number of assets increases, the dynamic estimate errors tend to decrease further, aligning closely with the errors generated by static models.

For the correlations, I also compute the average Mean Absolute Percentage Error (MAPE) and the results for all the performance measures are reported in Tables 3, 4 and 5.

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.0188	0.0119	0.0088	0.9011	0.9448	0.9713
CVC shrinkage	0.0209	0.0126	0.0091	1.0000	1.0000	1.0000
GCVC shrinkage	0.0190	0.0120	0.0088	0.9098	0.9525	0.9736
DCC-NL shrinkage	0.0235	0.0147	0.0108	1.1229	1.1666	1.1914
2F Copula	0.0748	0.0628	0.0630	3.5809	4.9855	6.9563
	N = 10					
Sample covariance	0.0266	0.0200	0.0173	0.9210	0.9860	0.9943
CVC shrinkage	0.0289	0.0202	0.0174	1.0000	1.0000	1.0000
GCVC shrinkage	0.0266	0.0199	0.0173	0.9212	0.9829	0.9932
DCC-NL shrinkage	0.0250	0.0190	0.0166	0.8672	0.9410	0.9531
2F Copula	0.0471	0.0408	0.0401	1.6321	2.0181	2.2978
	N = 20					
Sample covariance	0.0298	0.0260	0.0233	0.9671	1.0002	0.9970
CVC shrinkage	0.0308	0.0260	0.0234	1.0000	1.0000	1.0000
GCVC shrinkage	0.0297	0.0260	0.0233	0.9643	0.9975	0.9955
DCC-NL shrinkage	0.0278	0.0244	0.0220	0.9031	0.9377	0.9398
2F Copula	0.0381	0.0361	0.0328	1.2372	1.3874	1.4023

Table 3: Average RMSE of the estimated correlations in the static scenario

The analysis of performance measures for estimated correlations in the static scenario reveals similar results to those obtained for variances but with slightly higher error values. However, the 2-factor copula model exhibits a distinct behavior compared to the variance case. Notably, its error measures for correlations show significant spikes across all of them. These spikes tend to diminish with the increasing number of assets, as more information becomes available to estimate the dependency structure. In contrast, the

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.0188	0.0119	0.0088	0.9011	0.9448	0.9713
CVC shrinkage	0.0209	0.0126	0.0091	1.0000	1.0000	1.0000
GCVC shrinkage	0.0190	0.0120	0.0088	0.9098	0.9525	0.9736
DCC-NL shrinkage	0.0217	0.0135	0.0099	1.0373	1.0692	1.0932
2F Copula	0.0741	0.0624	0.0627	3.5454	4.9516	6.9246
	N = 10					
Sample covariance	0.0266	0.0200	0.0173	0.9210	0.9860	0.9943
CVC shrinkage	0.0289	0.0202	0.0174	1.0000	1.0000	1.0000
GCVC shrinkage	0.0266	0.0199	0.0173	0.9212	0.9829	0.9932
DCC-NL shrinkage	0.0244	0.0185	0.0162	0.8446	0.9156	0.9307
2F Copula	0.0466	0.0406	0.0399	1.6144	2.0042	2.2844
	N = 20					
Sample covariance	0.0298	0.0260	0.0233	0.9671	1.0002	0.9970
CVC shrinkage	0.0308	0.0260	0.0234	1.0000	1.0000	1.0000
GCVC shrinkage	0.0297	0.0260	0.0233	0.9643	0.9975	0.9955
DCC-NL shrinkage	0.0274	0.0242	0.0218	0.8908	0.9287	0.9314
2F Copula	0.0374	0.0358	0.0325	1.2156	1.3752	1.3879

Table 4: Average MAE of the estimated correlations in the static scenario

Average MAPE	Absolute value			Relative		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.0497	0.0302	0.0231	0.9340	0.9578	0.9796
CVC shrinkage	0.0532	0.0315	0.0235	1.0000	1.0000	1.0000
GCVC shrinkage	0.0499	0.0303	0.0231	0.9372	0.9634	0.9812
DCC-NL shrinkage	0.0568	0.0341	0.0259	1.0682	1.0823	1.1013
2F Copula	0.2222	0.1823	0.1863	4.1776	5.7899	7.9166
	N = 10					
Sample covariance	0.0741	0.0550	0.0501	0.9522	0.9920	0.9974
CVC shrinkage	0.0778	0.0555	0.0502	1.0000	1.0000	1.0000
GCVC shrinkage	0.0740	0.0549	0.0500	0.9513	0.9894	0.9962
DCC-NL shrinkage	0.0685	0.0516	0.0474	0.8803	0.9310	0.9453
2F Copula	0.1374	0.1178	0.1202	1.7766	2.1234	2.3950
	N = 20					
Sample covariance	0.0841	0.0764	0.0678	0.9854	1.0083	1.0001
CVC shrinkage	0.0853	0.0758	0.0678	1.0000	1.0000	1.0000
GCVC shrinkage	0.0838	0.0762	0.0677	0.9825	1.0053	0.9986
DCC-NL shrinkage	0.0790	0.0720	0.0644	0.9254	0.9506	0.9508
2F Copula	0.1105	0.1078	0.0961	1.2946	1.4223	1.4173

Table 5: Average MAPE of the estimated correlations in the static scenario

behavior of errors for the other models aligns with the variance trend, where an increase in dimensionality leads to larger errors.

In conclusion, the static models consistently provide the most accurate estimates for both variance and correlation, as expected. While the dynamic models are initially less accurate, they tend to converge towards the static models' accuracy as the number of assets increases.

3.2 Paths with Structural Breaks

Following the initial simulation of static paths for both variances and correlations, the simulation framework is extended to include structural breaks. These breaks entail intervals of static values interrupted by alterations, either increases or decreases, in values. This approach allows for the maintenance of the static scenario while accommodating fluctuations in volatility and correlation during particular periods.

The examination of these paths proceeds as previously described, with both the true path and estimated ones for the variances depicted in Figure 3.

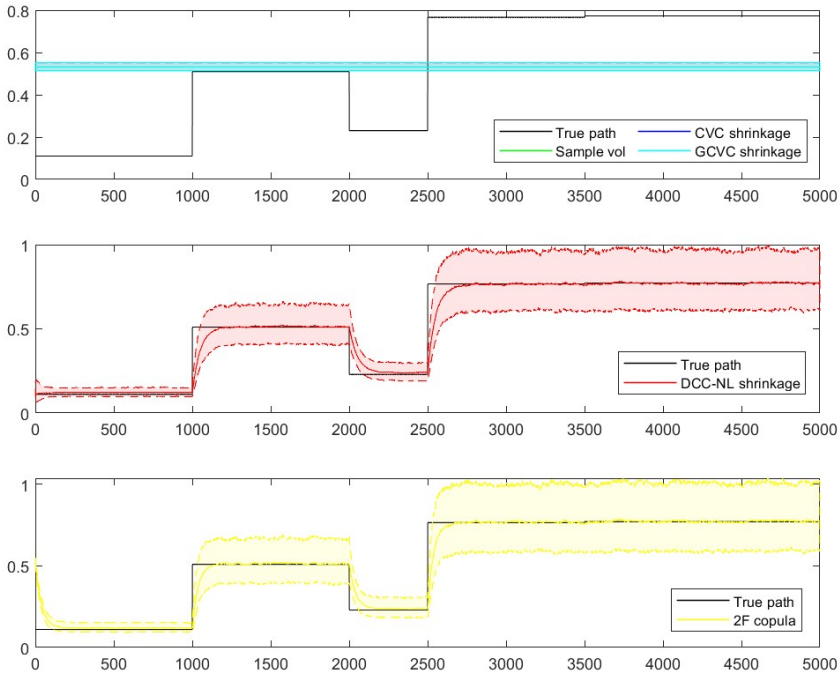

Figure 3: Estimated variances in the structural breaks scenario

The initial observation of the graph highlights an obvious shortcoming in estimating variances using static approaches. These models provide a single value for the variance with small MC confidence bands, overlooking variations indicated by structural breaks. In contrast, both the DCC-NL shrinkage and 2-factor copula approaches showcase adaptability

by considering such fluctuations in their estimations. However, there is a noticeable lag in capturing volatility changes, with the models needing multiple periods to converge to the true value. Another consideration with dynamic methods is the behavior of MC confidence bands. They tend to widen as volatility levels increase, contrasting with narrower intervals when volatility is low.

Figure 4 illustrates the true correlation path alongside the estimated ones for a representative example.

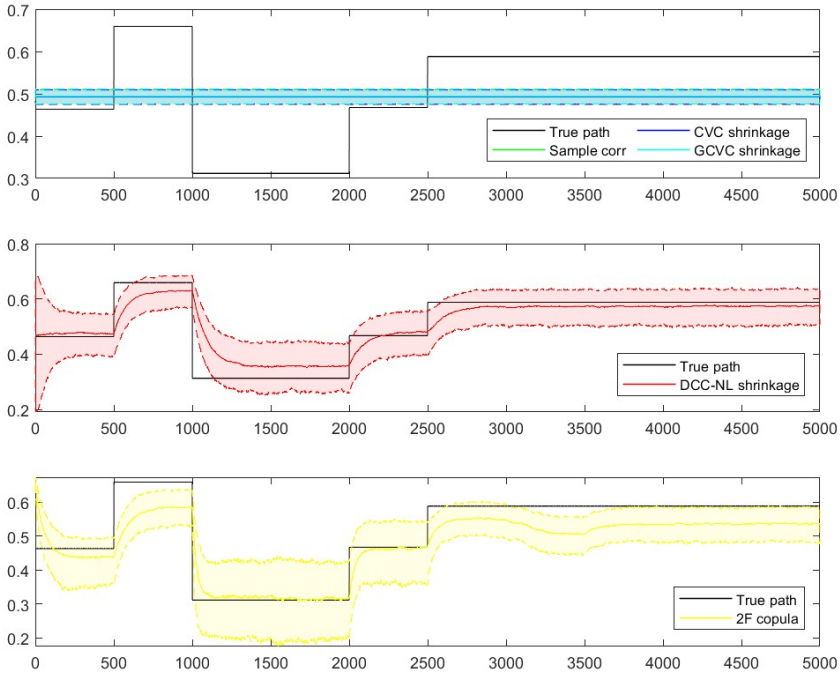

Figure 4: Estimated correlations in the structural breaks scenario

The correlation analysis mirrors the findings observed in the volatility analysis, indicating inaccuracies in estimated values for the static models due to their single-value outputs. Conversely, both the DCC-NL shrinkage and 2-factor copula methods show a superior estimation of true correlation. However, they present a delay in converging to the actual values in the initial periods after a change in correlation, where the models may fail to adjust promptly to the new dynamics. It is worth noting that the median values of estimated correlations through the DCC-NL shrinkage approach tend to underestimate the true path for higher values and overestimate it for lower ones. In contrast, the 2-factor copula exhibits a different trend, characterized by only underestimation in the case of higher correlation values with even the MC confidence bands hardly capturing the true path.

The examination of performance measures is also applied to the new paths. In particular, Tables 6 and 7 present the estimation errors for the variance.

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.1544	0.1535	0.1497	0.9963	0.9988	0.9995
CVC shrinkage	0.1549	0.1537	0.1498	1.0000	1.0000	1.0000
GCVC shrinkage	0.1544	0.1535	0.1497	0.9964	0.9989	0.9996
DCC-NL shrinkage	0.0924	0.0759	0.0632	0.5963	0.4942	0.4220
2F Copula	0.0969	0.0789	0.0657	0.6253	0.5136	0.4383
	N = 10					
Sample covariance	0.1495	0.1470	0.1477	0.9971	0.9987	0.9992
CVC shrinkage	0.1500	0.1472	0.1479	1.0000	1.0000	1.0000
GCVC shrinkage	0.1496	0.1470	0.1477	0.9971	0.9988	0.9992
DCC-NL shrinkage	0.0925	0.0748	0.0660	0.6168	0.5084	0.4463
2F Copula	0.0961	0.0780	0.0684	0.6409	0.5302	0.4628
	N = 20					
Sample covariance	0.1550	0.1534	0.1535	0.9963	0.9992	0.9994
CVC shrinkage	0.1555	0.1536	0.1536	1.0000	1.0000	1.0000
GCVC shrinkage	0.1550	0.1534	0.1535	0.9964	0.9992	0.9994
DCC-NL shrinkage	0.0984	0.0799	0.0678	0.6329	0.5201	0.4415
2F Copula	0.1032	0.0832	0.0706	0.6638	0.5415	0.4594

Table 6: Average RMSE of the estimated variances in the structural breaks scenario

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.1362	0.1357	0.1315	0.9923	0.9969	0.9988
CVC shrinkage	0.1373	0.1361	0.1317	1.0000	1.0000	1.0000
GCVC shrinkage	0.1363	0.1357	0.1315	0.9927	0.9971	0.9989
DCC-NL shrinkage	0.0660	0.0527	0.0432	0.4805	0.3872	0.3284
2F Copula	0.0695	0.0552	0.0454	0.5062	0.4053	0.3452
	N = 10					
Sample covariance	0.1305	0.1287	0.1297	0.9911	0.9962	0.9983
CVC shrinkage	0.1317	0.1292	0.1299	1.0000	1.0000	1.0000
GCVC shrinkage	0.1306	0.1287	0.1297	0.9914	0.9963	0.9984
DCC-NL shrinkage	0.0668	0.0527	0.0463	0.5070	0.4077	0.3561
2F Copula	0.0695	0.0553	0.0484	0.5275	0.4285	0.3727
	N = 20					
Sample covariance	0.1351	0.1347	0.1353	0.9902	0.9966	0.9983
CVC shrinkage	0.1364	0.1351	0.1356	1.0000	1.0000	1.0000
GCVC shrinkage	0.1351	0.1347	0.1353	0.9905	0.9967	0.9984
DCC-NL shrinkage	0.0717	0.0567	0.0475	0.5256	0.4196	0.3503
2F Copula	0.0753	0.0593	0.0498	0.5520	0.4392	0.3675

Table 7: Average MAE of the estimated variances in the structural breaks scenario

The obtained results confirm the earlier visual analysis. Notably, the three static approaches exhibit higher values for both metrics compared to the two dynamic approaches. Moreover, it is apparent that the three static methods yield very similar estimation errors, mirroring the findings observed also in the dynamic models. Furthermore, akin to the static scenario, an increase in the number of observations leads to a reduction in both average RMSE and MAE for the dynamic models. In contrast, the static ones do not exhibit a clear trend due to their inherent structure which is unable to adapt to changes in values. Additionally, while increasing dimensionality does not produce a clear pattern of increasing errors for static models, it does lead to a noticeable trend of increasing errors for dynamic ones.

A comprehensive overview of the performance metrics for estimated correlations is depicted in Tables 8, 9, and 10.

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.1334	0.1535	0.1297	0.9988	0.9991	0.9996
CVC shrinkage	0.1336	0.1537	0.1297	1.0000	1.0000	1.0000
GCVC shrinkage	0.1322	0.1535	0.1294	0.9895	0.9952	0.9976
DCC-NL shrinkage	0.0910	0.0759	0.0656	0.6810	0.5791	0.5060
2F Copula	0.1032	0.0789	0.0819	0.7726	0.6931	0.6314
	N = 10					
Sample covariance	0.1288	0.1279	0.1284	0.9949	0.9968	0.9985
CVC shrinkage	0.1295	0.1283	0.1286	1.0000	1.0000	1.0000
GCVC shrinkage	0.1288	0.1279	0.1285	0.9949	0.9968	0.9985
DCC-NL shrinkage	0.0881	0.0736	0.0669	0.6808	0.5736	0.5204
2F Copula	0.0903	0.0795	0.0778	0.6977	0.6197	0.6051
	N = 20					
Sample covariance	0.1284	0.1304	0.1249	0.9925	0.9977	0.9987
CVC shrinkage	0.1294	0.1307	0.1251	1.0000	1.0000	1.0000
GCVC shrinkage	0.1284	0.1304	0.1249	0.9924	0.9977	0.9987
DCC-NL shrinkage	0.0898	0.0782	0.0683	0.6941	0.5988	0.5462
2F Copula	0.0808	0.0751	0.0686	0.6244	0.5744	0.5482

Table 8: Average RMSE of the estimated correlations in the structural breaks scenario

Similar to the findings with variance estimates, the analysis of performance measures for the correlations shows that static approaches consistently have higher estimation errors compared to time-varying ones. The DCC-NL shrinkage approach remains the superior method for correlation estimation, as it consistently produces lower estimation errors across all three measures. However, it is worth noting that the increase in dimensionality produces a reduction in errors for all the models. Specifically, at the higher number of assets considered, the 2-factor copula model surpasses the accuracy performance of the

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.1155	0.1137	0.1120	0.9965	0.9982	0.9992
CVC shrinkage	0.1159	0.1140	0.1121	1.0000	1.0000	1.0000
GCVC shrinkage	0.1146	0.1134	0.1119	0.9881	0.9948	0.9975
DCC-NL shrinkage	0.0717	0.0584	0.0492	0.6187	0.5124	0.4389
2F Copula	0.0841	0.0734	0.0658	0.7257	0.6440	0.5869
	N = 10					
Sample covariance	0.1117	0.1119	0.1119	0.9925	0.9955	0.9980
CVC shrinkage	0.1125	0.1124	0.1121	1.0000	1.0000	1.0000
GCVC shrinkage	0.1117	0.1119	0.1119	0.9925	0.9955	0.9980
DCC-NL shrinkage	0.0692	0.0562	0.0505	0.6148	0.5001	0.4503
2F Copula	0.0713	0.0621	0.0612	0.6335	0.5528	0.5461
	N = 20					
Sample covariance	0.1123	0.1139	0.1090	0.9894	0.9963	0.9981
CVC shrinkage	0.1135	0.1144	0.1092	1.0000	1.0000	1.0000
GCVC shrinkage	0.1123	0.1139	0.1090	0.9894	0.9963	0.9981
DCC-NL shrinkage	0.0715	0.0606	0.0518	0.6294	0.5297	0.4746
2F Copula	0.0633	0.0583	0.0533	0.5573	0.5102	0.4884

Table 9: Average MAE of the estimated correlations in the structural breaks scenario

Average MAPE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.2494	0.2506	0.2469	1.0002	0.9995	0.9999
CVC shrinkage	0.2493	0.2507	0.2469	1.0000	1.0000	1.0000
GCVC shrinkage	0.2494	0.2507	0.2470	1.0005	0.9998	1.0002
DCC-NL shrinkage	0.1675	0.1391	0.1178	0.6717	0.5549	0.4771
2F Copula	0.1834	0.1617	0.1433	0.7355	0.6450	0.5805
	N = 10					
Sample covariance	0.2528	0.2504	0.2550	0.9987	0.9987	0.9993
CVC shrinkage	0.2531	0.2507	0.2552	1.0000	1.0000	1.0000
GCVC shrinkage	0.2528	0.2504	0.2550	0.9984	0.9986	0.9992
DCC-NL shrinkage	0.1713	0.1412	0.1319	0.6766	0.5632	0.5167
2F Copula	0.1888	0.1672	0.1706	0.7459	0.6668	0.6686
	N = 20					
Sample covariance	0.2594	0.2646	0.2546	0.9979	0.9997	1.0000
CVC shrinkage	0.2599	0.2647	0.2546	1.0000	1.0000	1.0000
GCVC shrinkage	0.2593	0.2646	0.2545	0.9976	0.9996	0.9999
DCC-NL shrinkage	0.1837	0.1585	0.1384	0.7069	0.5990	0.5438
2F Copula	0.1663	0.1569	0.1454	0.6399	0.5927	0.5713

Table 10: Average MAPE of the estimated correlations in the structural breaks scenario

DCC-NL for the smallest sample lengths, becoming the best model in this framework. Additionally, for the time-varying models, the increasing sample length produces a better accuracy of the models. However, for static models, this pattern is not apparent as also observed in the variance case.

Finally, the analysis of paths with structural breaks confirms the expected limitation of static models in capturing dynamic scenarios. Conversely, time-varying models demonstrate the ability to effectively handle such paths. Furthermore, for time-varying models, increasing the number of observations reduces the error for both variance and correlation. Interestingly, while increasing the number of assets also reduces the errors for correlation, it has the opposite effect on variance, leading to increased errors.

3.3 AR(1) Paths

The last paths simulated are time-varying paths constructed starting from an autoregressive (AR) process of order 1 but with two different specifications for the variances and correlations. The DGP for the variances σ_i^2 of asset returns is based on a Stochastic Volatility (SV) model, expressed as:

$$y_t = \epsilon_t = \eta_t \exp\left(\frac{1}{2}\sigma_t\right) ,$$

$$\sigma_{t+1} = \omega + \beta(\sigma_t - \omega) + \zeta_t ,$$

where $\eta_t \sim N(0,1)$, $\zeta_t \sim N(0, \sigma_\zeta)$ and the log volatility follows an AR(1) process. The parameter ω represents the unconditional mean of the log volatility, assumed to be $\log(15\%)$ based on the average volatility in the stock market. The values of β and σ_ζ are set to 0.99 and 0.10 to account for the high observed persistence of volatility.

The correlation paths are simulated using a logistic transformation of an AR(1) process, represented as:

$$X_{t+1} = \alpha + \phi(X_t - \alpha) + \gamma_t ,$$

where $\gamma \sim N(0, \sigma_\gamma^2)$. The value of the unconditional mean α is chosen to ensure that the correlations remain within the typical range of [0.2, 0.8], which is commonly observed for stock correlations. The coefficient ϕ is assumed to be positive and set to 0.995, while the shocks are assumed to have a variance σ_γ^2 of 0.15.

After simulating the paths, I conduct an analysis as before, utilizing graphs and tables to elucidate the relationship between the true path and the estimated ones. Figure 5 compares the true volatilities with the estimated ones. The visual analysis reveals that even in the scenario of a time-varying path for volatilities, the sample covariance, CVC, and GCVC approaches present similar values in terms of both median value and Monte

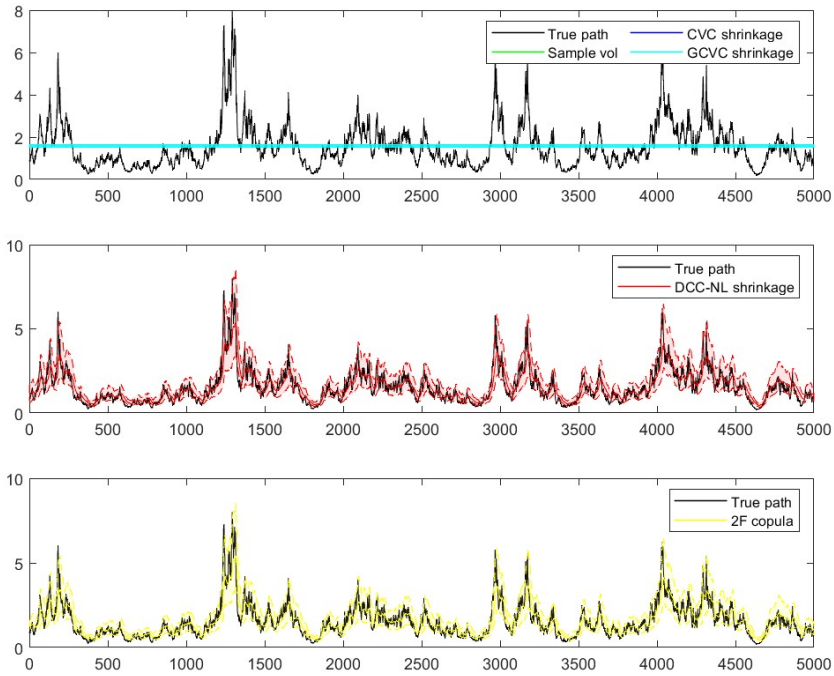

Figure 5: Estimated variances in the time-varying scenario

Carlo (MC) confidence bands. They tend to remain close to the unconditional mean value assumed for the path simulation. The two dynamic models, DCC-NL and 2-factor copula, effectively capture the variability of DGP and present the same results due to the models' construction. The MC confidence bands encompass nearly all fluctuations in value, with exceptions occurring only during rapid changes where the models may require some time to capture them, as shown in the paths with structural breaks.

The comparison between actual paths and estimated ones can also be extended to the correlations. Specifically, Figure 6 illustrates the estimated correlations among assets within the same group. The analysis of these estimated paths leads to the conclusion that the DCC-NL shrinkage model emerges as better equipped to capture the true paths, as its MC confidence interval encompasses them in most cases. However, challenges arise when the changes in values are particularly extreme, with the MC confidence bands initially failing to capture those values. Another significant consideration pertains to the width of these bands, which are wider when the correlation values are low and narrower when the values are high. On the other hand, the 2-factor copula model encounters more difficulties in capturing the true correlation path. Specifically, the median value finds difficulties in following the true path, and the MC confidence bands struggle to accurately include the true correlation in some cases. Additionally, it is worth mentioning that the model tends to overestimate the correlation values when they are low, while it tends to more accurately capture high correlation values, even though with tighter MC confidence bands.

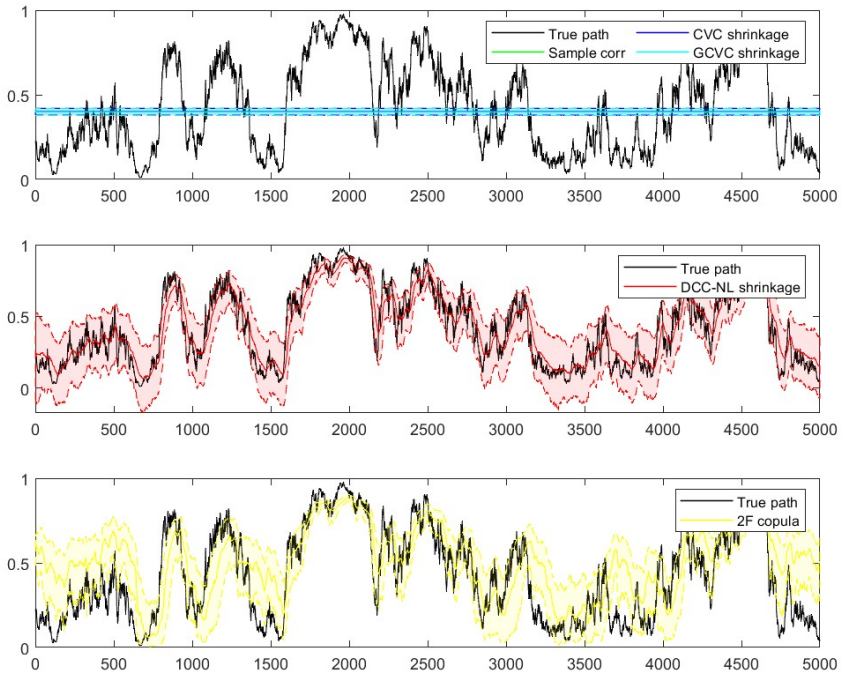

Figure 6: Estimated within-group correlations in the time-varying scenario

The graphical analysis can be also extended to the estimated correlations among assets in different groups, as depicted in Figure 7. The analysis highlights notable distinctions with the within-group correlations, primarily concerning the 2-factor copula approach. In periods of high persistence, the median values more effectively capture the true correlations. However, challenges arise when extreme values occur, with the MC confidence bands initially failing to capture significant changes. Whereas, the DCC-NL shrinkage approach in estimating between-group correlations yields results consistent with those observed for the within-group correlations. Additionally, it is also worth mentioning that this model tends to have wider MC confidence intervals than the ones observed for the 2-factor copula model.

To complete the analysis and better understand the performance of these models, the average Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) are computed for the volatility paths. The results of these errors are summarized in Tables 11 and 12. The analysis of these performance measures in the time-varying scenario reveals, in terms of model preference, results similar to those observed for the paths with structural breaks. Once again, the DCC-NL shrinkage method emerges as the top performer, producing the lowest amount of errors. It is followed by the 2-factor copula, which exhibits slightly higher errors, while the three static models lag behind, showing unsurprisingly higher values compared to the dynamic ones, with similar error rates among them. However, a notable difference in the time-varying scenario lies in the relationship between

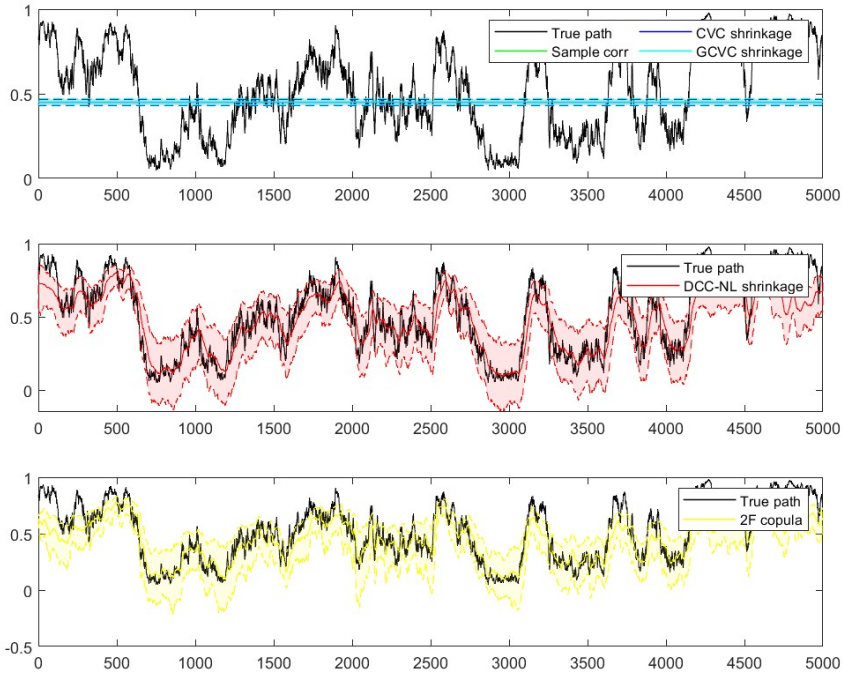

Figure 7: Estimated between-group correlations in the time-varying scenario

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.9819	1.1028	1.1429	0.9991	0.9997	0.9999
CVC shrinkage	0.9827	1.1032	1.1430	1.0000	1.0000	1.0000
GCVC shrinkage	0.9819	1.1028	1.1429	0.9992	0.9997	0.9999
DCC-NL shrinkage	0.6221	0.6571	0.6668	0.6330	0.5956	0.5834
2F Copula	0.6252	0.6582	0.6674	0.6362	0.5967	0.5838
	N = 10					
Sample covariance	0.9746	1.1076	1.1515	0.9984	0.9997	0.9999
CVC shrinkage	0.9762	1.1079	1.1516	1.0000	1.0000	1.0000
GCVC shrinkage	0.9746	1.1076	1.1515	0.9984	0.9997	0.9999
DCC-NL shrinkage	0.6387	0.6715	0.6862	0.6543	0.6061	0.5958
2F Copula	0.6431	0.6731	0.6869	0.6589	0.6075	0.5965
	N = 20					
Sample covariance	0.9918	1.1097	1.1424	0.9984	0.9996	0.9999
CVC shrinkage	0.9934	1.1102	1.1426	1.0000	1.0000	1.0000
GCVC shrinkage	0.9918	1.1097	1.1424	0.9984	0.9996	0.9999
DCC-NL shrinkage	0.6447	0.6849	0.6897	0.6489	0.6170	0.6037
2F Copula	0.6499	0.6864	0.6906	0.6542	0.6183	0.6044

Table 11: Average RMSE of the estimated variances in the time-varying scenario

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.7488	0.8069	0.8140	1.0000	0.9999	1.0002
CVC shrinkage	0.7488	0.8070	0.8138	1.0000	1.0000	1.0000
GCVC shrinkage	0.7488	0.8469	0.8140	1.0000	0.9999	1.0002
DCC-NL shrinkage	0.4312	0.4392	0.4378	0.5758	0.5442	0.5380
2F Copula	0.4326	0.4397	0.4381	0.5777	0.5448	0.5383
	N = 10					
Sample covariance	0.7544	0.8167	0.8315	0.9989	1.0001	1.0001
CVC shrinkage	0.7552	0.8167	0.8315	1.0000	1.0000	1.0000
GCVC shrinkage	0.7544	0.8167	0.8315	0.9989	1.0001	1.0001
DCC-NL shrinkage	0.4492	0.4533	0.4556	0.5947	0.5551	0.5480
2F Copula	0.4512	0.4540	0.4559	0.5974	0.5560	0.5483
	N = 20					
Sample covariance	0.7686	0.8230	0.8319	0.9989	1.0000	0.9999
CVC shrinkage	0.7694	0.8230	0.8320	1.0000	1.0000	1.0000
GCVC shrinkage	0.7686	0.8230	0.8319	0.9989	1.0000	0.9999
DCC-NL shrinkage	0.4521	0.4622	0.4599	0.5876	0.5617	0.5527
2F Copula	0.4546	0.4628	0.4602	0.5908	0.5624	0.5531

Table 12: Average MAE of the estimated variances in the time-varying scenario

errors and sample length. Unlike previous scenarios, here, an increase in sample length leads to higher estimation errors. This finding can be attributed to the increased likelihood of encountering peaks in the true path due to its time-varying nature, which consequently poses estimation challenges for the models. The analysis of errors across varying asset counts reveals different behaviors. The static model exhibits no clear pattern, with errors remaining relatively constant. Conversely, both dynamic models experience a slight increase in errors as dimensionality rises.

Additionally, for the correlations, the average Mean Absolute Percentage Error (MAPE) is computed and the results are reported separately for within and between-group correlations. Tables 13, 14 and 15 report the findings for correlations between assets in the same group. The examination of the corresponding performance measures reflects the findings obtained for estimated variances. The static models exhibit high estimation errors, with comparable values among them. Conversely, dynamic models, particularly the DCC-NL shrinkage method, demonstrate greater accuracy in estimating within-group correlations in the time-varying scenario, with slightly lower estimation errors. However, an increase in the dimensionality leads to a better performance of the 2-factor copula model in terms of all three measures. For this path, it is also explored the impact of a further increase in the number of assets (N = 50) on performance measures. Notably, all three measures continue to exhibit lower errors for the 2-factor copula model compared to the others.

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.2275	0.2534	0.2657	0.9921	0.9975	0.9983
CVC shrinkage	0.2293	0.2540	0.2662	1.0000	1.0000	1.0000
GCVC shrinkage	0.2275	0.2534	0.2657	0.9923	0.9975	0.9983
DCC-NL shrinkage	0.1379	0.1411	0.1418	0.6016	0.5553	0.5328
2F Copula	0.1798	0.1906	0.1925	0.7842	0.7506	0.7232
	N = 10					
Sample covariance	0.2259	0.2512	0.2631	0.9948	0.9979	0.9986
CVC shrinkage	0.2271	0.2517	0.2634	1.0000	1.0000	1.0000
GCVC shrinkage	0.2259	0.2512	0.2631	0.9947	0.9979	0.9986
DCC-NL shrinkage	0.1432	0.1476	0.1499	0.6305	0.5862	0.5691
2F Copula	0.1381	0.1423	0.1455	0.6082	0.5655	0.5524
	N = 20					
Sample covariance	0.2247	0.2534	0.2623	0.9956	0.9979	0.9987
CVC shrinkage	0.2257	0.2540	0.2627	1.0000	1.0000	1.0000
GCVC shrinkage	0.2247	0.2534	0.2623	0.9955	0.9979	0.9987
DCC-NL shrinkage	0.1474	0.1502	0.1521	0.6528	0.5912	0.5790
2F Copula	0.1259	0.1299	0.1329	0.5578	0.5116	0.5059
	N = 50					
Sample covariance	0.2255	0.2528	0.2615	0.9953	0.9980	0.9989
CVC shrinkage	0.2265	0.2533	0.2618	1.0000	1.0000	1.0000
GCVC shrinkage	0.2255	0.2528	0.2615	0.9952	0.9980	0.9989
DCC-NL shrinkage	0.1487	0.1520	0.1543	0.6562	0.6000	0.5894
2F Copula	0.1149	0.1211	0.1236	0.5070	0.4780	0.4719

Table 13: Average RMSE of the estimated within-group correlations in the time-varying scenario

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.1944	0.2166	0.2281	0.9890	0.9959	0.9978
CVC shrinkage	0.1966	0.2175	0.2287	1.0000	1.0000	1.0000
GCVC shrinkage	0.1944	0.2166	0.2281	0.9891	0.9959	0.9978
DCC-NL shrinkage	0.1095	0.1111	0.1111	0.5569	0.5111	0.4857
2F Copula	0.1445	0.1509	0.1509	0.7351	0.6941	0.6599
	N = 10					
Sample covariance	0.1938	0.2156	0.2262	0.9931	0.9973	0.9985
CVC shrinkage	0.1952	0.2162	0.2265	1.0000	1.0000	1.0000
GCVC shrinkage	0.1938	0.2156	0.2262	0.9931	0.9973	0.9985
DCC-NL shrinkage	0.1135	0.1156	0.1170	0.5817	0.5350	0.5167
2F Copula	0.1092	0.1106	0.1125	0.5597	0.5118	0.4965
	N = 20					
Sample covariance	0.1927	0.2182	0.2258	0.9936	0.9975	0.9987
CVC shrinkage	0.1939	0.2188	0.2261	1.0000	1.0000	1.0000
GCVC shrinkage	0.1927	0.2182	0.2258	0.9936	0.9975	0.9987
DCC-NL shrinkage	0.1170	0.1177	0.1186	0.6035	0.5382	0.5243
2F Copula	0.0990	0.1002	0.1018	0.5103	0.4582	0.4504
	N = 50					
Sample covariance	0.1946	0.2181	0.2252	0.9935	0.9975	0.9989
CVC shrinkage	0.1959	0.2186	0.2255	1.0000	1.0000	1.0000
GCVC shrinkage	0.1946	0.2181	0.2252	0.9935	0.9975	0.9989
DCC-NL shrinkage	0.1180	0.1187	0.1203	0.6023	0.5430	0.5334
2F Copula	0.0891	0.0920	0.0939	0.4551	0.4211	0.4162

Table 14: Average MAE of the estimated within-group correlations in the time-varying scenario

Average MAPE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.8652	1.0289	1.1029	0.9858	0.9952	0.9998
CVC shrinkage	0.8777	1.0339	1.1031	1.0000	1.0000	1.0000
GCVC shrinkage	0.8650	1.0289	1.1029	0.9855	0.9952	0.9998
DCC-NL shrinkage	0.4941	0.5003	0.5095	0.5630	0.4839	0.4618
2F Copula	0.8716	0.8983	0.8871	0.9931	0.8689	0.8041
	N = 10					
Sample covariance	0.9556	1.0480	1.1698	0.9968	1.0023	1.0031
CVC shrinkage	0.9586	1.0456	1.1662	1.0000	1.0000	1.0000
GCVC shrinkage	0.9554	1.0480	1.1697	0.9966	1.0022	1.0031
DCC-NL shrinkage	0.5957	0.5665	0.6009	0.6214	0.5418	0.5152
2F Copula	0.6640	0.6196	0.6647	0.6926	0.5925	0.5700
	N = 20					
Sample covariance	0.9592	1.1300	1.1832	0.9982	1.0038	1.0046
CVC shrinkage	0.9610	1.1257	1.1778	1.0000	1.0000	1.0000
GCVC shrinkage	0.9590	1.1300	1.1832	0.9980	1.0038	1.0046
DCC-NL shrinkage	0.6292	0.6305	0.6302	0.6547	0.5601	0.5351
2F Copula	0.5880	0.5928	0.5902	0.6119	0.5266	0.5011
	N = 50					
Sample covariance	1.0371	1.1057	1.1832	1.0033	1.0056	1.0047
CVC shrinkage	1.0336	1.0996	1.1776	1.0000	1.0000	1.0000
GCVC shrinkage	1.0369	1.1056	1.1832	1.0018	1.0055	1.0047
DCC-NL shrinkage	0.6778	0.6374	0.6524	0.6558	0.5797	0.5540
2F Copula	0.5581	0.5398	0.5577	0.5400	0.4909	0.4736

Table 15: Average MAPE of the estimated within-group correlations in the time-varying scenario

This finding suggests that this model is more accurate than others in estimating within-group correlations when the number of assets included is higher, likely due to the additional information provided about correlations between assets. Moreover, similarly to variances, the trend regarding sample length persists. An increase in the number of observations leads to higher estimation errors, meaning less accuracy of all models in estimating the true paths.

The same analysis is repeated also for the correlations between assets in different groups and it is shown in Tables 16, 17 and 18.

Average RMSE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.2330	0.2586	0.2721	0.9985	0.9998	1.0005
CVC shrinkage	0.2334	0.2586	0.2720	1.0000	1.0000	1.0000
GCVC shrinkage	0.2319	0.2581	0.2719	0.9935	0.9979	0.9996
DCC-NL shrinkage	0.1587	0.1623	0.1632	0.6798	0.6275	0.5999
2F Copula	0.1777	0.1842	0.1866	0.7613	0.7123	0.6859
	N = 10					
Sample covariance	0.2255	0.2747	0.2822	1.0000	1.0018	1.0018
CVC shrinkage	0.2254	0.2743	0.2817	1.0000	1.0000	1.0000
GCVC shrinkage	0.2543	0.2743	0.2820	0.9957	1.0002	1.0010
DCC-NL shrinkage	0.1994	0.2023	0.2051	0.7807	0.7378	0.7279
2F Copula	0.2075	0.2114	0.2148	0.8122	0.7709	0.7625
	N = 20					
Sample covariance	0.2616	0.2805	0.2909	1.0024	1.0031	1.0030
CVC shrinkage	0.2609	0.2797	0.2900	1.0000	1.0000	1.0000
GCVC shrinkage	0.2605	0.2801	0.2907	0.9982	1.0016	1.0022
DCC-NL shrinkage	0.2145	0.2184	0.2214	0.8221	0.7810	0.7634
2F Copula	0.2178	0.2227	0.2269	0.8348	0.7962	0.7822
	N = 50					
Sample covariance	0.2666	0.2898	0.2951	1.0031	1.0041	1.0035
CVC shrinkage	0.2658	0.2886	0.2941	1.0000	1.0000	1.0000
GCVC shrinkage	0.2655	0.2894	0.2949	0.9990	1.0027	1.0028
DCC-NL shrinkage	0.2225	0.2299	0.2283	0.8372	0.7965	0.7762
2F Copula	0.2320	0.2456	0.2472	0.8729	0.8507	0.8407

Table 16: Average RMSE of the estimated between-group correlations in the time-varying scenario

The study of these performance measures reveals again that dynamic models consistently outperform static models in terms of error. Specifically, the DCC-NL shrinkage estimator achieves again the lowest errors across all three metrics. An upward trend in errors is observed with an increasing number of observations, mirroring the findings of within-group correlations. However, a key difference emerges when analyzing dimensionality.

Average MAE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.1984	0.2201	0.2323	0.9868	0.9991	1.0000
CVC shrinkage	0.1990	0.2203	0.2323	1.0000	1.0000	1.0000
GCVC shrinkage	0.1977	0.2198	0.2321	0.9934	0.9977	0.9994
DCC-NL shrinkage	0.1289	0.1306	0.1308	0.6478	0.5929	0.5632
2F Copula	0.1442	0.1476	0.1486	0.7245	0.6700	0.6599
	N = 10					
Sample covariance	0.2175	0.2327	0.2386	0.9987	1.0009	1.0011
CVC shrinkage	0.2178	0.2325	0.2384	1.0000	1.0000	1.0000
GCVC shrinkage	0.2168	0.2324	0.2385	0.9954	0.9997	1.0006
DCC-NL shrinkage	0.1624	0.1619	0.1630	0.7459	0.6965	0.6839
2F Copula	0.1679	0.1677	0.1692	0.7708	0.7212	0.7099
	N = 20					
Sample covariance	0.2221	0.2364	0.2445	1.0010	1.0021	1.0022
CVC shrinkage	0.2219	0.2359	0.2439	1.0000	1.0000	1.0000
GCVC shrinkage	0.2213	0.2361	0.2443	0.9976	1.0009	1.0017
DCC-NL shrinkage	0.1750	0.1749	0.1757	0.7890	0.7413	0.7204
2F Copula	0.1762	0.1771	0.1791	0.7940	0.7507	0.7341
	N = 50					
Sample covariance	0.2260	0.2444	0.2470	1.0017	1.0032	1.0027
CVC shrinkage	0.2257	0.2436	0.2463	1.0000	1.0000	1.0000
GCVC shrinkage	0.2253	0.2441	0.2469	0.9983	1.0020	1.0022
DCC-NL shrinkage	0.1819	0.1842	0.1815	0.8061	0.7561	0.7367
2F Copula	0.1884	0.1959	0.1953	0.8347	0.8042	0.7928

Table 17: Average MAE of the estimated between-group correlations in the time-varying scenario

Average MAPE	Absolute value			Relative value		
Observations	1000	2500	5000	1000	2500	5000
	N = 4					
Sample covariance	0.8730	0.9744	1.0477	0.9925	0.9971	0.9978
CVC shrinkage	0.8796	0.9772	1.0500	1.0000	1.0000	1.0000
GCVC shrinkage	0.8863	0.9787	1.0498	1.0076	1.0015	0.9998
DCC-NL shrinkage	0.5422	0.5121	0.5037	0.6164	0.5240	0.4798
2F Copula	0.5469	0.5298	0.5309	0.6218	0.5422	0.5056
	N = 10					
Sample covariance	0.8522	0.9360	0.9924	0.9913	0.9952	0.9962
CVC shrinkage	0.8597	0.9406	0.9961	1.0000	1.0000	1.0000
GCVC shrinkage	0.8589	0.9385	0.9936	0.9991	0.9978	0.9975
DCC-NL shrinkage	0.5630	0.5441	0.5559	0.6549	0.5784	0.5581
2F Copula	0.5456	0.5326	0.5431	0.6346	0.5662	0.5452
	N = 20					
Sample covariance	0.7833	0.9274	0.9764	0.9906	0.9938	0.9957
CVC shrinkage	0.7907	0.9332	0.9806	1.0000	1.0000	1.0000
GCVC shrinkage	0.7883	0.9295	0.9774	0.9970	0.9960	0.9968
DCC-NL shrinkage	0.5620	0.5764	0.5759	0.7107	0.6176	0.5873
2F Copula	0.5333	0.5690	0.5864	0.6745	0.6097	0.5981
	N = 50					
Sample covariance	0.7878	0.8970	1.0029	0.9904	0.9940	0.9954
CVC shrinkage	0.7954	0.9024	1.0076	1.0000	1.0000	1.0000
GCVC shrinkage	0.7922	0.8987	1.0039	0.9959	0.9959	0.9964
DCC-NL shrinkage	0.5693	0.5754	0.6116	0.7157	0.6376	0.6070
2F Copula	0.5572	0.6006	0.6469	0.7005	0.6655	0.6420

Table 18: Average MAPE of the estimated between-group correlations in the time-varying scenario

Here, errors measured by RMSE and MAE increase, while MAPE shows a decreasing trend for all models except DCC-NL, which lacks a clear pattern. The previously observed advantage of the 2-factor copula for within-group correlations with increasing assets can be partially extended to between-group correlations. However, this benefit is only evident in MAPE for the smaller sample lengths. For the other two measures, the error estimations of the two dynamic models converge, with DCC-NL still maintaining a slight edge. Even with a further rise in the number of assets (N = 50), the 2-factor copula model struggles to outperform the DCC-NL shrinkage approach in terms of average RMSE and MAE. However, for the average MAPE, the 2-factor copula still exhibits lower errors at only lower sample lengths.

In conclusion, the findings confirm that dynamic models excel in time-varying environments. Both dynamic models exhibit similar accuracy in capturing variance due to their comparable underlying structure. However, for correlation estimation in high-dimensional settings, the 2-factor copula model provides superior precision for within-group correlations, while DCC-NL outperformed in estimating between-group correlations. Investigating the effect of further increases in dimensionality on model performance accuracy would be an interesting area for future research.

3.4 Computational Time

Selecting the optimal estimation method requires a trade-off, balancing the desired level of accuracy with the processing time required to estimate both variance and correlation. Understanding the computational costs associated with each method is crucial for making an informed decision. Table 19 shows each model's average time to generate estimates for various scenarios. These scenarios explore the impact of a changing number of assets and observations on processing time. This information empowers users to make a data-driven choice, selecting the model that best aligns with their specific needs of speed and precision. The analysis reveals that static models reign supreme in terms of estimation speed. Specifically, the sample covariance approach emerges as the quickest method, with computation time increasing minimally even with a growing number of assets and observations. The DCC-NL shrinkage model follows closely behind, exhibiting processing times ranging from a mere second to a maximum of 3.5 seconds in the extreme case. Instead, the 2-factor copula model presents a more significant trade-off. While it offers superior accuracy in some scenarios, this benefit comes at a substantial computational cost. Estimated times can range from a starting point of around 17 seconds up to 4 minutes in this analysis. Therefore, careful consideration must be given to weighing the 2-factor copula's increased accuracy against its computational demands. It is vital to ensure that the chosen method aligns seamlessly with the project's time constraints and available resources.

Computational time	Absolute value		
Observations	1000	2500	5000
	N = 4		
Sample covariance	0.0002	0.0004	0.0003
CVC shrinkage	0.0002	0.0004	0.0004
GCVC shrinkage	0.0003	0.0006	0.0005
DCC-NL shrinkage	0.3700	0.8900	1.0600
2F Copula	16.955	81.191	84.286
	N = 10		
Sample covariance	0.0002	0.0004	0.0018
CVC shrinkage	0.0004	0.0005	0.0028
GCVC shrinkage	0.0005	0.0006	0.0027
DCC-NL shrinkage	1.0000	1.0500	1.7900
2F Copula	19.325	64.687	106.71
	N = 20		
Sample covariance	0.0011	0.0006	0.0045
CVC shrinkage	0.0025	0.0010	0.0050
GCVC shrinkage	0.0027	0.0010	0.0058
DCC-NL shrinkage	1.6378	1.7368	3.5520
2F Copula	27.835	82.560	241.64

Table 19: Average computational time of the estimation models in seconds

4 Empirical Application

4.1 Data Collection

The discussed models are now applied to real-world financial data. Specifically, this research investigates the daily returns of stocks listed in the S&P500 index over a period of 20 years, from January 2, 2004, to December 29, 2023. I use the Standard Industrial Classification (SIC) code system to categorize the stocks into groups, which is necessary for two of the models considered. The first two digits of the SIC code, designating the major industry sector, allow the categorization of the stocks into groups according to their sector. I further refine the analysis by employing the third digit of the SIC code to sub-categorize stocks within each industry group. This enables a more granular examination of return dynamics at the sub-industry level. For each of these categories, I construct an equally weighted portfolio, which allows me to mitigate the idiosyncratic influence of individual stocks and gain a clearer picture of overall sub-industry return behavior. Finally, I focus on the portfolios that had consistent representation in the S&P500 index throughout the entire 20-year time frame, excluding the ones with sporadic appearances.

Following the data adjustments, I have a breakdown of 8 distinct major industry groups. These groups exhibit heterogeneity in the number of sub-industry portfolios they encompass. Notably, the total number of sub-industry portfolios is 42, with an average of 5.25 sub-industries per group. This average, however, masks the underlying variation, with a minimum of 2 sub-industries and a maximum of 9 within individual groups. The next subsections analyze the estimated variances and correlations for these portfolios.

4.2 Variance Estimation

The work proceeds with the estimation of the daily variances of sub-industry portfolios using the different models introduced in Section 2. A representative example of the estimated daily variances is depicted in Figure 8. The graph reveals that the two dynamic models, DCC-NL shrinkage and 2-factor copula, produce similar daily variance estimates. This is due to their inherent and similar design. Their paths indicate a relatively constant daily volatility at around 1%, with three distinct periods of significantly higher volatility. These three peaks correspond to the three major crises of the past two decades: the 2008 global financial crisis, the 2012 European sovereign debt crisis, and the 2020 COVID-19 market crash. The dynamic models effectively show the impact of these crises, with the first and last crisis causing daily variance to surge twenty to thirty times higher than in a low-volatility regime. Meanwhile, the 2012 crisis results in an increase in volatility of approximately 10%.

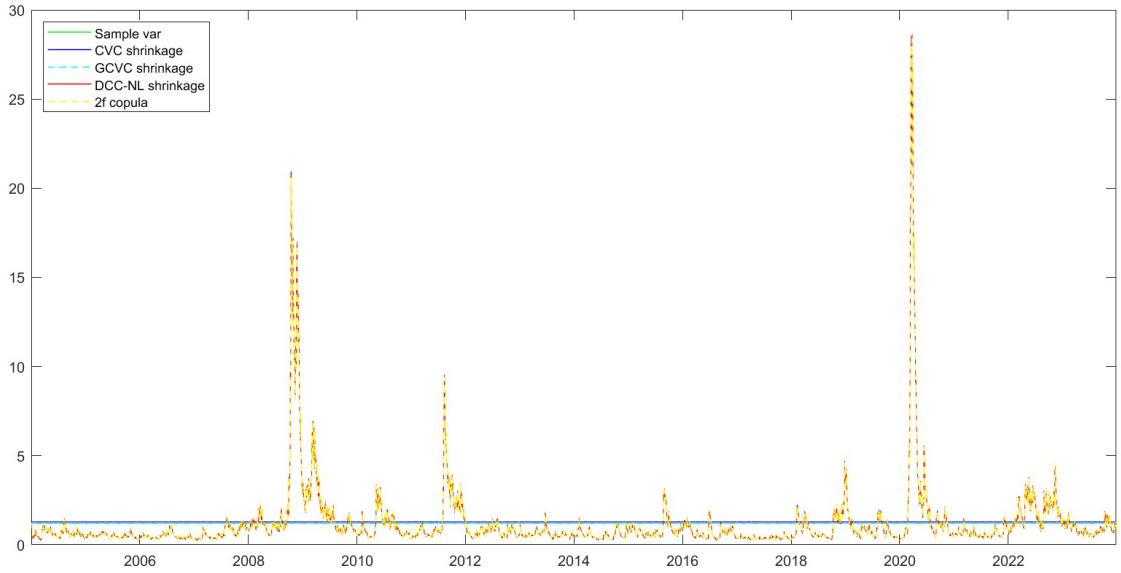

Figure 8: Daily variance estimates for a sub-industry portfolio from 2004 to 2023

The three static estimation methods (sample covariance, CVC and GCVC shrinkage) tend to exhibit a pattern of overestimating variance during normal periods and vastly underestimating variance during crisis periods. This is attributed to their static framework, which fails to fully account for the impact of crises. However, outside of these periods, they generally provide a satisfactory approximation of volatility estimation. In contrast, both dynamic models offer accurate estimations of asset return volatility. They effectively capture changes in values due to exogenous shocks, providing a more responsive and adaptable framework compared to static methods. Consequently, while dynamic models excel in capturing volatility dynamics, static methods remain valuable estimators in stable economic conditions.

In conclusion, dynamic models are adept at accurately estimating asset return volatility, particularly in response to exogenous shocks, while static methods offer reliable estimations in normal economic conditions, but may falter during periods of crisis.

4.3 Correlation Estimation

The estimation analysis progresses by examining correlations among the sub-industry portfolios. This entails assessing correlations within the same sector as well as those across different sectors. Figure 9 depicts the graphical representation of the estimated correlations between two representative portfolios in the same industry group. The analysis reveals that the three static models consistently produce comparable estimation values, with the CVC shrinkage approach estimating values slightly lower than the other two models. Similarly, the two dynamic models also demonstrate comparable values over

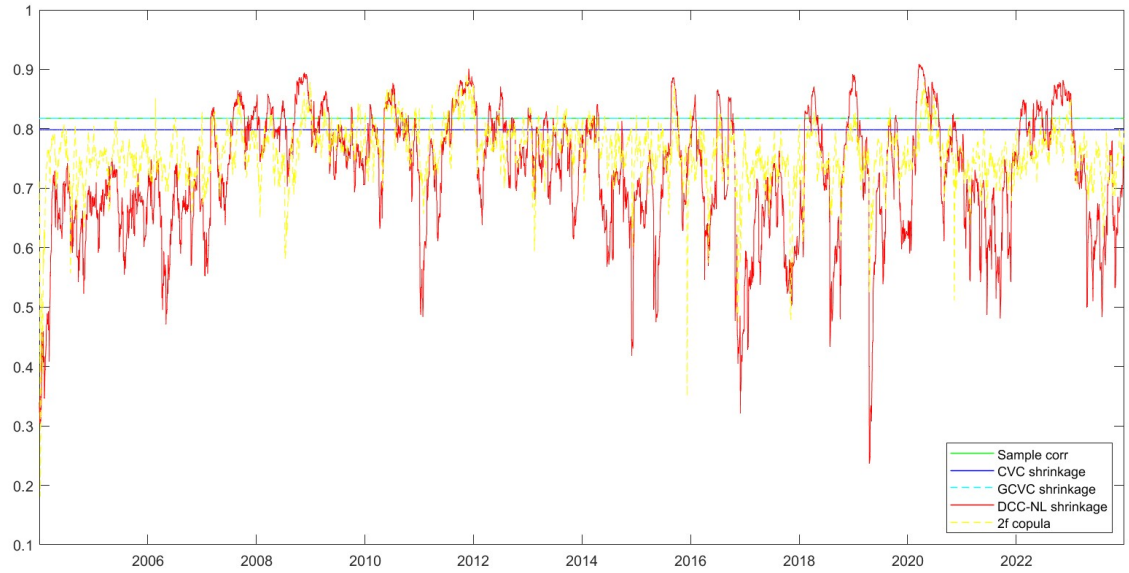

Figure 9: Daily within-group correlation estimates for a sub-industry portfolio from 2004 to 2023

periods among them. However, in the initial periods analyzed, the 2-factor copula model tends to estimate slightly higher values compared to the DCC-NL shrinkage approach. Furthermore, while both models show increased correlations after the occurrence of the 2008 financial crisis, the 2-factor copula model exhibits a slower adjustment to these changes. It also estimates larger decreases in correlation during periods of decline, suggesting a stronger response to significant market downturns. Notably, the correlation peaks at approximately 0.9 during crises, reflecting a trend observed also in the other major financial downturns of the past decade.

Following the 2008 financial crisis, both models exhibit slightly higher correlation estimates compared to pre-crisis levels. However, after 2015, these estimates became more variable, particularly for the DCC-NL model. In contrast, the 2-factor copula model maintained more persistent correlation estimates, closer to pre-crisis values. It is important to note that while the copula model captured changes in correlation, the magnitude of these changes was smaller compared to the DCC-NL model. Section 3 findings offer invaluable insights into these models. Monte Carlo simulations underscored wide confidence bands for both dynamic models in periods with low correlations, allowing me to conclude that despite the dynamic environment of the financial markets, static models remain robust estimation tools for correlation given their closeness to the dynamic model estimations.

Figure 10 visually represents the daily estimated correlation between sub-industry portfolios spanning diverse industry sectors (groups). The analysis of these estimated correlations yields results similar to those obtained for within-group correlations.

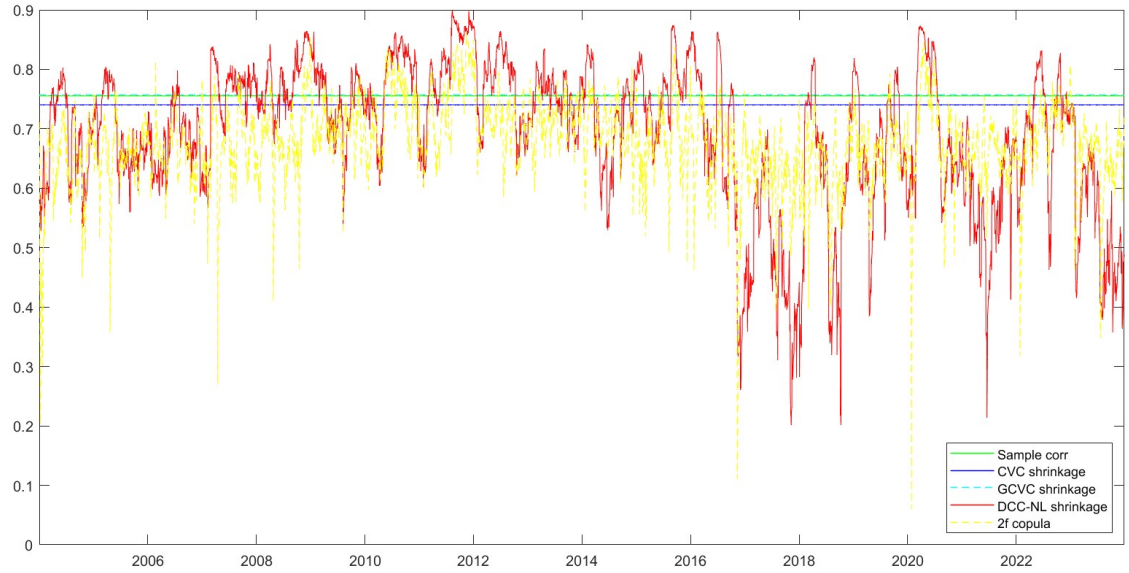

Figure 10: Daily between-group correlation estimates for a sub-industry portfolio from 2004 to 2023

However, an important difference lies in the initial estimation of correlations by the two dynamic models. While they converge on similar values in the beginning, the evolution of correlations over time reveals significant discrepancies. The 2-factor copula exhibits downtrends in correlation values, which suggests an increasing divergence between certain stock movements. This trend is not as strongly captured by the DCC-NL model, which remains more persistent in that period. As expected, the 2008 financial crisis, a period of heightened market stress, led to an increase in correlation values for both models. However, it is worth noting that the 2-factor copula model estimated slightly lower correlation values compared to the DCC-NL model during this time. This finding, along with the simulation results, points towards a potential overestimation of correlations by the DCC-NL model.

Significant differences become evident from 2017 onward, where both dynamic models estimate a decrease in values and an increase in the variability of the correlations. The DCC-NL shrinkage approach detects more pronounced changes compared to the 2-factor copula model, which generally estimates low but more stable values during this period. These observations, along with results from the simulation analysis, suggest that the DCC-NL model may place more importance on changes in values compared to the other dynamic model.

In contrast, the three static models fall back into a pattern of estimating similar values. The CVC shrinkage approach, however, tends to produce slightly lower correlation estimates compared to the other two methods. This similarity suggests a lack of adaptability

to changing market dynamics. Particularly noteworthy is how the static models, due to their inherent stability, mirror the average correlation values estimated by the dynamic models up until 2017. This is because the pre-2017 periods exhibited less variability in correlations. However, this mirroring becomes a weakness after that year. While the dynamic models capture the downward trend in correlations and the increased variability, the static models completely miss these crucial shifts. They remain stuck in their past estimates, failing to reflect the evolving market landscape. This comparison highlights the key advantage of dynamic models. Their ability to adapt to changing market conditions allows them to capture the nuanced fluctuations in correlations, especially during periods of increased volatility. Static models, while offering a simpler approach, may provide misleading information when market dynamics shift.

The comparison between within-group and between-group correlations reveals distinct patterns. Within-group correlations exhibit greater stability throughout the entire period, while between-group correlations show more variation, particularly after 2017. Interestingly, the two correlations diverge significantly following the initial high values due to the first two crises. The between-group correlation drops substantially to around 0.6 (as evident from the 2-factor copula), while the within-group correlation remains relatively stable at around 0.8. This trend is further supported by the static models where, as expected, the within-group correlation is consistently higher than the between-group correlation.

5 Discussion

The selection of the appropriate model for covariance matrix estimation presents a multi-faceted challenge for practitioners. While diverse models exist, each inherently embodies a trade-off between estimation accuracy and computational efficiency. This necessitates careful consideration of user needs and application context. This study explores this interplay, emphasizing how the optimal model selection depends on the specific user needs.

Entities that prioritize long-term stability, such as pension funds, may exhibit a tolerance for minor year-to-year fluctuations in covariance estimates. Here, static models, characterized by their computational efficiency and constant values, are well-suited for such applications. This property arises from their ability to provide a reasonable level of accuracy while maintaining rapid computational speed. In the context of pension fund management, where it is crucial to use a metric that is valid in most cases rather than seeking to outperform the market, static models offer a pragmatic solution for effective risk management strategies. Conversely, institutions like investment banks and hedge funds frequently prioritize high-precision tools for risk assessment. Dynamic models, despite their computational intensity, offer a superior level of accuracy in capturing the evolving nature of covariances. This characteristic makes them the preferred choice for these institutions, even at the cost of increased processing time.

However, the selection process within the realm of dynamic models becomes even more nuanced. The operational time horizon plays a crucial role in this context. For intraday trading activities, computational efficiency becomes crucial. The DCC-NL model, while computationally expensive in complex scenarios, might be a viable option due to its relative speed compared to the other dynamic model. In such high-frequency trading environments, the trade-off leans towards prioritizing speed over the most nuanced level of accuracy, as market movements can occur rapidly and require rapid decision-making capabilities. In contrast, for long-term investment strategies with an even larger number of asset classes, the 2-factor copula model emerges as a compelling choice. Its ability to handle a broader range of assets outweighs its potentially higher computational demands in these scenarios. The 2-factor copula model's ability to account for the complex relationships between multiple asset classes justifies the use of a more powerful model, even if it comes at the cost of increased computational resources.

In conclusion, the selection of the optimal model for covariance matrix estimation transcends a one-size-fits-all approach. Context remains essential, as the most suitable model depends on the specific user, their operational time frame, and the asset classes involved. By understanding these factors and the inherent trade-off between accuracy, computational efficiency, and user needs, practitioners can make informed decisions to effectively select their optimal covariance matrix estimation model.

6 Conclusion

This study investigates various methodologies for estimating covariance matrices in the context of financial markets, with a specific focus on their ability to capture the dynamic nature of asset returns. While the traditional sample covariance matrix serves as a foundational tool, its limitations in capturing the intricacies of market behavior are well-documented. As explored in this research, linear shrinkage techniques offer improvements over the sample covariance matrix, but their impact on estimation accuracy remains relatively modest.

The research examines more sophisticated approaches, specifically the DCC-NL shrinkage and the 2-factor copula models. These models offer significant advancements by effectively capturing the time-varying nature of variance and correlations between assets. Through a comprehensive analysis that combines simulations and empirical data from the S&P500, the study demonstrates the superior performance of these dynamic models, particularly during periods of heightened market volatility. The dynamic models exhibit a clear advantage in their ability to respond to external shocks and capture evolving market correlations, a capability lacking in static models.

These findings highlight the critical importance of employing dynamic models for financial analysis, particularly for institutions that require high precision in risk assessment and asset management. While dynamic models provide a more accurate and responsive framework for estimating covariance matrices, they come with increased computational demands. However, static models offer a compelling alternative in situations where computational efficiency is essential. While less accurate, they are significantly faster to compute, making them suitable for specific use cases. The choice of model ultimately rests with the user, who must carefully consider the trade-off between accuracy, computational resources, and the specific demands of its situation.

This study offers a clear understanding of the strengths and limitations of different covariance matrix estimation models. Moreover, further research in several key areas can enhance the understanding of these models. Firstly, an examination of how a broader range of assets impacts model accuracy could provide interesting insights into the behavior of these models in an even higher dimensional setting. Secondly, examining additional asset classes or diversified portfolios merits consideration to potentially extend the results beyond the stock market. Finally, the incorporation of more sophisticated models has the potential to yield a more nuanced understanding of variance and correlation in the financial markets. This would ultimately facilitate better decision-making capabilities for users navigating the financial landscape.

References

Creal, D., S. Koopman, and A. Lucas (2013). "Generalized Autoregressive Score Models With Applications". In: *Journal of Applied Econometrics* 28, pp. 777–795.

Creal, D. and R. Tsay (2015). "High Dimensional Dynamic Stochastic Copula Models". In: *Journal of Econometrics* 189, pp. 335–345.

De Nard, G. (2019). "Oops! I Shrunk the Sample Covariance Matrix Again: Blockbuster Meets Shrinkage". In: *Journal of Financial Econometrics* 20, pp. 569–611.

Engle, R.F. (2002). "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models". In: *Journal of Business & Economic Statistics* 20, pp. 339–350.

Engle, R.F., O. Ledoit, and M. Wolf (2019). "Large Dynamic Covariance Matrices". In: *Journal of Business & Economic Statistics* 37, pp. 363–375.

Jobson, J. D. and B.M. Korkie (1980). "Estimation for Markowitz Efficient Portfolios". In: *Journal of the American Statistical Association* 75, pp. 544–554.

Ledoit, O. and S. Péché (2011). "Eigenvectors of Some Large Sample Covariance Matrix Ensembles". In: *Probability Theory and Related Fields* 150, pp. 233–264.

Ledoit, O. and M. Wolf (2003). "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection". In: *Journal of Empirical Finance* 10, pp. 603–621.

— (2004a). "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices". In: *Journal of Multivariate Analysis* 88. b, pp. 365–411.

— (2004b). "Honey, I Shrunk the Sample Covariance Matrix". In: *The Journal of Portfolio Management* 30. a, pp. 110–119.

— (2012). "Nonlinear Shrinkage Estimation of Large-Dimensional Covariance Matrices". In: *Annals of Statistics* 40, pp. 1024–1060.

Lucas, A., B. Schwaab, and X. Zhang (2017). "Modeling Financial Sector Joint Tail Risk in the Euro Area". In: *Journal of Applied Econometrics* 32, pp. 171–191.

Marcenko, V. A. and L. A. Pastur (1967). "Distribution of Eigenvalues for Some Sets of Random Matrices". In: *Sbornik: Mathematics* 1, pp. 457–483.

Markowitz, H. (1952). "Portfolio Selection". In: *The Journal of Finance* 7, pp. 77–91.

Michaud, R. (1989). "The Markowitz Optimization Enigma: Is Optimized Optimal?" In: *Financial Analysts Journal* 45, pp. 31–42.

Oh, D. and A. Patton (2017). "Modeling Dependence in High Dimensions With Factor Copulas". In: *Journal of Business and Economic Statistics* 35, pp. 139–154.

— (2018). "Time-Varying Systemic Risk: Evidence From a Dynamic Copula Model of CDS Spreads". In: *Journal of Business and Economic Statistics* 36, pp. 181–195.

Opschoor, A. et al. (2021). "Closed-Form Multi-Factor Copula Models With Observation-Driven Dynamic Factor Loadings". In: *Journal of Business and Economic Statistics* 39.4, pp. 1066–1079.

Pakel, C. et al. (2021). "Fitting Vast Dimensional Time-Varying Covariance Models". In: *Journal of Business & Economic Statistics* 39.3, pp. 652–668.

Sharpe, W.F. (1963). "A Simplified Model for Portfolio Analysis". In: *Management Science* 9, pp. 277–293.