DSSR

Discussion Paper No. 152

Estimation of High-Dimensional Volatility Matrices with Dynamic Conditional Correlation-embedded Mixed Factor Structures

Runyu Dai and Yasumasa Matsuda

May 17, 2026

Data Science and Service Research
Discussion Paper

Center for Data Science and Service Research
Graduate School of Economic and Management
Tohoku University
27-1 Kawauchi, Aobaku
Sendai 980-8576, JAPAN

Estimation of High-Dimensional Volatility Matrices with Dynamic Conditional Correlation-embedded Mixed Factor Structures

Runyu Dai* Yasumasa Matsuda†

May 17, 2026

Abstract

Estimating large volatility matrices is essential to finance research in the era of big data. We propose a unified estimate that embeds a dynamic conditional correlation GARCH (DCC-GARCH) structure into a mixed factor model comprising both observable and weak latent factors, with the residual covariance estimated through an adaptively thresholded sparse matrix based on the extended Principal Orthogonal Complement Thresholding (ePOET) framework. The resulting method, termed DCC-ePOET, jointly captures pervasive signals from both types of factors and dynamic idiosyncratic co-volatilities. It resolves the singularity issue that arises in high-dimensional settings where the cross-sectional dimension N exceeds the serial dimension T, while remaining computationally feasible. Monte Carlo simulations confirm the good finite sample performance of DCC-ePOET across various dimensions. An out-of-sample minimum variance portfolio analysis using S&P 500 data demonstrates the usefulness of DCC-ePOET in practice.

Keywords: Volatility matrix; multivariate GARCH; sparsity-induced weak factor models; thresholding; high-dimensional data; portfolio allocation.

JEL Classification: C32, C55, C58, G11, G17.

*Corresponding author. Tohoku University, Graduate School of Economics and Management, 27-1 Kawauchi, Aoba-Ku, Sendai-city, Miyagi, 980-8576, JAPAN; Email: dai.runyu.r1@alumni.tohoku.ac.jp.
The reserach was supported by the Japan Society for the Promotion of Science (Grant-in-Aid for Research Activity Start-up 24K22603).

†Tohoku University, Graduate School of Economics and Management, 27-1 Kawauchi, Aoba, Sendai, Miyagi, 980-8576, JAPAN; E-mail: yasumasa.matsuda.a4@tohoku.ac.jp.
The research was supported by the Japan Society for the Promotion of Science (Grant Number 23K21647)

1 Introduction

Estimation of dynamic covariance matrices is an essential problem in finance research with wide applications recently (Bodilsen, 2025; Engle et al., 2019; Fan et al., 2024; Fiszeder et al., 2023; Moura et al., 2020). As large-scale economic and financial data continue to proliferate in this big data era, the number of assets available for investment and the number of firms and markets that should be jointly modeled have all expanded considerably, posing severe challenges to traditional low-dimensional dynamic modeling approaches. Accordingly, there has been a growing interest in estimating various large dynamic covariance matrices (Chen et al., 2019; Dendramis et al., 2021; Guo et al., 2017; Ke et al., 2022; Wang et al., 2021). See more in a recent survey by Li (2024).

The practical significance of covariance matrix estimation is underscored by its diverse applications in empirical finance and accounting research. In portfolio optimization, Lyle and Yohn (2021) utilized the nonlinear shrinkage estimator (Ledoit and Wolf, 2012, 2017) with fundamental analysis, showing that the quality of the covariance matrix estimate has a substantial impact on portfolio performance. In volatility spillover and systemic risk analysis, Kuo and Chiang (2025) employed multivariate GARCH-type models to quantify dynamic linkages across markets and sectors, and Joseph et al. (2020) used VECM-GARCH with time-varying copulas to characterize volatility co-movement among central and eastern European stock markets during crisis periods. In ESG and green finance research, researchers apply DCC-GARCH-type models to characterize dynamic dependence structures among green bonds, ESG indices, conventional equities, and government bonds, providing guidance for ESG-based investment strategies and risk management (Arouri et al., 2025; Wu and Qin, 2024; Zhang et al., 2022). However, when these directions are extended to large-scale dynamic systems with hundreds or thousands of assets, the volatility matrix estimator must capture the time-varying dependencies while remaining statistically robust and computationally tractable. This is the problem addressed in the present paper. In this study, we examine a typical large dynamic covariance matrix, namely the volatility matrix, which is used for modeling the co-volatile structures of many financial time series.

Various approaches are used to model volatility matrices. A comprehensive overview can be found in the textbook by Francq and Zakoian (2019). Among these methods, the multivariate (generalized) autoregressive conditional heteroskedasticity (MGARCH), a multivariate extension of the well-known univariate ARCH (Engle, 1982) and GARCH (Bollerslev, 1986), stands out as a natural choice. However, standard MGARCH models are notoriously affected by the curse of dimensionality, in which the number of parameters increases explosively with the cross-sectional dimension (the number of series) N, leading to extremely high computational demands. To mitigate this issue, for moderate dimensions, Bollerslev (1990) introduced the constant conditional correlation GARCH (CCC-GARCH) model, which assumes that the MGARCH framework comprises N univariate GARCH processes linked by a constant conditional correlation matrix. In the early 2000s, Engle (2002) extended the CCC-GARCH model to the dynamic conditional correlation (DCC)-GARCH model, which allows conditional correlations among series to vary over time. For high-dimensional data, namely, N increases as serial dimension T increases or even $N \geqslant T$, Engle et al. (2019) recently introduced the DCC-nonlinear (NL) model, which employs the NL shrinkage method proposed by Ledoit and Wolf (2012) to estimate the target matrix of DCC. Due to its flexibility and tractability, the DCC framework has become one of the most widely used volatility modeling tools in empirical finance (Allen et al., 2012; Brownlees and Engle, 2017; Chou et al., 2009; Engle and Siriwardane, 2018; Girardi and Ergün, 2013).

From a different perspective, namely dimension reduction, Ding (1994) proposed the Principal Component (PC)-GARCH model, designed to capture co-volatility using several PCs. The PC-GARCH model provides a simple low-rank framework for parameterizing MGARCH models and has been extended into various variants, such as the orthogonal (O)-GARCH and generalized orthogonal (GO)-GARCH models (Alexander, 2000, 2002; Alexander and Chibumba, 1996; Van der Weide, 2002). The PC-GARCH model is considered part of the Factor-GARCH model family, which broadly assumes that data are influenced by a set of common factors, either observable or latent. Vrontos et al. (2003) introduced the Full-factor GARCH model, which assumes that the number of conditionally uncorrelated factors equals N, while the method of Lanne and Saikkonen (2007) allows for a smaller number of conditionally heteroskedastic factors together with idiosyncratic shocks. Hafner and Preminger (2009) investigated the asymptotic properties of the quasi-maximum likelihood estimation (QMLE) for general Factor-GARCH models. In the same year, Zhang and Chan (2009) incorporated a DCC structure into a Factor-GARCH model to better capture the dynamic relationships among factors.

For high-dimensional data, Fan et al. (2008), Fan et al. (2011), and Fan et al. (2013)

consecutively proposed a structure of common factors plus a noise covariance matrix to model the static covariance matrix. The first two approaches, namely the well-known Fama-French three- and five-factor models (FF3 and FF5 hereafter, respectively), assume that the factors are known or observable. The third approach identifies common latent factors using PCs. They further assumed that idiosyncratic error terms were weakly correlated after extracting latent factors, leading to a conditional sparsity structure in the error covariance matrix; the related empirical evidence has been documented in Jurado et al. (2015), for instance. By combining PCs with a sparse error covariance matrix, Fan et al. (2013) developed the Principal Orthogonal Complement Thresholding (POET) framework, which has been widely applied in the literature (Choi and Kim, 2023, 2025; Ding et al., 2021; Fan and Kim, 2018; Wang et al., 2021). Building on the series of contributions by Fan and co-authors (Fan et al., 2008, 2011, 2013) and the Factor-GARCH framework, Li et al. (2022b) proposed the CCC-Diag estimator, which embeds observable factors within a CCC-GARCH structure and models the idiosyncratic component as a diagonal matrix of univariate error volatilities. To further capture cross-sectional correlations in error terms, Li et al. (2022a) replaced the diagonal error volatility matrix in CCC-Diag with a static POET estimate (CCC-POET). However, relying solely on observable factors for covariance matrix estimation raises significant concerns regarding omitted variables; for instance, see the discussion on omitted factors in asset pricing by Giglio and Xiu (2021) and the research on factor zoos by Feng et al. (2020). A natural and intuitive solution is to detect and utilize potential latent factors in the residuals of linear observable factor models. Shi et al. (2022) augmented the Fama-French models of observed factors by further extracting latent factors from residuals to estimate large static covariance matrices. Their approach inherently assumed that the latent factors in the residuals (LFR) are obtained by the PC method as strong factors; specifically, each LFR diverges at the rate of N. However, as pointed out in the extended POET (ePOET) model (Dai et al., 2024), the latent factors in the residuals of the Fama-French models are usually relatively weak, such that each latent factor may diverge more slowly than N. The scree plot of an illustrative example using S&P500 return data (Fig. 1) reveals that the sample eigenvalues of residuals decrease relatively smoothly.[1] This behavior easily violates the strong factor assumption, which requires a large gap between the

[1]This illustrative example utilizes a dataset comprising daily returns of 331 companies listed in the S&P500, along with FF5 data, spanning from April 2, 2002, to March 31, 2003. Notably, the edge method (Onatski, 2010) robustly identifies six latent factors in the residuals of FF5.

eigenvalues.

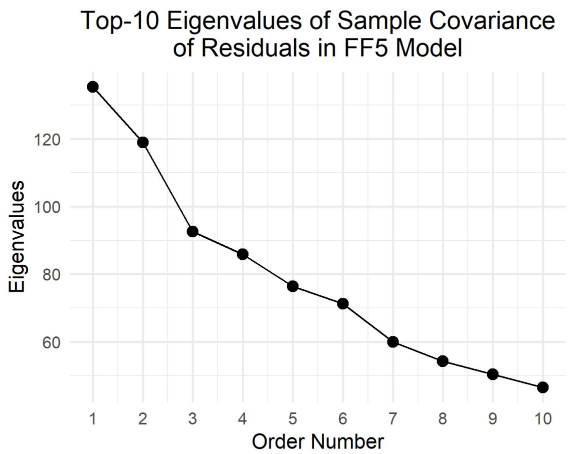

Fig. 1 Eigenvalues plot of sample covariance of FF5 model

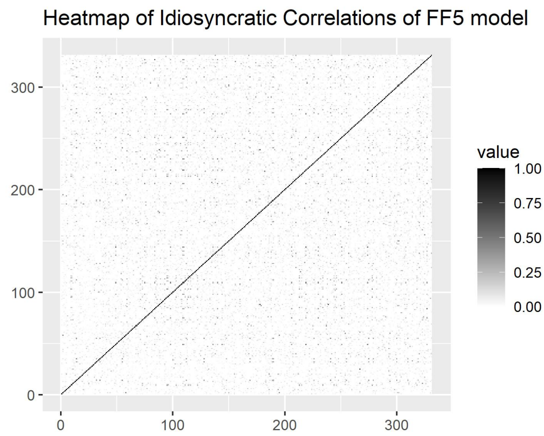

Fig. 2 Heatmap of absolute values of devolatilized idiosyncratic correlations from DCC-ePOET

Building on these discussions, we propose a novel estimate for a large volatility matrix that combines the Factor-GARCH framework and weak latent factors in the residuals of linear observable factor models. Weak latent factors are estimated using the sparsity-induced weak factor (sWF) framework introduced by Uematsu and Yamagata (2023), which employs the Sparse Orthogonal Factor Regression (SOFAR) to achieve sparse factor loadings with diverging rates slower than N. To effectively capture factor volatilities, we assume that both observable and latent factors follow a DCC-GARCH structure, drawing on the inspiration from the DCC-embedded factor framework proposed by Zhang and Chan (2009). Furthermore, we refine the volatility structure of the error terms by integrating a static conditionally sparse matrix into the DCC-GARCH model, based on the assumption that the devolatilized error terms exhibit weak correlations. Fig. 2 illustrates the plot of the example, using the same data as in Fig. 1. The proposed approach extends the original ePOET method from static covariance matrices to volatility matrices by embedding DCC structures into both mixed factors and idiosyncratic errors. This extension, referred to as DCC-ePOET, addresses the issue of information loss of unobservable factors inherent in the approaches of Fan et al. (2011), Li et al. (2022a), and Li et al. (2022b), and overcomes the limitations associated with relying solely on static error covariance, as seen in Li et al. (2022b).

The contributions of this study can be summarized in the following aspects. (i) We propose the DCC-ePOET estimator, which embeds DCC-GARCH structures simultaneously at three levels: observable factors, weak latent factors, and idiosyncratic error terms. Compared with the CCC-Diag estimator of Li et al. (2022b) and the CCC-POET estimator of Li et al. (2022a), our approach addresses both the information loss arising from reliance on observable factors alone and the dynamic information loss entailed by using only a static error covariance. (ii) Computationally, compared with standard MGARCH approaches, DCC-ePOET remains feasible in high-dimensional settings where N is allowed to grow with or even exceed T. We further provide a user-friendly flowchart of the algorithm (Fig. A.1) for practitioners. (iii) Numerically, Monte Carlo simulations confirm the good finite sample performance of DCC-ePOET across various dimensions. In an out-of-sample minimum-variance portfolio analysis using S&P 500 data, DCC-ePOET delivers the lowest risk among nine candidate estimators, providing direct empirical support for its practical usefulness.

1.1 Notations and Organization

Throughout this study, $\mathbf{I}_N$ is an $N \times N$ identity matrix, and $\mathbf{0}$ is a zero matrix. For any square matrix $\mathbf{A}$, we denote the kth largest, the largest, and the smallest eigenvalues by $\lambda_k(\mathbf{A})$, $\lambda_{\max}(\mathbf{A})$, and $\lambda_{\min}(\mathbf{A})$, respectively. $|\mathbf{A}|$ denotes the determinant of $\mathbf{A}$. For any matrix $\mathbf{M} = (m_{t,i}) \in \mathbb{R}^{T\times N}$, we define the l_2 norm (spectral), the Frobenius norm, the entry-wise maximum norm, and the l_1 norm as $\|\mathbf{M}\|_2 = \lambda_{\max}^{1/2}(\mathbf{M}'\mathbf{M})$, $\|\mathbf{M}\|_{\mathrm{F}} = (\sum_{t,i} m_{t,i}^2)^{1/2}$, $\|\mathbf{M}\|_{\max} = \max_{t,i}|m_{t,i}|$, and $\|\mathbf{M}\|_1 = \sum_{t,i}|m_{t,i}|$, respectively. Let $\lesssim$ and $\gtrsim$ represent $\leqslant$ and $\geqslant$ up to a positive constant factor, respectively. For two positive values x and y, let $x \wedge y =: \min\{x, y\}$ and $x \vee y =: \max\{x, y\}$. Finally, for two positive sequences, a_n and b_n, depending on n, we have $a_n \asymp b_n$ if $a_n \gtrsim b_n$ and $a_n \lesssim b_n$. The notation $\mathrm{Cov}(\cdot)$ means the covariance matrix of a random vector. The terms "errors", "noise", and "idiosyncratic terms/errors" in factor models have the same meaning in the paper.

The rest of this paper is organized as follows. Section 2 introduces the DCC-ePOET model. The estimation methodologies and theoretical discussions are presented in Section 3. The Monte Carlo experiments are presented in Section 4. Section 5 compares the DCC-ePOET model with other candidate models in a portfolio analysis. Section 6 concludes the paper. Finally, we provide the algorithm flowchart and a computational speed comparison with several competitors in Online Appendix A.

2 Model descriptions

We estimate the volatility matrix of N-dimensional financial series vector $\mathbf{y}_t$ generated by the linear factor model as follows:

$$\mathbf{y}_t = \mathbf{A}^0\mathbf{x}_t + \mathbf{u}_t, \tag{2.1}$$

where $\mathbf{x}_t$ and $\mathbf{A}^0 = (\mathbf{a}_1^0, \dots, \mathbf{a}_r^0)$ represent the r-dimensional observable factor and its factor loadings, respectively. Furthermore, the error term $\mathbf{u}_t$ has the following latent factor structure:

$$\mathbf{u}_t = \mathbf{B}^0\mathbf{f}_t^0 + \mathbf{e}_t, \tag{2.2}$$

where $\mathbf{f}_t^0$ and $\mathbf{B}^0 = (\mathbf{b}_1^0, \dots, \mathbf{b}_K^0)$ represent the K-dimensional latent (unobservable) factor and its factor loadings, respectively, and $\mathbf{e}_t$ is the idiosyncratic error term. We can also rewrite the models in matrix form: $\mathbf{Y} = \mathbf{X}\mathbf{A}' + \mathbf{U}$ and $\mathbf{U} = \mathbf{F}\mathbf{B}' + \mathbf{E}$. The number of observable factors r and the number of latent factors K are assumed to be finite and small. To separately estimate $\mathbf{A}^0$ and $\mathbf{B}^0$, we further assume that $\mathbf{x}_t$ and the latent parts $\mathbf{u}_t$ are uncorrelated. One may choose a more rigorous model that considers the endogeneity issue between observable and latent factors, such as Bai (2009). Notably, to avoid the omitted variable problem, Model (2.1) is allowed to include many observable factors; however, we do not pursue these directions to avoid technical and computational issues. Throughout the paper, the following conditions are proposed without loss of generality to identify the latent factors and loadings:

$$E[\mathbf{f}_t^0\mathbf{f}_t^{0'}] = \mathbf{I}_K \quad \text{and} \quad \mathbf{B}^{0'}\mathbf{B}^0 \text{ diagonal.} \tag{2.3}$$

For the unobservable weak factors in (2.2), we formally introduce the sWF model following Uematsu and Yamagata (2023). We assume that $\lambda_{\min}(\mathbf{B}^{0'}\mathbf{B}^0)$ is bounded away from zero and $\lambda_{\max}(\mathbf{\Sigma}_e)$ is bounded from infinity. Then for $k = 1, \dots, K$, the latent factor model in (2.2) is called the weak factor model if

$$\lambda_k(\mathbf{B}^{0'}\mathbf{B}^0) = \lambda_k(\mathbf{B}^0\mathbf{B}^{0'}) \asymp N^{\alpha_k}, \quad k = 1, \dots, K \tag{2.4}$$

for some $\alpha_k \in (0, 1]$. The weak factor model is realized by assuming that $\mathbf{B}^0$ is potentially sparse. Specifically, for each $k = 1, \ldots, K$, suppose $\mathbf{b}_k^0$ has $N_k := \lfloor N^{\alpha_k} \rfloor$ nonzero elements for some constant $\alpha_k \in (0, 1]$, and this is so-called sparsity-induced WF (sWF). Notably, the exponent α_k being close to zero means that a factor is extremely weak, while $\alpha_k = 1$ implies a so-called strong factor (Bai and Ng, 2002; Fan et al., 2013). In practice, α is measured by $\hat{\alpha} = \log(\hat{N}_k) / \log(N)$, where $\hat{N}_k$ is the number of nonzero elements in $\hat{\mathbf{b}}_k$.

Under the orthogonal condition between $\mathbf{x}_t$ and $\mathbf{u}_t$, the volatility matrix of $\mathbf{y}_t$ and $\mathbf{\Sigma}_y(t) = \mathrm{Cov}(\mathbf{y}_t | \mathcal{F}_{t-1})$ is then given by

$$\mathbf{\Sigma}_y(t) = \mathbf{A}^0 \mathbf{\Sigma}_x(t) \mathbf{A}^{0'} + \mathbf{B}^0 \mathbf{\Sigma}_f(t) \mathbf{B}^{0'} + \mathbf{\Sigma}_e(t), \tag{2.5}$$

where $\mathbf{\Sigma}_x(t) = \mathrm{Cov}(\mathbf{x}_t | \mathcal{F}_{t-1})$ is the volatility matrix of observable factors $\mathbf{x}_t$, $\mathbf{\Sigma}_f(t) = \mathrm{Cov}(\mathbf{f}_t | \mathcal{F}_{t-1})$ is the volatility matrix of latent factors $\mathbf{f}_t$, and $\mathbf{\Sigma}_e(t) = \mathrm{Cov}(\mathbf{e}_t | \mathcal{F}_{t-1})$ is the volatility matrix of the idiosyncratic errors $\mathbf{e}_t$. To capture the volatility of the observable factors, we assume $\mathbf{x}_t$ has the DCC-GARCH(p, q) structure (Engle, 2002) as follows. For $l \in (1, ..., r)$,

$$\begin{aligned}
\mathbf{\Sigma}_x(t) &= \mathbf{D}_t^x \mathbf{R}_t^x \mathbf{D}_t^x, \\
\mathbf{R}_t^x &= (\mathrm{diag}\, \mathbf{Q}_t^x)^{-1/2} \mathbf{Q}_t^x (\mathrm{diag}\, \mathbf{Q}_t^x)^{-1/2}, \\
\mathbf{Q}_t^x &= (1 - a - b)\mathbf{S}_x + a \boldsymbol{\eta}_{t-1}^x \boldsymbol{\eta}_{t-1}^{x'} + b \mathbf{Q}_{t-1}^x, \\
{\sigma_{t,l}^x}^2 &= \omega_l^x + \sum_{i=1}^q \alpha_{i,l}^x x_{t-i,l}^2 + \sum_{i=1}^p \beta_{i,l}^x {\sigma_{t-i,l}^x}^2,
\end{aligned} \tag{2.6}$$

where $\mathbf{D}_t^x = \mathrm{diag}\left(\sigma_{t,1}^x, ..., \sigma_{t,r}^x\right)$, $\boldsymbol{\eta}_t^x = \left(x_{t,1}/\sigma_{t,1}^x, ..., x_{t,r}/\sigma_{t,r}^x\right)$ is the vector of devolatilized factors, and $\mathbf{S}_x$ is a positive-definite unconditional covariance (target) matrix. The latent factors $\mathbf{f}_t$ are assumed to a DCC-GARCH structure of the same form as in (2.6). If no latent factor exists, the second term on the RHS of (2.5) will disappear. Finally, in a manner similar to the two types of factors, the idiosyncratic error term $\mathbf{e}_t$ also follows the standard DCC-GARCH model as:

$$\begin{aligned}
\mathbf{\Sigma}_e(t) &= \mathbf{D}_t^e \mathbf{R}_t^e \mathbf{D}_t^e, \\
\mathbf{R}_t^e &= (\mathrm{diag}\, \mathbf{Q}_t^e)^{-1/2} \mathbf{Q}_t^e (\mathrm{diag}\, \mathbf{Q}_t^e)^{-1/2}, \\
\mathbf{Q}_t^e &= (1 - a - b)\mathbf{S}_e + a \boldsymbol{\eta}_{t-1}^e \boldsymbol{\eta}_{t-1}^{e'} + b \mathbf{Q}_{t-1}^e, \\
{\sigma_{t,j}^e}^2 &= \omega_j^e + \sum_{i=1}^q \alpha_{i,j}^e e_{t-i,j}^2 + \sum_{i=1}^p \beta_{i,j}^e {\sigma_{t-i,j}^e}^2, \ j = 1, ..., N,
\end{aligned} \tag{2.7}$$

where $\mathbf{D}_t^e = \text{diag}\left(\sigma_{t,1}^e, ..., \sigma_{t,N}^e\right)$, $\boldsymbol{\eta}_t^e = \left(e_{t,1}/\sigma_{t,1}^e, ..., e_{t,N}/\sigma_{t,N}^e\right)$. Notably, after removing the pervasive observable and (weak) latent factors from the target variables, it is reasonable to think that the devolatilized error terms are weakly cross-sectionally-correlated. Following Cai and Liu (2011), Fan et al. (2011), and Fan et al. (2013), we treat the unconditional target covariance matrix of the devolatilized errors, $\mathbf{S}_e$, as conditionally sparse to reflect the weak correlations. That is $\mathbf{S}_e \in \Upsilon(m_N)$, where

$$\Upsilon(m_N) = \left\{ \mathbf{S}_e = (s_{ij}^e)_{N\times N} \succ 0 : \max_{i\leqslant N} \sum_{j=1}^{N} |s_{ij}^e|^q (s_{ii}^e s_{jj}^e)^{(1-q)/2} \leqslant m_N \right\} \tag{2.8}$$

for the sparsity measure $m_N = \max_{i\leqslant N} \sum_{j\leqslant N} |s_{ij}^e|^q$ with some constant $q \in [0,1]$. Intuitively, this condition requires each row of $\mathbf{S}_e$ to contain only a small number of non-negligible entries.

3 Estimation and theoretical discussions

First, we apply ordinary least squares (OLS) to (2.1) to estimate the loadings, $\hat{\mathbf{A}}$, for the observable factors and obtain the initial residuals, $\hat{\mathbf{U}}$. Next, we estimate the latent factors and their loadings from $\hat{\mathbf{U}}$ and calculate the corresponding errors, $\hat{\mathbf{E}}$. Subsequently, we estimate the volatility matrices for the observable factors, the latent factors, and the errors.

3.1 Estimation of the volatility matrix of observable factors

As the observable factors and the number of factors, r, are known, only the r-dimensional DCC-GARCH model is estimated at this stage. The estimation is based on the standard QMLE but is not computationally challenging due to the small r. We follow the usual DCC-GARCH estimation procedure proposed by Engle (2002). The log-likelihood function is given by

$$\begin{aligned}
\mathbf{L}^x &= -\frac{1}{2}\sum_{t=1}^{T} r\log(2\pi) + \log|\boldsymbol{\Sigma}_x(t)| + \mathbf{x}_t'\boldsymbol{\Sigma}_x(t)^{-1}\mathbf{x}_t \\
&= -\frac{1}{2}\sum_{t=1}^{T} r\log(2\pi) + 2\log|\mathbf{D}_t^x| + \mathbf{x}_t'\mathbf{D}_t^{x\,-2}\mathbf{x}_t - \boldsymbol{\eta}_t^{x'}\boldsymbol{\eta}_t^x + \log|\mathbf{R}_t^x| + \boldsymbol{\eta}_t^{x'}\mathbf{R}_t^{x\,-1}\boldsymbol{\eta}_t^x.
\end{aligned}$$

We observe that $\mathbf{L}^x$ can be decomposed into two sub-functions, the univariate volatility one and the conditional correlation one. We denote θ_g^x as the univariate GARCH parameter for observable factors and $\phi_c = (a,b)$ as the correlation parameter. The usual stationarity conditions: $0 < \sum_{i=1}^{q} \alpha_{i,l} + \sum_{i=1}^{p} \beta_{i,l} < 1, \omega_l > 0$ for GARCH(p,q) and $a,b \geqslant 0, a+b<1$ for DCC are satisfied. Then, we have

$$\begin{aligned}
\mathbf{L}^x(\theta_g^x, \phi_c) &= -\frac{1}{2}\sum_{t=1}^{T} r\log(2\pi) + 2\log|\mathbf{D}_t^x| + \mathbf{x}_t'\mathbf{D}_t^{x\,-2}\mathbf{x}_t \\
&\quad -\frac{1}{2}\sum_{t=1}^{T} \boldsymbol{\eta}_t^{x\prime}\boldsymbol{\eta}_t^x + \log|\mathbf{R}_t^x| + \boldsymbol{\eta}_t^{x\prime}\mathbf{R}_t^{x\,-1}\boldsymbol{\eta}_t^x \\
&= \mathbf{L}_g^x(\theta_g^x) + \mathbf{L}_c^x(\theta_g^x, \phi_c). \qquad (3.9)
\end{aligned}$$

The first term on the RHS of (3.9) represents the likelihood of the volatility components and can be expressed as the sum of the univariate GARCH likelihoods:

$$\mathbf{L}_g^x(\theta_g^x) = -\frac{1}{2}\sum_{t=1}^{T}\sum_{l=1}^{r} \log(2\pi) + \log\left(\sigma_{t,l}^{x\;2}\right) + \frac{x_{t,l}^2}{\sigma_{t,l}^{x\;2}}. \qquad (3.10)$$

The second term on the RHS of (3.9) is the likelihood for estimating the conditional correlation matrix. It is obvious that given θ_g^x, the maximum of $\mathbf{L}_c^x$ can be achieved, leading to a two-step QMLE. The final estimator of DCC-GARCH is then obtained by

$$\hat{\theta}_g^x = \arg\min -\mathbf{L}_g(\theta_g^x), \;\; \hat{\phi}_c = \max_{\phi_c} \mathbf{L}_c(\hat{\theta}_g^x, \phi_c). \qquad (3.11)$$

Recalling the DCC-GARCH model, we also need an estimate of the unconditional covariance matrix $\mathbf{S}_x$ for estimating the dynamic correlation $\mathbf{R}_t^x$. We denote the volatility estimator and correlation estimators as $\hat{\sigma}_{t,l}^x$ and $\hat{\mathbf{R}}_t^x$, respectively. Using $\hat{\mathbf{S}}_x = T^{-1}\sum_{t=1}^{T} \hat{\boldsymbol{\eta}}_t^x \hat{\boldsymbol{\eta}}_t^{x\prime}$, where $\hat{\boldsymbol{\eta}}_t^x = \left(x_{t,1}/\hat{\sigma}_{t,1}^x, ..., x_{t,r}/\hat{\sigma}_{t,r}^x\right)$, we can solve $\mathbf{Q}_t^x$ in (2.6) recursively to obtain $\hat{\mathbf{R}}_t^x$. The final estimator of $\boldsymbol{\Sigma}_x(t)$ is given by $\hat{\boldsymbol{\Sigma}}_x(t) = \hat{\mathbf{D}}_t^x \hat{\mathbf{R}}_t^x \hat{\mathbf{D}}_t^x$, where $\hat{\mathbf{D}}_t^x = \mathrm{diag}\left(\hat{\sigma}_{t,1}^x, ..., \hat{\sigma}_{t,r}^x\right)$. Note that in this study, we assume r is finite and small, and T is very large. Thus, compared to the sample covariance matrix in high-dimensional cases, the sample version $\hat{\mathbf{S}}_x$ here is not a concern. If r is allowed to be large or grow comparably to N and T, another issue of estimating a large dimensional covariance matrix will arise.

3.2 Estimation of latent factors and their volatility matrices

To estimate the sWF model, we first obtain the Sparse Orthogonal Factor Regression (SOFAR) estimator (Uematsu et al., 2019; Uematsu and Yamagata, 2023) of $(\mathbf{B}^0, \mathbf{F}^0)$ by

$$(\hat{\mathbf{B}}, \hat{\mathbf{F}}) = \underset{(\tilde{\mathbf{B}},\tilde{\mathbf{F}}) \,\in\, \mathbb{R}^{N\times\hat{K}}\times\mathbb{R}^{T\times\hat{K}}}{\arg\min} \frac{1}{2}\|\hat{\mathbf{U}} - \tilde{\mathbf{F}}\tilde{\mathbf{B}}'\|_{\mathrm{F}}^2 + \eta_{NT}\|\tilde{\mathbf{B}}\|_1 \quad (3.12)$$

$$\text{subject to } \ \tilde{\mathbf{F}}'\tilde{\mathbf{F}}/T = \mathbf{I}_{\hat{K}} \ \text{ and } \ \tilde{\mathbf{B}}'\tilde{\mathbf{B}} \text{ diagonal,}$$

where $\hat{K}$ is the estimated number of factors, and $\eta_{NT} > 0$ is a penalty coefficient. Note that the latent factor components, $\mathbf{C}^0 = \mathbf{F}^0\mathbf{B}^{0\prime}$, can be singular value decomposed (SVD) into $\mathbf{C}^0 = \mathbf{U}^0\mathbf{D}^0\mathbf{V}^{0\prime}$. Here, $\mathbf{D}^0 = \mathrm{diag}(d_1, \ldots, d_K)$, where $(d_1, \ldots, d_K)$ are the scaled singular values of $\mathbf{C}^0$. The matrices $\mathbf{U}^0 \in \mathbb{R}^{T\times K}$ and $\mathbf{V}^0 \in \mathbb{R}^{N\times K}$ represent the scaled left- and right-singular vectors, respectively. These matrices satisfy the identification restrictions $\mathbf{U}^0\mathbf{U}^{0\prime}/T = \mathbf{I}_K$ and $\mathbf{V}^{0\prime}\mathbf{V}^0 = \mathrm{diag}(N_1, \ldots, N_K)$, allowing us to treat $\mathbf{F}^0 = \mathbf{U}^0$ and $\mathbf{B}^0 = \mathbf{V}^0\mathbf{D}^0$. Given $\hat{K}$, the SOFAR method employs the augmented Lagrangian algorithm combined with block coordinate descent to solve the regularized optimization problem (3.12), yielding $(\hat{\mathbf{F}}, \hat{\mathbf{B}}) = (\hat{\mathbf{U}}, \hat{\mathbf{V}}\hat{\mathbf{D}})$.

Using the estimated latent factors $\hat{\mathbf{F}}$, we fit the same DCC-GARCH model as for the observable factors in Section 3.1 to obtain the univariate volatility estimates $(\hat{\sigma}^f_{t,1}, ..., \hat{\sigma}^f_{t,\hat{K}})$ and the conditional correlation matrix $\hat{\mathbf{R}}^f_t$. Then the estimator of $\mathbf{\Sigma}_f(t)$ is given by $\hat{\mathbf{\Sigma}}_f(t) = \hat{\mathbf{D}}^f_t\hat{\mathbf{R}}^f_t\hat{\mathbf{D}}^f_t$, where $\hat{\mathbf{D}}^f_t = \mathrm{diag}\left(\hat{\sigma}^f_{t,1}, ..., \hat{\sigma}^f_{t,\hat{K}}\right)$.

3.2.1 Determining the number of latent factors

In practice, it is essential to estimate the number of latent factors K. During the past two decades, many techniques have been developed, including the influential method proposed by Bai and Ng (2002), which employs information criteria (IC). However, this method is specifically designed for strong factor models in which all K signal eigenvalues diverge proportionally to N. The simulation studies of Uematsu and Yamagata (2023) demonstrated that the $\hat{K}$ estimated by IC could be far from the true K, especially when there are weak factors. Thus, we recommend using the approach proposed by Onatski (2010), which is more appropriate to accommodate weak factors. Basically, this method determines the number of (weak) factors by $\hat{K} = \hat{K}(\delta)$ with

$$\hat{K}(\delta) = \{k = 1, \ldots, k_{\max} - 1 : \lambda_k - \lambda_{k+1} \geqslant \delta\},$$

where $\delta > 0$ is a fixed constant, $k_{\max} > K$ is an integer, and λ_k represents the kth largest eigenvalue of $(N \vee T)^{-1}\mathbf{U}\mathbf{U}'$. Empirically, δ is predetermined through calibrations using the edge distribution (ED) method in Onatski (2010). Notably, underestimating the number of latent factors can severely distort model performance, as it may result in the loss of significant systemic signals. In contrast, overestimating the number of factors generally causes only minor discrepancies. Therefore, we recommend adopting a slightly higher estimate of the number of latent factors in practical applications. See also the simulation evidence reported in Section 4.4.

3.3 Unconditional covariance matrix of devolatilized errors via POET

Since the systemic information is captured by the observable and latent factors, the error terms can reasonably be assumed to be weakly correlated. To account for these weak correlations, we impose conditional sparsity (2.8) on the unconditional target covariance matrix of the devolatilized errors, $\mathbf{S}_e$, in model (2.7). Specifically, we denote the residuals as $\hat{\mathbf{e}}_t = \hat{\mathbf{u}}_t - \hat{\mathbf{B}}\hat{\mathbf{f}}_t$ and fit a univariate GARCH model on $\hat{\mathbf{e}}_t = \{\hat{e}_{t,i}\}_{i=1,\ldots,N}$ for each i to obtain the univariate volatility estimate ${\hat{\sigma}^{e}_{t,i}}^2$. Subsequently, we devolatilize the residuals as $\hat{\eta}^e_{t,i} = \hat{e}_{t,i}/\hat{\sigma}^e_{t,i}$. We then estimate $\mathbf{S}_e$ by employing adaptive thresholding techniques (Cai and Liu, 2011; Fan et al., 2013) on the sample covariance matrix of $\hat{\boldsymbol{\eta}}^e_t = \{\hat{\eta}^e_{t,i}\}_{i=1,\ldots,N}$ as follows.

$$\hat{\mathbf{S}}^{\tau}_e = (\hat{s}^{\tau}_{i,j})_{N\times N} \text{ with } \hat{s}^{\tau}_{i,j} = \begin{cases} \hat{s}^e_{i,i} & \text{for } i = j, \\ \mathbb{T}(\hat{s}^e_{i,j}) & \text{for } i \neq j. \end{cases} \tag{3.13}$$

where $\hat{s}^e_{i,j}$ is the (i,j)-th entry of the sample covariance matrix of $\hat{\boldsymbol{\eta}}_t$, $\mathbb{T}(\cdot)$ is a soft thresholding function given by

$$\mathbb{T}(\hat{s}^e_{i,j}) = \text{sign}(\hat{s}^e_{i,j}) \max(0, |\hat{s}^e_{i,j}| - \tau_{i,j})$$

with

$$\tau_{ij} = C_\tau \omega_{NT} (\hat{\theta}^e_{i,j})^{1/2} \text{ and } \hat{\theta}^e_{i,j} = T^{-1} \sum_{t=1}^{T} (\hat{\eta}^e_{t,i} \hat{\eta}^e_{t,j} - \hat{s}^e_{i,j})^2 \tag{3.14}$$

for some sequence $\omega_{NT} > 0$ depending on N & T, and a sufficiently large constant $C_\tau > 0$. Other thresholding techniques, such as hard thresholding, can also be applied here; however, our preliminary research suggests that soft thresholding generally performs more robustly. Equations (3.13) and (3.14) together yield the target matrix estimator $\hat{\mathbf{S}}^\tau_e$. We use it in the DCC recursion (2.7) to get $\hat{\mathbf{Q}}^e_t$. The conditional correlation matrix estimator is then given by

$$\hat{\mathbf{R}}^e_t = (\operatorname{diag} \hat{\mathbf{Q}}^e_t)^{-1/2} \hat{\mathbf{Q}}^e_t (\operatorname{diag} \hat{\mathbf{Q}}^e_t)^{-1/2}.$$

Combining $\hat{\mathbf{R}}^e_t$ with the diagonalized univariate volatilities $\hat{\mathbf{D}}^e_t$ gives the volatility matrix estimator of $\boldsymbol{e}_t$, denoted by $\hat{\boldsymbol{\Sigma}}^\tau_e(t) = \hat{\mathbf{D}}^e_t \, \hat{\mathbf{R}}^e_t \, \hat{\mathbf{D}}^e_t$.

3.4 Volatility Matrix Estimate of $\mathbf{y}_t$

Finally, the volatility matrix of $\boldsymbol{y}_t$ is estimated by aggregating the components obtained in Sections 3.1–3.3 as follows:

$$\hat{\boldsymbol{\Sigma}}_y(t) \;=\; \hat{\mathbf{A}} \hat{\boldsymbol{\Sigma}}_x(t) \hat{\mathbf{A}}' + \hat{\mathbf{B}} \hat{\boldsymbol{\Sigma}}_f(t) \hat{\mathbf{B}}' + \hat{\boldsymbol{\Sigma}}^\tau_e(t). \tag{3.15}$$

Note that we provide a detailed algorithm flowchart in Fig. A.1, which summarizes the main steps of the proposed DCC-ePOET estimation procedure leading to the estimator in (3.15).

Remark 3.1 *If no observable factors are available, one may directly extract all latent factors from* $\mathbf{y}_t$ *via the approach in Section 3.2; see Section 5 for empirical performance. We do not consider a large number of observable factors (e.g., "factor zoos") to avoid the additional complexity of estimating a high-dimensional factor covariance matrix and the risk of overfitting.*

Remark 3.2 *In practice, it is necessary to determine the rate* ω_{NT} *in equation* (3.14)*. Dai et al.* (2024) *prove that, under certain mild assumptions, for the static error covariance matrix* $\boldsymbol{\Sigma}_e$ *and* *its ePOET estimator $\hat{\boldsymbol{\Sigma}}_e^\tau$, we have $\|\hat{\boldsymbol{\Sigma}}_e^\tau - \boldsymbol{\Sigma}_e\|_2 \lesssim \omega_{NT}^{1-q} m_N$ with high probability, where*

$$\omega_{NT} = T^{-1/2} \frac{N_1^{3/2}(N_1 \vee T)^{1/2}}{N_K^{3/2}(N_K \wedge T)^{1/2}} \log^{1/2}(N \vee T) \log^{1/2} T. \tag{3.16}$$

Given this rate, we can optimally select the threshold constant C_τ by multi-fold cross-validation (CV) as discussed in Appendix A.2[2].

3.5 Discussion on theoretical properties

Although a complete asymptotic analysis of the DCC-ePOET estimator is beyond the scope of this paper, in this subsection, we discuss the theoretical properties of each building component separately, drawing on the existing literature and clarifying the level at which results are currently available.

The estimator $\hat{\boldsymbol{\Sigma}}_y(t)$ in (3.15) is constructed from three building parts. We discuss the theoretical properties of them in turn.

(I) Observable factor part. Under standard regularity conditions for the DCC-GARCH model, including stationarity, finite moments of the innovations, a compact parameter space, and identifiability, as commonly imposed in the literature (see, e.g., Engle, 2002; Engle and Sheppard, 2001; Francq and Zakoian, 2019), the two-step QMLE yields consistent and asymptotically normal estimators of the DCC-GARCH parameters at the standard rate. Aielli (2013) identifies a consistency issue in the unconditional target matrix estimator and proposes a corrected (cDCC) variant with a heuristic consistency argument. That argument, however, is also confined to the parameter level. In our framework, the DCC-GARCH model is applied to the r-dimensional observable factors $\mathbf{x}_t$ with r fixed, so that the standard large sample properties hold. Consistency of $\hat{\mathbf{A}}$ at the standard parametric rate follows from OLS on (2.1) under the usual exogeneity and moment conditions without further difficulties.

(II) Latent factor part. In the absence of observable factors, Uematsu and Yamagata (2023) establish the convergence rate of the SOFAR estimator $(\hat{\mathbf{B}}, \hat{\mathbf{F}})$ under some standard assumptions (e.g., sub-Gaussian errors, identification restrictions, and sparsity levels $N_k = \lfloor N^{\alpha_k} \rfloor$ with $\alpha_k \in (0, 1]$). In the presence of observable factors, Dai et al. (2024)

[2] Note that other similar rates, such as $T^{-1/2} \log^{1/2} N$, are also feasible in practice, because the optimal choice of C_τ can automatically adjust during the CV procedure according to different values of ω_{NT}, yielding identical covariance matrix estimates.

further show that the estimator satisfies

$$\|\hat{\mathbf{B}} - \mathbf{B}^0\|_{\mathrm{F}} = O_p\left(N_1^{1/2}\omega_{NT}\right), \qquad \|\hat{\mathbf{F}} - \mathbf{F}^0\|_{\mathrm{F}} = O_p(T^{1/2}\omega_{NT}).$$

Here, ω_{NT} is the rate defined in (3.16), which incorporates an additional cost from the first-stage OLS regression relative to the pure latent sparsity-induced weak factor model of Uematsu and Yamagata (2023). Given the estimated latent factors $\hat{\mathbf{F}}$, the second-stage DCC-GARCH on $\hat{\mathbf{f}}_t$ inherits the parameter-level consistency discussed in part (I) with diminishing cost from estimating the latent factor part. Consistency of the estimated number of latent factors $\hat{K}$ under the edge-distribution method is established in Onatski (2010) and remains valid under the weak factor model (Uematsu and Yamagata, 2023).

(III) Idiosyncratic part. The adaptive thresholding method of Cai and Liu (2011) and Fan et al. (2013) controls the spectral norm of high-dimensional sparse covariance estimators in static factor settings. Under the conditional sparsity class (2.8), Dai et al. (2024) establish that the static ePOET estimator of the unconditional idiosyncratic covariance matrix $\mathbf{\Sigma}_e$ satisfies

$$\|\hat{\mathbf{\Sigma}}_e^{\tau} - \mathbf{\Sigma}_e\|_2 = O_p\left(\omega_{NT}^{1-q} m_N\right). \tag{3.17}$$

Our DCC-ePOET estimator extends this step to a time-varying covariance matrix through the decomposition $\mathbf{\Sigma}_e(t) = \mathbf{D}_t^e \mathbf{R}_t^e \mathbf{D}_t^e$. The asymptotic properties of DCC-GARCH follow the literature.

A further concern is the consistency of the estimated unconditional target matrix $\mathbf{S}_e^{\tau}$ that enters the DCC recursion on the devolatilized residuals (the same issue arises for the DCC modeling on factor part (I) and (II)). As pointed out by Aielli (2013), the unconditional target matrix in the original DCC formulation does not consistently estimate the unconditional second moment of the devolatilized innovations, i.e., $\mathbf{S}_e \neq \mathbb{E}(\boldsymbol{\eta}_t^e \boldsymbol{\eta}_t^{e\prime})$. They provide a corrected DCC (cDCC) variant to resolve this issue at the parameter level by reformulating the process as

$$\mathbf{Q}_t^e = (1 - a - b)\bar{\mathbf{S}}_e + a\boldsymbol{\eta}_{t-1}^{e*}\boldsymbol{\eta}_{t-1}^{e*\prime} + b\mathbf{Q}_{t-1}^e,$$

with $\boldsymbol{\eta}_{t-1}^{e*} = \left(\operatorname{diag}\mathbf{Q}_{t-1}^e\right)^{1/2}\boldsymbol{\eta}_{t-1}^e$, so that $\bar{\mathbf{S}}_e = \mathbb{E}(\boldsymbol{\eta}_t^{e*}\boldsymbol{\eta}_t^{e*\prime})$. Combining (3.17) with the consistency of the univariate volatility estimators, the consistency of the thresholded target matrix follows by standard arguments. In addition, we have verified in preliminary studies that the numerical difference between DCC and cDCC is negligible for our setting, consistent with Engle et al. (2019). We therefore keep the original DCC formulation in the present paper for simplicity.

Remarks on the joint asymptotic behavior of DCC-ePOET. A unified high-dimensional matrix-norm rate for $\hat{\boldsymbol{\Sigma}}_y(t)$ that simultaneously accounts for the DCC-GARCH plug-in factor model estimation errors in parts (I) and (II), and the plug-in POET estimation error in part (III), is, to our knowledge, currently beyond the reach of the large dimensional DCC-GARCH literature. Both the original DCC of Engle (2002) and the corrected cDCC theory of Aielli (2013) establish only parameter-level consistency. Our DCC-ePOET estimator, at the parameter level, should achieve similar asymptotic properties with extra costs from the aforementioned plug-in errors in the factor models and the POET framework. The recent DCC-NL framework of Engle et al. (2019) formulates the large sample properties of a DCC-NL estimator through a portfolio-level minimum variance loss rather than a matrix-norm rate. This contrasts with the related CCC-embedded Factor framework of Li et al. (2022b) and Li et al. (2022a), where a matrix-norm rate for the factor volatility matrix follows from the QMLE asymptotics of Francq and Zakoian (2019) under the CCC structure, leading to a matrix-norm error rate. The present framework upgrades the factor part from CCC to DCC, but the corresponding high-dimensional matrix-norm theory for DCC volatility matrices remains an open problem. Accordingly, following Engle et al. (2019), we evaluate the proposed estimator empirically through the quintessential risk loss in the following simulation studies (Section 4.3), where DCC-ePOET exhibits favorable performance.

4 Monte Carlo simulations

In this section, we investigate the finite sample perfromance of DCC-ePOET through Monte Carlo experiments, comparing several Factor-GARCH models to examine the distortions that arise when latent factors are ignored and the dynamics of the error terms are not adequately captured. Throughout the experiments, we fix the serial dimension at $T = 200$, while the cross-sectional dimension takes values $N = 100,\ 200$ and 300.

4.1 Design 1: data are generated from observable and latent factors

In this experiment, we consider a data generating process (DGP) from the DCC-ePOET model in Section 2. Recall the proposed linear factor model:

$$\mathbf{y}_t = \mathbf{A}^0\mathbf{x}_t + \mathbf{B}^0\mathbf{f}_t^0 + \mathbf{e}_t.$$

We consider that $\mathbf{x}_t$ follows a standard DCC-GARCH process in model (2.6). In this experiment, we construct the true $\mathbf{S}_x$ from the sample covariance matrix of the devolatized FF3 factor (Mkt, SMB, and HML), spanning from April 2, 2002 to March 30, 2023[3]. Each univariate volatility component in (2.6) follows a GARCH(1,1) process, namely $p = 1$ and $q = 1$, with $\omega = 0.02$, $\alpha = 0.05$, and $\beta = 0.9$. For the (weak) latent factors, we assume the number of factors $K^0 = 2$ with factor strengths $(0.8, 0.8)$. The factors follow a nearly identical DGP process as the observable factors but with the true unconditional covariance (for the DCC structure) derived from the sample covariance of the last two (devolatized) factors in the FF5 model, RMW and CMA. For the idiosyncratic term $\mathbf{e}_t = (e_{t,i})_{N\times 1}$, we assume that it follows a DCC-GARCH(1,1) structure in which the unconditional covariance matrix of the devolatilized error, $\mathbf{S}_e$, is block-diagonal. Each block of $\mathbf{S}_e$ contains 5×5 nonzero elements, with the within-block correlation set to 0.25. The GARCH and DCC parameters are identical to those of the factor components.

4.2 Design 2: data are generated from observable factors and only partial factors are known

In this experiment, we let

$$\mathbf{y}_t = \mathbf{A}^0\mathbf{x}_t + \mathbf{e}_t,$$

where the observable factors $\mathbf{x}_t$ follow the DCC-GARCH model with the same parameters as in Design 1. The number of factors $\mathbf{x}_t$ is set to five, and all factors are exact strong factors, that is, all the loadings are dense. The true unconditional target covariance matrix $\mathbf{S}_x$ is the sample covariance matrix of five devolatized FF5 factors spanning from April 2, 2002, to March 30, 2023. We retain the DGP process of the error term $\mathbf{e}_t$ from

[3] See https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Design 1. Most importantly, we assume that only the first three factors are observable in this experiment, while the remaining two are latent. This setting is more in line with real world financial data.

4.3 Results of Design 1 and Design 2

We use the following standard statistical losses to evaluate the accuracy of the estimate.

(a) Relative error: $Loss_R(t) = N^{-1/2}\|\boldsymbol{\Sigma}(t)^{-1/2}\hat{\boldsymbol{\Sigma}}(t)\boldsymbol{\Sigma}(t)^{-1/2} - \mathbf{I}_N\|_{\mathrm{F}}$.

(b) Quintessential risk loss:

$$Loss_Q(t) = \frac{\mathrm{Tr}\left(\hat{\boldsymbol{\Sigma}}(t)^{-1}\boldsymbol{\Sigma}(t)\hat{\boldsymbol{\Sigma}}(t)^{-1}\right)/N}{\left(\mathrm{Tr}(\hat{\boldsymbol{\Sigma}}(t)^{-1})/N\right)^2} - \frac{N}{\mathrm{Tr}(\boldsymbol{\Sigma}(t)^{-1})}.$$

Loss (a) is a relative error measure for covariance matrix estimates, as used in Fan et al. (2011) and Fan et al. (2013). Loss (b) relates to the out-of-sample portfolio variance risk proposed and theoretically analyzed by Engle et al. (2019). We compare our method with the CCC-POET approach by Li et al. (2022a), which models the volatility matrix with observable factors using CCC-GARCH and a static POET noise covariance matrix. Additionally, we consider an incomplete version of our DCC-ePOET method, incorporating only observable factors and an error term with a DCC and POET structure but excluding latent factors (referred to as DCC-ObsFactor inspired by Zhang and Chan (2009)). The simulation results for the average losses, $Loss_R = T^{-1}\sum_{t=1}^{T} Loss_R(t)$ and $Loss_Q = T^{-1}\sum_{t=1}^{T} Loss_Q(t)$, based on 100 experiments, are shown in Table 1 and 2. The estimated number of latent factors is presented in Table 3.

The simulation results reveal several key findings. (i) The proposed DCC-ePOET method exhibits robust and consistent performance across various scenarios, regardless of whether the observable factors are fully known, partially observed, $N < T$, or $N \geqslant T$. Even when the factors are strong, our method remains effective. (ii) The number of latent factors is estimated with high accuracy in all cases. Although slight overestimation occasionally occurs, it does not affect model performance, as overestimated factors generally have minimal impact (see additional details in Section 4.4). (iii) Our approach consistently outperforms the recently developed CCC-POET (Li et al., 2022a) and DCC-ObsFactor in all evaluation criteria. Its superiority becomes more pronounced as the cross-sectional dimension N increases, demonstrating scalability and robustness

Table 1 Performance of the DCC-ePOET, DCC-ObsFactor and CCC-POET estimates (**Design 1**)

N	Methods	$Loss_R$		$Loss_Q$	
		Mean	**Median**	**Mean**	**Median**
100	**DCC-ePOET**	1.047	1.043	0.194	0.192
	DCC-ObsFactor	1.340	1.334	1120.711	0.373
	CCC-POET	1.794	1.467	3.889	0.501
200	**DCC-ePOET**	1.308	1.304	0.221	0.222
	DCC-ObsFactor	2.078	2.075	0.557	0.497
	CCC-POET	2.888	2.854	7.766	4.587
300	**DCC-ePOET**	1.438	1.434	0.225	0.226
	DCC-ObsFactor	2.464	2.449	0.665	0.600
	CCC-POET	3.267	3.221	15.214	8.827

Table 2 Performance of the DCC-ePOET, DCC-ObsFactor and CCC-POET estimates when only partial factors are observable (**Design 2**)

N	Methods	$Loss_R$		$Loss_Q$	
		Mean	**Median**	**Mean**	**Median**
100	**DCC-ePOET**	1.036	1.033	0.193	0.193
	DCC-ObsFactor	1.526	1.453	245.354	0.499
	CCC-POET	4.621	4.520	6.153	3.574
200	**DCC-ePOET**	1.328	1.321	0.227	0.224
	DCC-ObsFactor	2.716	2.703	0.944	0.726
	CCC-POET	6.075	6.107	48.038	23.522
300	**DCC-ePOET**	1.452	1.448	0.253	0.230
	DCC-ObsFactor	3.198	3.187	1.522	0.920
	CCC-POET	6.864	6.861	155.873	64.811

Table 3 The number of estimated latent factors of DCC-ePOET

Design 1: True $K^0 = 2$					**Design 2**: True $K^0 = 2$				
N	**Mean**	**S.D.**	**Min**	**Max**	N	**Mean**	**S.D.**	**Min**	**Max**
100	2	0	2	2	100	2.05	0.26	2	4
200	2	0	2	2	200	2.04	0.20	2	3
300	2	0	2	2	300	2.01	0.10	2	3

in high-dimensional settings. Although CCC-POET and DCC-ObsFactor achieve acceptable relative errors (scaled by the cross-sectional dimension), they fail significantly in terms of the quintessential risk loss since they are unable to capture critical unobservable systemic information from latent factors or dynamic correlations in the residuals. For example, for DCC-ObsFactor, this failure arises because the omitted systematic signals of latent factors leak into the idiosyncratic residuals. The subsequent DCC fit on these contaminated residuals occasionally produces ill-conditioned $\hat{\boldsymbol{\Sigma}}_e(t)$, whose inversion further induces sharp spikes in the portfolio-level loss. The empirical consequences of these limitations are further illustrated in the minimum variance portfolio analysis of Section 5. (iv) As shown in Table 3, the number of latent factors is estimated with high accuracy in both designs. In Design 2, which violates the assumption that latent and observable factors are conditionally uncorrelated, only a slight overestimation is occasionally observed. More importantly, no underestimation is detected.

Remark 4.3 *We also conduct a computational speed comparison of our method against several alternative candidates. Detailed results are reported in Appendix A.3. As expected, DCC-ePOET is computationally slower than DCC-ObsFactor and CCC-POET, since DCC-ePOET involves an extensive regularization step in which the SOFAR procedure runs over a default grid of 100 candidate penalty terms to estimate the sparsity-induced weak factor loadings, whereas the latter two methods involve no regularization. Nevertheless, our method remains of a similar order of magnitude in computational cost as recent non-linear shrinkage estimators, in which regularization steps likewise contributes the dominant computational cost.*

4.4 Robustness check of number of latent factors

We also conduct simulation studies for cases in which the number of latent factors is known, underestimated, or overestimated. The data are generated based on Design 1 in Section 4.1, with the true number of latent factors set to $K^0 = 3$. As summarized in Table 4, underestimating the number of latent factors ($K = K^0 - 1$) leads to significant distortion and instability, as reflected in substantially higher values of $Loss_Q$. For instance, when $N = 100$, the mean $Loss_Q$ increases dramatically to 110.447, compared with only 0.193 when $K = K^0$. In contrast, overestimating the number of latent factors yields results that closely align with those obtained using the true K^0, indicating that overestimation does not significantly affect performance. These robust performances suggest that DCC-ePOET is able to handle the overestimation of latent factors while maintaining accuracy and reliability.

Table 4 Robustness check of DCC-ePOET estimate under different number of latent factors K

K		$N=100$		$N=200$		$N=300$	
		$Loss_R$	$Loss_Q$	$Loss_R$	$Loss_Q$	$Loss_R$	$Loss_Q$
$K=K^0-1=2$	**Mean**	1.285	110.447	1.958	1.623	2.296	1.720
	Median	1.280	1.205	1.958	1.170	2.301	1.229
$K=K^0=3$	**Mean**	1.034	0.193	1.312	0.221	1.426	0.226
	Median	1.031	0.192	1.309	0.220	1.419	0.226
$K=K^0+2=5$	**Mean**	1.269	0.200	1.414	0.225	1.536	0.230
	Median	1.260	0.201	1.414	0.224	1.538	0.229
$K=K^0+7=10$	**Mean**	1.272	0.200	1.407	0.224	1.522	0.229
	Median	1.263	0.200	1.403	0.224	1.521	0.229

5 Empirical studies

To evaluate the goodness of the estimation of volatility matrix in practice, we carry out a minimum variance portfolio (MVP) analysis of S&P 500 data based on several volatility models. MVP analysis aims to allocate N financial assets in such a way that the portfolio risk $\mathbf{w}'\tilde{\mathbf{\Sigma}}_t\mathbf{w}$ is minimized. Here, $\mathbf{w}$ is a vector of weights and $\tilde{\mathbf{\Sigma}}_t$ is a volatility matrix estimate of the given assets at time t. In detail, the MVP solves the following optimization problem:

$$\min_{\mathbf{w}} \mathbf{w}'\tilde{\mathbf{\Sigma}}_t\mathbf{w} \text{ subject to } \mathbf{w}'\mathbf{1}_N = 1, \tag{5.18}$$

where $\mathbf{1}_N = (1,\dots,1)'$. We allow short sales and ignore any transaction cost for simplicity. The optimal weight $\mathbf{w}^*$ is obtained by the quadratic problem (5.18) and the corresponding risk $\mathbf{R}_t^*$ is computed as

$$\mathbf{w}^* = \frac{\tilde{\mathbf{\Sigma}}_t^{-1}\mathbf{1}_N}{\mathbf{1}_N'\tilde{\mathbf{\Sigma}}_t^{-1}\mathbf{1}_N}, \quad \mathbf{R}_t^* = \mathbf{w}^{*\prime}\tilde{\mathbf{\Sigma}}_t\mathbf{w}^*. \tag{5.19}$$

We obtain the S&P 500 data from Yahoo Finance, comprising 5040 daily excess returns spanning from April 2, 2002, to April 5, 2022. This period covers approximately 20 years of trading, with an average of 21 trading days per month. In addition to the S&P

500 data, we also collect the FF5 factors data from the Kenneth R. French Data Library for the same time period.

The trading strategy is standard in the literature as follows. On the first trading day of each month, we construct an optimal portfolio using a candidate volatility matrix estimate of historical data from the preceding T days. The cross-sectional dimension N is fixed at 331, which is the maximum number of stocks available without missing values in the dataset. We set $T = 252$, approximately one year of trading days, reflecting a high-dimensional setting ($N > T$). Under a rolling window scheme, the vector of optimal portfolio weights ($\mathbf{w}_t^*$) is updated monthly for constructing the next month's portfolios until April 5, 2022. Once all out-of-sample portfolios are obtained, we calculate the out-of-sample variance (**Var**), the total out-of-sample excess returns (**TR**), and the mean Sharpe ratio (**SR**) following the standard procedure of DeMiguel et al. (2009) and Lam (2016). Also, drawing from existing literature such as Engle et al. (2019), we include the following two additional criteria:

- **Annual-SD**: We calculate the standard deviation of the out-of-sample returns for 19 trading years, then multiply by 252 to annualize.
- **Annual-IR**: We calculate the annualized information ratio by dividing the annualized average return (**AV**) by **Annual-SD**.

Comparison of all large-dimensional volatility matrix estimators in the literature is clearly beyond the scope of this study. Therefore, we focus on discussing the following representative candidates, including typical methods published in the past few years.

(i) **DCC-ePOET-1**: Our proposed method.

(ii) **DCC-ePOET-2**: This method corresponds to a special case of **DCC-ePOET-1** in which no observable factor is provided.

(iii) **DCC-ObsFactor**: This method is **DCC-ePOET-1** without latent (weak) factors, inspired by Zhang and Chan (2009), but includes a sparse noise covariance matrix to tackle the singularity issues arising from large N.

(iv) **ePOET**: A static covariance matrix proposed by Dai et al. (2024), whereby the covariance matrix consists of observable factors, latent factors, and the sparse noise covariance matrix. This estimator is the winner in their study.

(v) **CCC-POET**: This method is proposed by Li et al. (2022a), in which the volatility matrix is composed of FF5 factors embedded with CCC-GARCH, plus the (static) sparse noise covariance matrix estimate according to Fan et al. (2011).

(vi) **CCC-Diag**: This is similar to (v), but the noise component is diagonal to the univariate volatilities of the residuals.

(vii) **DCC-NL**: This method is a modified DCC-GARCH that employs a nonlinear shrinkage technique on the unconditional covariance matrix to address singularity issues (Engle et al., 2019).

(viii) **AFM1-DCC-NL**: This method is proposed by De Nard et al. (2021), whereby the factor part is captured only by the market factor in FF5, and DCC-NL is conducted on residuals.

(ix) **AFM5-DCC-NL**: Similar to AFM1-DCC-NL, while all five factors in FF5 are used.

We employ GARCH(1,1) in each univariate volatility estimate as Hansen and Lunde (2005) suggested that GARCH(1,1) is typically adequate to capture the clustered nature of volatilities observed in the data for most scenarios. In each window, the number of latent factors is determined according to the method of Onatski (2010) and the optimal threshold tuning parameters are chosen using CV.

Table 5 Performance of the nine methods in out-of-sample minimum variance portfolio analysis.

Model \ Criteria	Var	Annual-SD	SR	Annual-IR	TR
DCC-ePOET-1	**0.445**	**10.840**	**2.281**	0.945	194.633
DCC-ePOET-2	0.464	11.049	2.192	**0.952**	199.862
DCC-ObsFactor	0.666	13.298	2.008	0.618	156.077
ePOET	0.483	11.278	2.228	0.911	195.270
CCC-POET	1.469	19.569	2.257	0.744	**278.329**
CCC-Diag	0.653	13.187	1.824	0.675	169.228
DCC-NL	0.559	12.110	2.196	0.838	192.845
AFM1-DCC-NL	0.488	11.347	1.872	0.744	160.364
AFM5-DCC-NL	0.461	11.046	2.099	0.877	184.060

Table 5 presents the out-of-sample MVP performance for each method. The primary performance metrics are the out-of-sample variance (**Var**) and the annualized standard deviation (**Annual-SD**), given that our main objective is to minimize portfolio risk. The mean Sharpe ratios (**SR**) and annualized information ratio (**Annual-IR**) are considered as the secondary criteria. Table 5 highlights several key findings: (i) Our method demonstrates exceptional performance in risk reduction, surpassing other approaches in two key metrics, **Var** and **Annual-SD**, while maintaining a well-balanced trade-off between return and risk, as evidenced by **SR** and **Annual-IR**. (ii) Although using only latent factors produces acceptable results, incorporating some known information, namely observable factors, into the DCC-ePOET framework further enhances risk reduction. This highlights the advantage of combining observable and latent factors to create a more comprehensive model. (iii) Models relying solely on observable factors risk losing critical latent factor information, making it challenging to fully capture systemic signals. Such limitations may result in the failure of the volatility matrix estimation in practice, consistent with the quintessential loss results in previous simulation studies. Fig. 3 shows the number of latent factors estimated within the rolling windows for the residuals of FF5, revealing that nearly every window detects at least one latent factor, with up to ten latent factors identified in certain periods. This underscores the importance of considering latent factors to improve the robustness and accuracy of the model.

Fig. 3 Number of estimated latent factors over time in the residuals of FF5 model

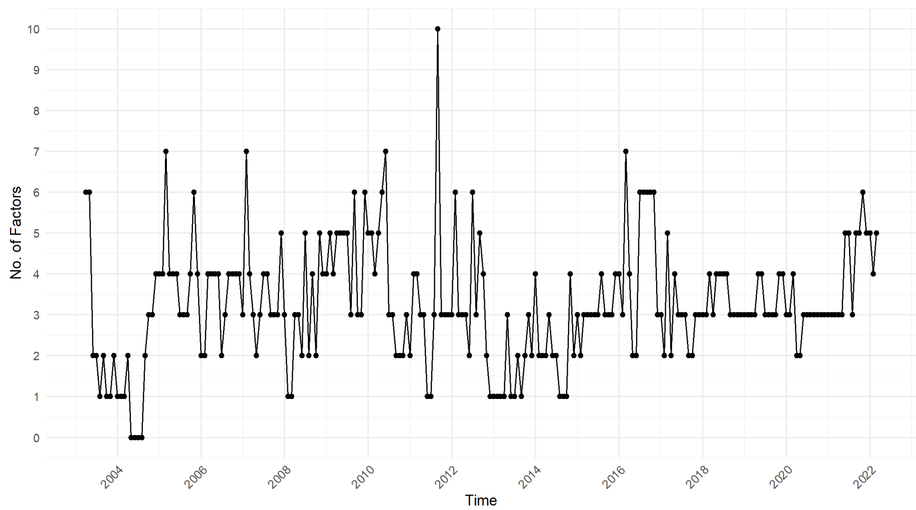

6 Conclusion remarks

This study introduces a concise yet comprehensive Factor-GARCH model, namely, DCC-ePOET, designed to estimate large volatility matrices. It addresses the challenges arising from omitted variables and singularity issues within the Factor-GARCH family in high-dimensional settings. The DCC-GARCH framework is embedded into both observable and latent factors, and additional covariance information in the residuals is captured through sWFs and a sparse noise target matrix constructed using ePOET techniques. The simulation results demonstrate that the proposed estimator performs well even when the observable factors are only partially known. The estimation of the number of latent factors shows high accuracy and applicability. Additionally, the out-of-sample MVP analysis validates the robustness of the DCC-ePOET estimator, which consistently outperforms other candidate models, particularly in the context of risk minimization.

Future research could extend this work in several directions: (i) The current model uses the simple GARCH model for the univariate component. Future studies could build on this by exploring more advanced extensions within the univariate GARCH family, such as GJR-GARCH that accounts for asymmetric volatility responses to positive and negative shocks (Glosten et al., 1993), better suited to more complex volatility dynamics. Furthermore, incorporating parameter-scaled QMLE for GARCH models with heavy-tailed likelihoods, as proposed by Fan et al. (2014), could provide a possible improvement of our method when the factor and noise innovations exhibit heavy tails. (ii) This study assumes that factor loadings are time-invariant. To enhance the model's flexibility, future research could allow the loadings of observable factors to vary over time. The investigation of time-varying loadings for sWFs and their application to volatility matrix estimation remains an open question worthy of further exploration. (iii) Finally, this work can be applied to various high-dimensional financial studies, such as modeling large-scale financial asset connectedness and assessing value at risk.

Statements and Declarations

The authors declare that there are no conflicts of interest to disclose. They also declare no known competing financial interests or personal relationships that might have affected the work presented in this study.

References

Aielli, G. P. (2013). Dynamic conditional correlation: on properties and estimation. *Journal of Business & Economic Statistics*, 31(3):282–299.

Alexander, C. (2000). A primer on the orthogonal garch model. Manuscript ISMA Centre, University of Reading, UK, 2.

Alexander, C. (2002). Principal component models for generating large garch covariance matrices. *Economic Notes*, 31(2):337–359.

Alexander, C. and Chibumba, A. M. (1996). Multivariate orthogonal factor garch. Mimeo, University of Sussex.

Allen, L., Bali, T. G., and Tang, Y. (2012). Does systemic risk in the financial sector predict future economic downturns? *The Review of Financial Studies*, 25(10):3000–3036.

Arouri, M., Mhadhbi, M., and Shahrour, M. H. (2025). Dynamic connectedness and hedging effectiveness between green bonds, esg indices, and traditional assets. *European Financial Management*, 31(5):1704–1719.

Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221.

Bodilsen, S. T. (2025). Large-dimensional portfolio selection with a high-frequency-based dynamic factor model. *Journal of Financial Econometrics*, 23(2).

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch model. *The Review of Economics and Statistics*, pages 498–505.

Brownlees, C. and Engle, R. F. (2017). Srisk: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies*, 30(1):48–79.

Cai, T. and Liu, W. (2011). Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*, 106(494):672–684.

Chen, J., Li, D., and Linton, O. (2019). A new semiparametric estimation approach for large dynamic covariance matrices with multiple conditioning variables. *Journal of Econometrics*, 212(1):155–176.

Choi, S. H. and Kim, D. (2023). Large volatility matrix analysis using global and national factor models. *Journal of Econometrics*, 235(2):1917–1933.

Choi, S. H. and Kim, D. (2025). Large global volatility matrix analysis based on observation structural information. *Econometric Theory*, 41(6):1452–1467.

Chou, R. Y., Wu, C.-C., and Liu, N. (2009). Forecasting time-varying covariance with a range-based dynamic conditional correlation model. *Review of Quantitative Finance and Accounting*, 33(4):327–345.

Dai, R., Uematsu, Y., and Matsuda, Y. (2024). Estimation of large covariance matrices with mixed factor structures. *The Econometrics Journal*, 27(1):62–83.

De Nard, G., Ledoit, O., and Wolf, M. (2021). Factor models for portfolio selection in large dimensions: The good, the better and the ugly. *Journal of Financial Econometrics*, 19(2):236–257.

DeMiguel, V., Garlappi, L., Nogales, F. J., and Uppal, R. (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management science*, 55(5):798–812.

Dendramis, Y., Giraitis, L., and Kapetanios, G. (2021). Estimation of time-varying covariance matrices for large datasets. *Econometric Theory*, 37(6):1100–1134.

Ding, Y., Li, Y., and Zheng, X. (2021). High dimensional minimum variance portfolio estimation under statistical factor models. *Journal of Econometrics*, 222(1):502–515.

Ding, Z. (1994). *Time series analysis of speculative returns*. PhD thesis, University of California, San Diego, Department of Economics.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, pages 987–1007.

Engle, R. F., Ledoit, O., and Wolf, M. (2019). Large dynamic covariance matrices. *Journal of Business & Economic Statistics*, 37(2):363–375.

Engle, R. F. and Sheppard, K. (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate garch.

Engle, R. F. and Siriwardane, E. N. (2018). Structural garch: the volatility-leverage connection. *The Review of Financial Studies*, 31(2):449–492.

Fan, J., Fan, Y., and Lv, J. (2008). High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1):186–197.

Fan, J. and Kim, D. (2018). Robust high-dimensional volatility matrix estimation for high-frequency factor model. *Journal of the American Statistical Association*, 113(523):1268–1283.

Fan, J., Liao, Y., and Mincheva, M. (2011). High dimensional covariance matrix estimation in approximate factor models. *Annals of Statistics*, 39(6):3320.

Fan, J., Liao, Y., and Mincheva, M. (2013). Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, 75(4).

Fan, J., Qi, L., and Xiu, D. (2014). Quasi-maximum likelihood estimation of garch models with heavy-tailed likelihoods. *Journal of Business & Economic Statistics*, 32(2):178–191.

Fan, Q., Wu, R., Yang, Y., and Zhong, W. (2024). Time-varying minimum variance portfolio. *Journal of Econometrics*, 239(2):105339.

Feng, G., Giglio, S., and Xiu, D. (2020). Taming the factor zoo: A test of new factors. *The Journal of Finance*, 75(3):1327–1370.

Fiszeder, P., Fałdziński, M., and Molnár, P. (2023). Modeling and forecasting dynamic conditional correlations with opening, high, low, and closing prices. *Journal of Empirical Finance*, 70:308–321.

Francq, C. and Zakoian, J.-M. (2019). *GARCH models: Structure, statistical inference and financial applications*. John Wiley & Sons.

Giglio, S. and Xiu, D. (2021). Asset pricing with omitted factors. *Journal of Political Economy*, 129(7):1947–1990.

Girardi, G. and Ergün, A. T. (2013). Systemic risk measurement: Multivariate garch estimation of covar. *Journal of Banking & Finance*, 37(8):3169–3180.

Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):1779–1801.

Guo, S., Box, J. L., and Zhang, W. (2017). A dynamic structure for high-dimensional covariance matrices and its application in portfolio allocation. *Journal of the American Statistical Association*, 112(517):235–253.

Hafner, C. M. and Preminger, A. (2009). Asymptotic theory for a factor garch model. *Econometric Theory*, 25(2):336–363.

Hansen, P. R. and Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a garch (1, 1)? *Journal of Applied Econometrics*, 20(7):873–889.

Joseph, N. L., Vo, T. T. A., Mobarek, A., and Mollah, S. (2020). Volatility and asymmetric dependence in central and east european stock markets. *Review of Quantitative Finance and Accounting*, 55(4):1241–1303.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.

Ke, Y., Lian, H., and Zhang, W. (2022). High-dimensional dynamic covariance matrices with homogeneous structure. *Journal of Business & Economic Statistics*, 40(1):96–110.

Kuo, C.-Y. and Chiang, S.-M. (2025). Spillover effects and network connectedness among stock markets: evidence from the us and asia. *Review of Quantitative Finance and Accounting*, 64(1):1–52.

Lam, C. (2016). Nonparametric eigenvalue-regularized precision or covariance matrix estimator. *Annals of Statistics*, 44(3):928–953.

Lanne, M. and Saikkonen, P. (2007). A multivariate generalized orthogonal factor garch model. *Journal of Business & Economic Statistics*, 25(1):61–75.

Ledoit, O. and Wolf, M. (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices. *The Annals of Statistics*, 40(2):1024.

Ledoit, O. and Wolf, M. (2017). Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets goldilocks. *The Review of Financial Studies*, 30(12):4349–4388.

Li, D. (2024). Estimation of large dynamic covariance matrices: A selective review. *Econometrics and Statistics*, 29:16–30.

Li, X., Li, Y., Pan, J., and Zhang, X. (2022a). A factor-garch model for high dimensional volatilities. *Acta Mathematicae Applicatae Sinica, English Series*, 38(3):635–663.

Li, X., Zhang, X., and Li, Y. (2022b). High-dimensional conditional covariance matrices estimation using a factor-garch model. *Symmetry*, 14(1):158.

Lyle, M. R. and Yohn, T. L. (2021). Fundamental analysis and mean-variance optimal portfolios. *The Accounting Review*, 96(6):303–327.

Moura, G. V., Santos, A. A., and Ruiz, E. (2020). Comparing high-dimensional conditional covariance matrices: Implications for portfolio selection. *Journal of Banking & Finance*, 118:105882.

Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *The Review of Economics and Statistics*, 92(4):1004–1016.

Shi, F., Shu, L., and Gu, X. (2022). An enhanced factor model for portfolio selection in high dimensions. *Journal of Financial Econometrics*.

Uematsu, Y., Fan, Y., Chen, K., Lv, J., and Lin, W. (2019). SOFAR: large-scale association network learning. *IEEE transactions on information theory*, 65(8):4924–4939.

Uematsu, Y. and Yamagata, T. (2023). Estimation of sparsity-induced weak factor models. *Journal of Business & Economic Statistics*, 41(1):213–227.

Van der Weide, R. (2002). Go-garch: A multivariate generalized orthogonal garch model. *Journal of Applied Econometrics*, 17(5):549–564.

Vrontos, I. D., Dellaportas, P., and Politis, D. N. (2003). A full-factor multivariate garch model. *The Econometrics Journal*, 6(2):312–334.

Wang, H., Peng, B., Li, D., and Leng, C. (2021). Nonparametric estimation of large covariance matrices with conditional sparsity. *Journal of Econometrics*, 223(1):53–72.

Wu, R. and Qin, Z. (2024). Asymmetric volatility spillovers among new energy, esg, green bond and carbon markets. *Energy*, 292:130504.

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

Zhang, W., He, X., and Hamori, S. (2022). Volatility spillover and investment strategies among sustainability-related financial indexes: Evidence from the dcc-garch-based dynamic connectedness and dcc-garch t-copula approach. *International Review of Financial Analysis*, 83:102223.

A Online Appendix

A.1 Flowchart pipeline for DCC-ePOET estimation

For ease of demonstration, we provide a simple pipeline illustrating each step of estimating DCC-ePOET as in the following Fig. A.1.

Fig. A.1 Estimation pipeline of the DCC-ePOET framework

Step 1. Input $\mathbf{y}_t \in \mathbb{R}^N$

$\downarrow$

Step 2. Observable factor regression (OLS)

$$\mathbf{y}_t = \hat{\mathbf{A}}\,\mathbf{x}_t + \hat{\mathbf{u}}_t \quad \longrightarrow \quad \hat{\mathbf{A}},\ \hat{\mathbf{U}}$$

$\downarrow$

Step 3. Weak latent factor estimation (SOFAR)

· Determine $\hat{K}$ in $\hat{\mathbf{U}}$ via the edge distribution method (Onatski, 2010).

· Estimate $(\hat{\mathbf{B}}, \hat{\mathbf{F}})$ by sparse orthogonal factor regression (SOFAR) (Uematsu et al., 2019):

$$\min_{\tilde{\mathbf{B}}, \tilde{\mathbf{F}}} \tfrac{1}{2}\|\hat{\mathbf{U}} - \tilde{\mathbf{F}}\tilde{\mathbf{B}}'\|_F^2 + \eta_{NT}\|\tilde{\mathbf{B}}\|_1.$$

Obtain $\hat{\mathbf{B}}$, $\hat{\mathbf{F}}$, and errors $\hat{\mathbf{e}}_t = \hat{\mathbf{u}}_t - \hat{\mathbf{B}}\hat{\mathbf{f}}_t$.

$\downarrow$

Step 4a. Factor volatility via DCC-GARCH

· $\mathbf{x}_t$: fit GARCH$\rightarrow \hat{\sigma}^x_{t,l} \rightarrow \hat{\mathbf{D}}^x_t$

· Devolatilize: $\hat{\eta}^x_{t,l} = x_{t,l}/\hat{\sigma}^x_{t,l}$

· $\hat{\mathbf{S}}^x = T^{-1}\sum_{t=1}^T \hat{\boldsymbol{\eta}}^x_t \hat{\boldsymbol{\eta}}^{x\prime}_t$

· $\hat{\mathbf{Q}}^x_t = (1-a-b)\hat{\mathbf{S}}^x + a\hat{\boldsymbol{\eta}}^x_{t-1}\hat{\boldsymbol{\eta}}^{x\prime}_{t-1} + b\hat{\mathbf{Q}}^x_{t-1}$

· $\hat{\mathbf{R}}^x_t = (\operatorname{diag}\hat{\mathbf{Q}}^x_t)^{-1/2}\,\hat{\mathbf{Q}}^x_t\,(\operatorname{diag}\hat{\mathbf{Q}}^x_t)^{-1/2}$

· $\hat{\boldsymbol{\Sigma}}_x(t) = \hat{\mathbf{D}}^x_t\,\hat{\mathbf{R}}^x_t\,\hat{\mathbf{D}}^x_t$

· $\hat{\mathbf{f}}_t$: same DCC as above $\rightarrow \hat{\boldsymbol{\Sigma}}_f(t)$

$\downarrow$

Factor component

$$\hat{\mathbf{A}}\,\hat{\boldsymbol{\Sigma}}_x(t)\,\hat{\mathbf{A}}' + \hat{\mathbf{B}}\,\hat{\boldsymbol{\Sigma}}_f(t)\,\hat{\mathbf{B}}'$$

Step 4b. Noise volatility via DCC & POET

· $\hat{e}_{t,i}$: fit GARCH$\rightarrow \hat{\sigma}^e_{t,i} \rightarrow \hat{\mathbf{D}}^e_t$

· Devolatilize: $\hat{\eta}^e_{t,i} = \hat{e}_{t,i}/\hat{\sigma}^e_{t,i}$

· $\hat{\mathbf{S}}^e = T^{-1}\sum_{t=1}^T \hat{\boldsymbol{\eta}}^e_t \hat{\boldsymbol{\eta}}^{e\prime}_t$

· Adaptive thresholding (C_τ by CV) $\rightarrow \hat{\mathbf{S}}^\tau_e$

· $\hat{\mathbf{Q}}^e_t = (1-a-b)\hat{\mathbf{S}}^\tau_e + a\hat{\boldsymbol{\eta}}^e_{t-1}\hat{\boldsymbol{\eta}}^{e\prime}_{t-1} + b\hat{\mathbf{Q}}^e_{t-1}$

· $\hat{\mathbf{R}}^e_t = (\operatorname{diag}\hat{\mathbf{Q}}^e_t)^{-1/2}\hat{\mathbf{Q}}^e_t(\operatorname{diag}\hat{\mathbf{Q}}^e_t)^{-1/2}$

$\downarrow$

Error component

$$\hat{\boldsymbol{\Sigma}}^\tau_e(t) = \hat{\mathbf{D}}^e_t\,\hat{\mathbf{R}}^e_t\,\hat{\mathbf{D}}^e_t$$

$\searrow \quad \swarrow$

Step 5. Final estimator (DCC-ePOET)

$$\hat{\boldsymbol{\Sigma}}_y(t) = \hat{\mathbf{A}}\,\hat{\boldsymbol{\Sigma}}_x(t)\,\hat{\mathbf{A}}' + \hat{\mathbf{B}}\,\hat{\boldsymbol{\Sigma}}_f(t)\,\hat{\mathbf{B}}' + \hat{\boldsymbol{\Sigma}}^\tau_e(t)$$

$\downarrow$

Output: $\hat{\boldsymbol{\Sigma}}_y(t) \in \mathbb{R}^{N\times N}$

A.2 Choice of the threshold constant for DCC-ePOET

In practice, we have to determine the tuning constant C_τ in the threshold level $\tau_{ij} = C_\tau \omega_{NT}(\theta_{ij})^{1/2}$. Following the procedures of Fan et al. (2013), we use the multiple cross-validation method to select C_τ as follows. Repeat Steps 1–2 below H times and denote each time as h.

Step 1. Obtain the devolatilized residuals $\{\hat{\boldsymbol{\eta}}_t^e\}_{t=1}^T$ based on our two-step regressions.

Step 2. Randomly divide $\{\hat{\boldsymbol{\eta}}_t^e\}_{t=1}^T$ into two sets, denoted by M_1 and M_2. Let M_1 be the training set $\{\hat{\boldsymbol{\eta}}_t^e\}_{t\in M_1}$, and M_2 be the validation set $\{\hat{\boldsymbol{\eta}}_t^e\}_{t\in M_2}$. The training set has size $T(M_1)$ and the validation set has size $T(M_2)$, where we simply set $T(M_1) = \lfloor T(1-\log^{-1} T)\rfloor$ and $T(M_1) + T(M_2) = T$.

Step 3. Obtain the optimal tuning parameter C_τ^* according to the Frobenius loss:

$$C_\tau^* = \underset{C_\tau \in [C^{\min},\ \bar{C}]}{\arg\min} \frac{1}{H} \sum_{h=1}^{H} \left\| \hat{\mathbf{S}}_e^\tau (C_\tau)^{M_1,h} - \hat{\mathbf{S}}_e^{M_2,h} \right\|_F^2 .$$

For each h, $\hat{\mathbf{S}}_e^\tau (C_\tau)^{M_1,h}$ is the ePOET estimator of $\{\hat{\boldsymbol{\eta}}_t^e\}_{t\in M_1}^h$ with the threshold constant C_τ. $\hat{\mathbf{S}}_e^{M_2,h}$ is the sample covariance matrix of $\{\hat{\boldsymbol{\eta}}_t^e\}_{t\in M_2}^h$. Note that C_τ^* is selected from an interval $[C^{\min},\ \bar{C}]$, where $C^{\min}$ is the minimum constant to ensure the positive definiteness and $\bar{C}$ is some large constant set by users.

A.3 Discussion on computational speed of DCC-ePOET

We have conducted a comprehensive comparison of computational time across the main candidate methods for $N \in \{100, 200, 300\}$ with $T = 200$, based on 10 independent simulation replications. The average computational time per experiment, measured in CPU seconds, is reported in Table A.1.

The findings can be summarized as follows. Methods that rely solely on factor extraction, such as CCC-POET and DCC-ObsFactor, are computationally the fastest. This is because these methods do not involve an additional regularization step. By contrast, the proposed DCC-ePOET requires a regularization step on the latent factor loadings within the estimation pipeline. Specifically, the SOFAR procedure in DCC-ePOET recovers the sparse factor loadings, with the optimal regularization penalty selected from 100 candidate values by default. Shrinkage-based methods, such as DCC-NL (Engle

Table A.1 Computational time comparison (average in CPU seconds)

Method	$N = 100$	$N = 200$	$N = 300$
DCC-ePOET	6.911	13.864	22.472
DCC-ObsFactor	0.322	1.005	2.103
CCC-POET	0.320	0.986	2.087
DCC-NL	3.464	7.365	12.619
AFM5-DCC-NL	3.151	8.025	14.739

Notes: The candidates reported here are selected to cover the full range of computational patterns appearing in Section 5. Several methods used in Section 5 are omitted here because they share the same dominant computational step as one of the reported methods. For example, methods differing only in the factor loading estimator (e.g., PC versus SOFAR, and 1 versus 5 observable factors) but retaining identical DCC and regularization steps, and their running times are therefore expected to be of the same order of magnitude as those of their structural counterparts. This comparison study is conducted on an Apple Silicon Mac with 32GB RAM.

et al., 2019) and AFM5-DCC-NL (De Nard et al., 2021), do not require such an extensive search over tuning parameters. This explains why DCC-ePOET is slower than these methods by roughly a factor of two, while still remaining of a similar order across all values of N.

We view this additional computational cost of DCC-ePOET as a natural and acceptable consequence of the richer and general model structure. Even when $N = 300$, a single DCC-ePOET estimate is obtained in approximately 22 seconds, which remains well within practical limits for typical financial applications. It is also worth noting that the SOFAR procedure adopted in DCC-ePOET is based on an augmented Lagrangian with block coordinate descent (Uematsu et al., 2019). Recent advances in large-scale optimization may offer further computational speed-ups, which we leave for future research, as it is beyond the scope of the present paper.