TI 2026-061/III
Tinbergen Institute Discussion Paper

Composite Univariate Modeling of Realized Covariance Matrix Dynamics and Volatility-at-Risk

Justus Holman[1]
Andre Lucas[1]
Anne Opschoor[1]

1 Vrije Universiteit Amsterdam and Tinbergen Institute

Composite Univariate Modeling of Realized Covariance Matrix Dynamics and Volatility-at-Risk

Justus Holman* Andre Lucas* Anne Opschoor*

August, 2026

Abstract

We propose a new model for realized covariance matrix dynamics using a composite of univariate time series models for its realized eigenvalues, linked by a copula function. The dynamics of each eigenvalue are based on a conditional F distribution, thus allowing for fat-tailedness and outliers in the realized covariance matrices. Given its composition from univariate elements, the static parameters of the new model can be estimated efficiently by maximum likelihood. In an empirical application, we show that the new model outperforms relevant recent benchmarks in an extensive portfolio (Conditional) Volatility-at-Risk (VolaR) application.

Key Words: Realized Covariance Matrix, HAR model, Dynamic Conditional F Distribution, Dynamic Eigenvalues, Volatility-at-Risk (VolaR).

JEL Classification: C32, C58, C22.

1. Introduction

Portfolio volatility is a fundamental driver of asset pricing, risk management and volatility trading in financial institutions (Christoffersen and Diebold, 2000; Andersen et al., 2007; Sinclair, 2013; Caporin et al., 2017b). Today's availability of high-frequency data allows for volatility modeling based on realized measures of the covariance matrix (see, for instance, Jin and Maheu, 2013; Noureldin et al., 2012; Opschoor et al., 2018).

Recent literature explores the use of decomposition techniques and factors to model the most important aspects of the time-variation in realized measures, trying to limit the proliferation of parameters and increase estimation efficiency (Bauwens et al., 2012; Bauwens and Xu, 2023; Vassallo et al., 2021). For realized covariance

*Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands. Holman and Opschoor thank the Dutch National Science Foundation (NWO) for financial support under grant VI.VIDI.201.079.

matrices the eigenvalue-eigenvector-decomposition is arguably one of the most important decomposition methods, as the eigenvalues and eigenvectors summarize the fundamental characteristics of a matrix (Laloux et al., 2000; Horn and Johnson, 2012; Chamberlain and Rothschild, 1983). Following this idea, Hetland et al. (2023) introduce a model in which the eigenvalues of the conditional covariance matrix of daily returns evolve over time using the score-driven dynamics of Creal et al. (2013) and Harvey (2013). Gribisch and Stollenwerk (2020), by contrast, build a model on the spectral decomposition of realized covariance matrices. More specifically, in their DPC-CAW model, latent eigenvalues of realized covariance matrices are modeled by a Gamma distribution and GARCH type dynamics. The Gamma distribution, however, is less suitable for dealing with outliers and incidental large observations, which are often encountered for realized (co)variances. Moreover, eigenvalues may not move independently, but exhibit co-movement over time.

In this paper, we develop a new model focusing directly on the realized eigenvalues using a combination of univariate (conditional) marginal models and a copula. Jointly, the marginal models and the copula make up a much more flexible model class for realized covariance matrix data, one that turns out to be more in line with the empirical data. For the dynamics of each of the separate eigenvalues, we adopt a HAR specification with a conditional dynamic fat-tailed F distribution, following evidence that both heterogeneous dynamics (Corsi, 2009) and fat-tailed conditional distributions (Opschoor et al., 2018) are important stylized facts of realized (co)variances. By allowing each eigenvalue to follow univariate dynamics, the new model is flexible, numerically efficient, easily scalable, and straightforward to estimate using a standard maximum likelihood estimator (MLE). As the realized eigenvalues may be conditionally dependent (Zumbach, 2011), we complete the model by adding a copula that ties together the different marginal conditional distributions of the eigenvalues. We label the new model the Realized λ-HAR (ReλHAR) model, following the λ-GARCH label for the model of Hetland et al. (2023).

We estimate our model parameters in a two-step approach following Joe and Xu (2005) and Patton (2006) and show consistency and asymptotic normality of marginal models. In a controlled setting, we show that the finite sample properties of the parameters of the marginal models can be well approximated by the theoretical asymptotic results.

We apply our new model to daily realized (co)variances of five U.S. financial institutions during 2002-2023. To test whether the new ReλHAR model succeeds in capturing the large, unpredictable increments and volatility bursts that characterize typical realized volatility time series (Caporin et al., 2017a; Caporin et al., 2017b), we evaluate the model in terms of its portfolio Volatility-at-Risk (VolaR; see Caporin et al., 2017b) and conditional portfolio VolaR (or volatility-expected-shortfall), rather than in terms of its conditional mean forecasts only (as in, e.g., Vassallo et al., 2021; Zhou et al., 2022; Opschoor et al., 2025). Using the VolaR and conditional VolaR, we compare the new ReλHAR model to several benchmarks from the recent literature, including a version of the DPC-CAW model of Gribisch and Stollenwerk (2020), the Matrix-F model of Opschoor et al. (2018), and a DCC Conditional Autoregressive F model. Despite the fact that we improve each of these benchmark models with additional empirically congruent features, we find that the ReλHAR model performs better: it attains accurate nominal VolaR estimates across different quantile levels and has the lowest Tick-Loss. In addition, the conditional VolaR estimates are superior against all competitors. This suggests that the composite modeling strategy of the ReλHAR model forecasts the tails of the portfolio variance distribution more accurately than its peers, both conditionally and unconditionally.

Interestingly, we empirically find a positive conditional correlation between the eigenvalues. Including this via the copula part of the ReλHAR model, we can further improve the shape of the portfolio variance forecast distribution, though not always: in some cases a restricted model with conditional independence between the eigenvalues performs better in conditional backtests due to the more conservative eigenvalue forecasts it produces.

This paper touches several strands of the literature. Many earlier contributions focus on modeling the dynamics of the full realized covariance matrix, such as the conditional Wishart-type models (Gouriéroux et al., 2009; Noureldin et al., 2012) and their generalizations such as the matrix F and F-Riesz type models (Stollenwerk, 2026; Opschoor et al., 2025). In contrast to these approaches, the ReλHAR model focuses on the realized eigenvalues as the dominant source of time variation in the realized covariance matrices. Moreover, the composite approach in the ReλHAR model of univariate models for each of the eigenvalues allows eigenvalue-specific tail behavior, in contrast to, e.g., the matrix-F formulations of Opschoor et al. (2018). The latter impose a common tail behavior across all elements of the entire covariance matrix. Allowing for more tail heterogeneity aligns better with the empirical evidence of heterogeneous tail risk across volatility components (see, e.g., Näf et al., 2019; Hediger et al., 2023). The new model does so in a way that is invariant to the ordering of the variables in the system, in contrast to the heterogeneous tail behavior allowed for in, for instance, the F-Riesz distribution (Stollenwerk, 2026; Opschoor et al., 2025). The ReλHAR also allows for additional flexibility in the dependence structure between eigenvalues through the choice of a copula.

The rest of this paper is set up as follows. Section 2 introduces the model, while Section 3 discusses its properties. Section 4 discusses our *VolaR* and conditional *VolaR* backtesting procedures, which are applied in Section 5 to evaluate the performance of the different models for US stock data. Section 6 concludes. As a notational convention, we use symbols like a, $\boldsymbol{a}$, and $\boldsymbol{A}$ to denote scalars, vectors, and matrices, respectively.

2. Model

Consider a time series of realized covariance matrices $\boldsymbol{RC}_t \in \mathbb{R}^{k\times k}$ for k assets, observed for $t = 1, \ldots, T$, and computed, for instance, using 5-minute intraday returns based on high-frequency data. We assume that $\boldsymbol{RC}_t$ has full rank for every t and admits an eigenvalue-eigenvector decomposition

$$\boldsymbol{RC}_t = \boldsymbol{Q}_t \boldsymbol{\Lambda}_t \boldsymbol{Q}_t^\top, \tag{1}$$

where $\boldsymbol{\Lambda}_t \in \mathbb{R}^{k\times k}$ is a diagonal matrix containing the ordered eigenvalues, $\lambda_{1,t} > \ldots > \lambda_{k,t} > 0$, and $\boldsymbol{Q}_t \in \mathbb{R}^{k\times k}$ is an orthonormal matrix containing the corresponding eigenvectors. Note that the realized covariance matrix $\boldsymbol{RC}_t$ is observed and that, therefore, also $\boldsymbol{\Lambda}_t$ and $\boldsymbol{Q}_t$ are observed for every $t = 1, \ldots, T$.

As motivated later on, in this paper we focus on modeling the dynamics of the eigenvalues. Let $\boldsymbol{\lambda}_t$ denote the diagonal of $\boldsymbol{\Lambda}_t$. We assume that the conditional density of $\boldsymbol{\lambda}_t$ is given by

$$\begin{aligned} g\left(\boldsymbol{\lambda}_t \mid \boldsymbol{\mu}_t, \mathcal{F}_{t-1}; \boldsymbol{\theta}^c, \boldsymbol{\nu}\right) &= \left(\prod_{i=1}^{k} g_i\left(\lambda_{i,t} \mid \mu_{i,t}, \mathcal{F}_{t-1}; \boldsymbol{\nu}_i\right)\right) \times c\left(\boldsymbol{u}_t; \boldsymbol{\theta}^c\right), \\ \boldsymbol{u}_t = (u_{1,t}, \ldots, u_{k,t}), &\qquad u_{i,t} = G_i\left(\lambda_{i,t} \mid \mu_{i,t}, \mathcal{F}_{t-1}, \boldsymbol{\nu}_i\right), \end{aligned} \tag{2}$$

where $g_i\left(\lambda_{i,t} \mid \mu_{i,t}, \mathcal{F}_{t-1}; \boldsymbol{\nu}_i\right)$ is the pdf of a scaled conditional F distribution with HAR dynamics for its mean $\mu_{i,t}$, such that

$$g_i\left(\lambda_{i,t} \mid \mu_{i,t}, \mathcal{F}_{t-1}, \boldsymbol{\nu}_i\right) = \frac{\Gamma((\nu_{1,i}+\nu_{2,i})/2)}{\Gamma(\nu_{1,i}/2)\Gamma(\nu_{2,i}/2)} \times \frac{\lambda_{i,t}^{(\nu_{1,i}-2)/2}\left(\dfrac{\nu_{1,i}}{(\nu_{2,i}-2)\,\mu_{i,t}}\right)^{\nu_{1,i}/2}}{\left(1+\dfrac{\nu_{1,i}}{\nu_{2,i}-2}\,\dfrac{\lambda_{i,t}}{\mu_{i,t}}\right)^{(\nu_{1,i}+\nu_{2,i})/2}}, \tag{3}$$

$$\mu_{i,t+1} = \omega_i + \alpha_i^d \cdot \lambda_{i,t} + \alpha_i^w \cdot \bar{\lambda}_{i,t}^w + \alpha_i^m \cdot \bar{\lambda}_{i,t}^m + \beta_i \cdot \mu_{i,t}, \tag{4}$$

$$\bar{\lambda}_{i,t}^w = P_5(L)\lambda_{i,t}, \qquad \bar{\lambda}_{i,t}^m = P_{22}(L)\lambda_{i,t}, \qquad P_n(L) = n^{-1}\sum_{i=1}^{n} L^i,$$

where L denotes the lag operator $L\,\lambda_{i,t} = \lambda_{i,t-1}$, $\mathcal{F}_t = \{\boldsymbol{\lambda}_s\}_{s\leq t}$ is a filtration, $G_i\left(\lambda_{i,t} \mid \mu_{i,t}, \mathcal{F}_{t-1}; \boldsymbol{\nu}_i\right)$ denotes the cdf of the F distribution, $\boldsymbol{\nu}_i = (\nu_{1,i}, \nu_{2,i})^\top$ holds the two degrees of freedom parameters for the F distribution, $\boldsymbol{\nu}^\top = (\boldsymbol{\nu}_1^\top, \ldots, \boldsymbol{\nu}_k^\top)$ gathers all the degrees of freedom parameters into a single vector, and $\boldsymbol{u}_t$ holds the Probability Integral Transforms (PITs) of $\lambda_{i,t}$. Finally, $c(\,\cdot\,; \boldsymbol{\theta}^c)$ is a copula density, characterized by the parameter vector $\boldsymbol{\theta}^c$.

The model in Eq. (2)–(4) is very flexible and specifies the dynamics of each $\lambda_{i,t}$ by its marginal conditional distribution, and their dependence by a copula. It also makes the model numerically efficient. In particular, each of the marginal models is specified using a conditional F distribution (3) with a scale parameter $\mu_{i,t}$ as in a Multiplicative Error Model (MEM) of Engle (2002b). By using the F distribution, the estimation becomes more robust to outliers and incidental large values in $\boldsymbol{RC}_t$ and, thus, $\boldsymbol{\lambda}_t$ (see also, e.g., Opschoor et al., 2018; Gribisch and Stollenwerk, 2020; Stollenwerk, 2026; Opschoor et al., 2025, for evidence on the importance of using such fat-tailed distributions for realized measures). The scale parameter $\mu_{i,t}$ follows the HAR dynamics of Corsi (2009), using daily ($\lambda_{i,t}$), weekly ($\bar{\lambda}_{i,t}^w$), and monthly ($\bar{\lambda}_{i,t}^m$) averages of past observed realized eigenvalues. This ensures that the model can capture the typical hyperbolic decay of the realized variance autocorrelations observed for empirical data. Given these HAR dynamics for the realized eigenvalues $\lambda_{i,t}$, we refer to the model from now on as the Realized λ-HAR (ReλHAR) model.

The marginals are linked together in Eq. (2) using the copula density $c(\boldsymbol{u}_t; \boldsymbol{\theta}^c)$. Such dependence further enhances the stochastic structure of the underlying realized covariance matrix $\boldsymbol{RC}_t$. Different choices are possible here, including the Gaussian, Student's t, Archimedean, or other copula specifications. We use a simple t copula in our empirical application, such that $\boldsymbol{\theta}^c$ contains the strict lower triangular half of a correlation matrix and shape parameter ν. In the empirical application in Section 5 we show that there exists positive dependence between the eigenvalues.

Estimation of the copula parameter $\boldsymbol{\theta}^c$ can be done jointly with the marginal parameters, or using a two-step approach. The latter is typically numerically more efficient, and typically comes at a limited cost in terms of statistical efficiency if the sample size T is sufficiently large; see for instance Joe and Xu (2005) and Patton (2006). For two-stage estimation, one first estimates the marginal parameters $\boldsymbol{\theta}_i = (\nu_{1,i}, \nu_{2,i}, \omega_i, \alpha_i^d, \alpha_i^w, \alpha_i^m, \beta_i)^\top$ using the specification in (3)–(4) and a standard maximum likelihood procedure. In a second stage, one computes the PITs $\boldsymbol{u}_t$ based on the first-stage estimates, and estimates the copula parameters using these PITs and the MLE. Given its numerical efficiency, this is also the approach we follow in this paper.

The model above has several advantages compared to earlier proposals in the literature. By concentrating on the eigenvalues, the model focuses on the most dominant part of the time variation in $\boldsymbol{RC}_t$ compared to approaches that model the full realized covariance matrix dynamics in one step, such as the conditional Wishart model of Gribisch and Stollenwerk (2020), Noureldin et al. (2012), and Gouriéroux et al. (2009), the matrix-F model of Opschoor et al. (2018), or the Riesz, t-Riesz, and Riesz-F models of Stollenwerk (2026) and Opschoor et al. (2025); compare also the arguments in Hetland et al. (2023). Unlike the λ-GARCH model of Hetland et al. (2023), we model the realized eigenvalues directly, rather than inferring the eigenvalue dynamics from daily returns. Also, unlike the Wishart model of Gribisch and Stollenwerk (2020), Noureldin et al. (2012), and Gouriéroux et al. (2009), the F distributions in (3) directly allow for fat-tails and influential observations in the observed $\lambda_{i,t}$, a feature that is empirically important following the arguments in Opschoor et al. (2018). By allowing for different tail parameters $\boldsymbol{\nu}_i$ for each marginal distribution, the model in (2)–(4) is also much more flexible than the matrix F model for $\boldsymbol{RC}_t$ of Opschoor et al. (2018), which uses a single ν_1 and ν_2 to capture the fat-tailedness of all assets. This flexibility is empirically relevant based on the tail heterogeneity results of Näf et al. (2019), Stollenwerk (2026), Hediger et al. (2023), and Opschoor et al. (2025). However, unlike the Riesz related distributions used in some of these papers, the new model above (i) does *not* depend on the (arbitrary) ordering of the variables in the system, and (ii) has additional flexibility via the choice of the copula density function in (2).

Note that the specification for the dynamics of $\mu_{i,t}$ can easily be generalized even further. One can, for instance, include explanatory variables in (4), or use the score-driven approach of Creal et al. (2013) to drive the dynamics of $\mu_{i,t}$. In our empirical application on (conditional) *VolaR* forecasting in Section 5, this did not produce any significant improvements, which is why we restrict ourselves in the current paper to the simpler specification in (4).

As a final step in modeling $\boldsymbol{RC}_t$, we consider the eigenvectors $\boldsymbol{Q}_t$ in (1). Given the orthonormality of $\boldsymbol{Q}_t$, the size of these vectors remains constant at one. The main variation in the size of the elements of $\boldsymbol{RC}_t$ is therefore captured by the dynamics of the eigenvalues $\boldsymbol{\lambda}_t$ rather than the elements in $\boldsymbol{Q}_t$. We therefore propose a simple recursive forecasting procedure for $\boldsymbol{Q}_t$ using the matrix $\hat{\boldsymbol{Q}}_t$, which holds the eigenvectors of the recursive mean of the realized covariances, $(t-1)^{-1}\sum_{s=1}^{t-1}\boldsymbol{RC}_s$ for $t>1$, with initial value $\hat{\boldsymbol{Q}}_1 = \boldsymbol{Q}_1$. Alternatives based on rolling rather than recursive estimators can of course also be used, as well as random walk forecasts of $\boldsymbol{Q}_t$. Preliminary analyses for our empirical application show that such variations matter little, as long as some stability in the forecasts of $\boldsymbol{Q}_t$ is guaranteed, e.g., using the ones based on the simple recursive means above.

A similar recursive recipe does not suffice for the eigenvalues $\lambda_{i,t}$, however, as these vary much more and have a much larger impact on the prediction of $\boldsymbol{RC}_t$. Algorithm 1 summarizes the full model and its estimation.

3. Estimating the marginal models

Algorithm 1 describes the full estimation algorithm of the ReλHAR model. The estimates of the parameters $\boldsymbol{\theta}_i = (\nu_{1,i}, \nu_{2,i}, \omega_i, \alpha_i^d, \alpha_i^w, \alpha_i^m, \beta_i)^\top$ of the ith marginal distribution are

Algorithm 1 Estimation of the ReλHAR model with t copula

$\boldsymbol{RC}_t = \boldsymbol{Q}_t \boldsymbol{\Lambda}_t \boldsymbol{Q}_t^\top$ for $t = 1, \ldots, T$, with $\boldsymbol{\Lambda}_t = \operatorname{diag}(\boldsymbol{\lambda}_t)$ and orthonormal $\boldsymbol{Q}_t$
$\hat{\boldsymbol{Q}}_t \leftarrow$ the eigenvectors of $(t-1)^{-1} \sum_{s=1}^{t-1} \boldsymbol{RC}_s$ for $t = 2, \ldots, T$
for $i = 1, \ldots, k$ **do**
    $\boldsymbol{\theta}_i \leftarrow (\nu_{1,i}, \nu_{2,i}, \omega_i, \alpha_i^d, \alpha_i^w, \alpha_i^m, \beta_i)^\top$
    using (3)–(4), compute $\hat{\boldsymbol{\theta}}_i$ by maximizing the marginal log likelihood in (5)
    with $\hat{\boldsymbol{\theta}}_i$, compute $\hat{\mu}_{i,t}$ using (4)
    $u_{i,t} \leftarrow G_i\left(\lambda_{i,t} \mid \hat{\mu}_{i,t}, \mathcal{F}_{t-1}; \hat{\boldsymbol{\nu}}_i\right)$
end for
$\boldsymbol{z}_t = \left(t^{-1}(u_{1,t}, \nu), \ldots, t^{-1}(u_{k,t}, \nu)\right)^\top$ where $t^{-1}(\,\cdot\,, \nu)$ is the inverse standard t cdf
estimate $\boldsymbol{\theta}^c$ on $\boldsymbol{z}_t$ for $t = 1, \ldots, T$ using, for example, MLE where the correlation matrix R is estimated by the sample correlation matrix.

obtained by maximizing the log-likelihood function

$$\hat{\boldsymbol{\theta}}_i = \arg\max_{\boldsymbol{\theta}_i \in \Theta} \sum_{t=1}^{T} \log g_i\left(\lambda_{i,t} \mid \mu_{i,t}, \mathcal{F}_{t-1}; \boldsymbol{\nu}_i\right), \tag{5}$$

for $i = 1, \ldots, k$, where Θ denotes the common parameter space for each $\boldsymbol{\theta}_i$. In a second step, we estimate the copula parameters using the probability integral transforms obtained from estimating the first-stage (marginal) models. For this approach, the generic two-stage copula argument is developed by Joe and Xu (2005), while Patton (2006) provides the corresponding arguments for conditional time-series models. Related dynamic copula results are given by Chen et al. (2006). The first-stage estimators require a bit more attention given the use of the conditional F-distribution for the realized eigenvalues and the dynamics of $\mu_{i,t}$ in (4). We therefore treat the copula step as standard under the usual copula regularity conditions and only investigate the asymptotic and finite-sample properties for the first-stage marginal estimators.

We first establish the consistency and asymptotic normality of the first-stage ML Estimators (MLEs) of $\boldsymbol{\theta}_i$ for the marginal models and for $i = 1, \ldots, k$. We obtain the following result.

***Theorem* 1** (Consistency and asymptotic normality of the first-stage MLE)**.** *Let the data be generated by (2)–(4) for some $\theta_{0,i} \in \operatorname{int}(\Theta)$, $i = 1, \ldots, k$, where Θ is compact and coincides with the closure of its open interior. For every $\theta_i \in \Theta$, all elements of $\theta_i = (\nu_{1,i}, \nu_{2,i}, \omega_i, \alpha_i^d, \alpha_i^w, \alpha_i^m, \beta_i)^\top$ are strictly positive, $\nu_{2,i} > 4$, and $\alpha_i^d + \alpha_i^w + \alpha_i^m + \beta_i < 1$. Then, for every $i = 1, \ldots, k$, $\hat{\theta}_i \xrightarrow{a.s.} \theta_{0,i}$.*

Moreover, let $\ell_{i,t}(\theta_i) := \log g_i\big(\lambda_{i,t} \mid \mu_{i,t}(\theta_i), \mathcal{F}_{t-1}; \boldsymbol{\nu}_i\big)$ denote the marginal log-likelihood *contribution in (5), and define the score $s_{i,t}(\theta_i) := \partial \ell_{i,t}(\theta_i)/\partial\theta_i$. Define*

$$V_i(\theta_i) := \mathbb{E}\Big[s_{i,t}(\theta_i)\, s_{i,t}(\theta_i)^\top\Big], \qquad H_i(\theta_i) := -\,\mathbb{E}\left[\frac{\partial^2 \ell_{i,t}(\theta_i)}{\partial\theta_i\,\partial\theta_i^\top}\right]. \tag{6}$$

If $H_i(\theta_i)$ is non-singular in an ε-neighborhood of $\theta_{0,i}$, then

$$\sqrt{T}\,(\hat{\theta}_i - \theta_{0,i}) \ \xrightarrow{d}\ \mathcal{N}\Big(0,\ H_i(\theta_{0,i})^{-1}\, V_i(\theta_{0,i})\, H_i(\theta_{0,i})^{-1}\Big).$$

The proof provided in Appendix A directly follows the literature (see, e.g., Bollerslev, 1986; Berkes and Horváth, 2004; Straumann and Mikosch, 2006). In particular, we impose the existence of an unconditional first and second moment of $\mu_{i,t}(\boldsymbol{\theta}_i)$ via the condition $\nu_{2,i} > 4$ and the familiar condition $\alpha_i^d + \alpha_i^w + \alpha_i^m + \beta_i < 1$. The existence of a stationary and ergodic data generating process (dgp) and filter invertibility then follow along the same lines as in the GARCH literature (see, e.g., Francq and Zakoian, 2019). The rest of the proof then uses the specific form of the conditional F distribution in (3). As a result, not only the static parameters converge to their true values, but also the filtered paths of $\mu_{i,t}$ converge to their true counterparts exponentially fast almost surely (Straumann and Mikosch, 2006).

To see whether the first step of the estimation procedure provides a good approximation in finite samples, we perform a Monte-Carlo experiment (MC) based on 50,000 replications. We simulate data from model (3)–(4) with parameter values $\boldsymbol{\theta} = (\alpha_d, \alpha_w, \alpha_m, \beta, \nu_1, \nu_2) = (0.12, 0.1, 0.05, 0.7, 10, 10)$ and subsequently estimate the parameters using Maximum Likelihood. The parameter values are chosen in line with empirical observations. Table 1 shows the results for various sample sizes $T = 1000, 2000, 5000$, in line with typical sample sizes in the empirical application. For each parameter, we report the true value and the Monte-Carlo mean of its ML estimate. We also provide the Monte-Carlo standard deviation of the estimates (MC St.dev.), and the Monte-Carlo mean of the standard errors based on Theorem 1 (Mean SE).

Table 1 shows that the ML estimator is sufficiently accurate. The Bias and MSE become smaller as T increases and are generally small. We also note that the Mean SE and MC St.dev. are close, implying that the asymptotic result in Theorem 1 provides a good approximation in finite samples. The latter is also confirmed by the kernel density plots of the t-statistics as shown in Figure 1. In most cases, the normality assumption provides a reasonable approximation to the empirical distribution, particularly if the sample size grows large. The densities of the t-test statistics for α_d and β seem to take somewhat longer to converge, as the estimator has difficulty distinguishing the long memory features caused by α_d from those caused by β. As a result, one of the distributions is slightly skewed to the left, whereas the other is slightly right-skewed.

Table 1: Finite Sample Estimation Properties.

This table shows the simulation results for the Monte-Carlo experiment using Model (3)–(4) for the parameter set $\boldsymbol{\theta}$. Panels A, B, and C show the results for T = 1000, 2000, 5000, respectively. Column 'True' shows the true parameter values, Column 'Mean' shows the mean of the MLE across the MC experiment, Column 'Mean SE' shows the mean of the ML standard errors across the MC experiment, Column 'MC Std.dev' shows the MC standard deviation of the ML estimates, and the last two columns report the Bias and MSE, respectively. The results are based on 50000 MC repetitions.

	True	Mean	Mean SE	MC Std.dev.	Bias	MSE
Panel A: T = 1000						
α_d	0.12	0.111	0.039	0.039	-0.009	0.002
α_w	0.10	0.122	0.088	0.096	0.022	0.010
α_m	0.05	0.063	0.056	0.066	0.013	0.005
β	0.70	0.667	0.122	0.141	-0.033	0.021
ν_1	10	10.284	1.419	1.471	0.284	2.243
ν_2	10	10.176	1.316	1.362	0.176	1.886
Panel B: T = 2000						
α_d	0.12	0.115	0.029	0.029	-0.005	0.001
α_w	0.10	0.113	0.071	0.074	0.013	0.006
α_m	0.05	0.057	0.046	0.052	0.007	0.003
β	0.70	0.681	0.102	0.111	-0.019	0.013
ν_1	10	10.137	0.957	0.969	0.137	0.957
ν_2	10	10.079	0.894	0.905	0.079	0.825
Panel C: T = 5000						
α_d	0.12	0.118	0.019	0.019	-0.002	0.000
α_w	0.10	0.104	0.050	0.049	0.004	0.002
α_m	0.05	0.052	0.033	0.033	0.002	0.001
β	0.70	0.694	0.073	0.073	-0.006	0.005
ν_1	10	10.048	0.589	0.593	0.048	0.355
ν_2	10	10.031	0.554	0.593	0.031	0.352

Given the overall adequate behavior of the model and the estimation procedure, we feel comfortable in applying the model to empirical data. In particular, we investigate whether the model's performance is also good for out-of-sample forecasting. We do so in Sections 4 and 5.

4. Empirical Application Set-up

In our empirical application, we use the ReλHAR model from Section 2 to predict portfolio Volatility-at-Risk (VolaR), i.e., quantiles of the realized volatility distribution. In addition, we also predict portfolio conditional VolaR, i.e., the average realized portfolio volatility beyond these quantiles (comparable to the Expected Shortfall measure for returns). We first introduce both concepts, the data used, and the benchmark models included in our

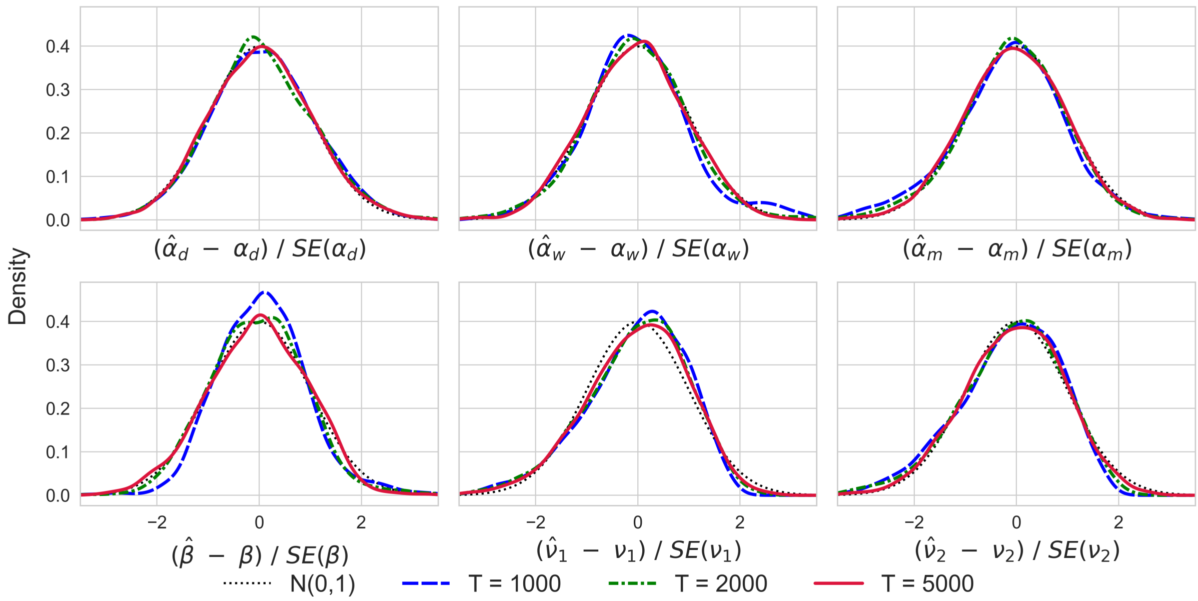

Figure 1: Finite Sample Parameter Distribution.

This figure shows Kernel Density Estimates (KDE) of studentized estimation errors based on Theorem 1 for each of the parameters using a Monte-Carlo experiment. The panels show the KDE for α_d, α_w, α_m, β, ν_1, ν_2. In each subfigure, the KDEs are shown for T = 1000, 2000, 5000 in blue, green, and red, respectively.

analysis. We then present the empirical results in Section 5.

4.1. (Conditional) Volatility-at-Risk

VolaR (Caporin et al., 2017b) assesses the risk of high levels of volatility of a given portfolio. It has important implications for financial risk management and particularly during periods of financial turmoil (Huang et al., 2019). During such periods the correlation between asset returns tends to increase, causing assets to cluster and diversification benefits to fade, possibly increasing the variance of a portfolio more than expected. *VolaR* measures such risks and is computed as an upper quantile of the portfolio variance forecast density. In particular, the $q\%$ *VolaR* is defined as

$$\mathbb{P}\left(rc_{t+1} > VolaR^q_{t+1} \mid \mathcal{F}_t\right) = q, \tag{7}$$

where rc_{t+1} denotes the realized portfolio variance, computed as

$$rc_t = \boldsymbol{w}^\top \boldsymbol{RC}_t\, \boldsymbol{w}, \tag{8}$$

where $\boldsymbol{w} \in \mathbb{R}^{k\times 1}$ is a vector of portfolio weights. In most of our analysis, we use an equally weighted portfolio $\boldsymbol{w} = \iota_k/k$, where ι_k denotes a $k \times 1$ vector of ones, but results are robust to other choices of $\boldsymbol{w}$, see Appendix B.1. The *VolaR* thus measures the uncertainty in the risk of a portfolio in a similar way as the Value-at-Risk measures the uncertainty of the

Algorithm 2 Simulation of the time $T+1$ Portfolio Variance Distribution

$\hat{\boldsymbol{\theta}}_i \leftarrow$ from Algorithm 1 $\forall\ i = 1, \dots, k$
$\hat{\boldsymbol{Q}}_{T+1} \leftarrow$ the eigenvectors of $(1/T)\sum_{t=1}^{T} \boldsymbol{RC}_t$
$\hat{\boldsymbol{R}}$ from $\hat{\boldsymbol{\theta}}^c$ and Algorithm 1
for $m = 1, \dots, M$ **do**
  simulate k-dimensional vector $\tilde{\boldsymbol{z}}_{T+1,m} = (\tilde{z}_1, \dots, \tilde{z}_k)^\top_{T+1,m}$ from $t(\mathbf{0}_k, \hat{\boldsymbol{R}}, \nu)$
  for $i = 1, \dots, k$ **do**
    compute $\tilde{\lambda}_i$ for $\tilde{z}_i$ using (10)
  end for
  $\hat{\boldsymbol{\Lambda}}_{T+1,m} \leftarrow$ a diagonal matrix containing $(\tilde{\lambda}_1, \dots, \tilde{\lambda}_k)^\top_{T+1,m}$
  $\widetilde{\boldsymbol{RC}}_{T+1,m} = \hat{\boldsymbol{Q}}_{T+1}\hat{\boldsymbol{\Lambda}}_{T+1,m}\hat{\boldsymbol{Q}}^\top_{T+1}$
  $\tilde{rc}_{T+1,m} = \mathbf{w}^\top \widetilde{\boldsymbol{RC}}_{T+1,m}\mathbf{w}$, with $\mathbf{w}$ a vector that sums up to unity.
end for

portfolio returns themselves. More recently, Expected Shortfall (ES) has taken a more prominent position in regulation (Tasche, 2002) as a measure of portfolio risk. In our context, ES (conditional *VolaR*) reflects the expected magnitude of extreme volatility, not just the threshold exceeded. This ensures better sensitivity to tail risks in volatility, which is crucial for robust risk management in volatile markets. We define the $q\%$ conditional *VolaR* as

$$\mathbb{E}\left(rc_{t+1} \mid rc_{t+1} > VolaR^q_{t+1}, \mathcal{F}_t\right). \tag{9}$$

We assess the tail shape of the portfolio variance distribution at q levels of 5%, 2.5% and 1%, respectively, and compute 1-step ahead forecasts of the (conditional) *VolaR* by using simulation techniques. The univariate distribution of the portfolio variance σ^2_{pf} is directly obtained from the (multivariate) covariance matrix forecast distribution, which, in turn, can be constructed from the univariate eigenvalue distributions and the copula. Each of the eigenvalue processes can be simulated using the conditional F distribution. We first simulate a copula draw $\tilde{u}_{i,t}$ from the estimated copula density $c(\boldsymbol{u}_t; \boldsymbol{\theta}^c)$. Next, we compute the simulated eigenvalues as

$$\tilde{\lambda}_{t+1,i} = \hat{\mu}_{t+1,i} \cdot \frac{\hat{\nu}_{2,i} - 2}{\hat{\nu}_{2,i}} \cdot F^{-1}_{\mathrm{CDF}}\left(\tilde{u}_{i,t}; \hat{\nu}_{1,i}, \hat{\nu}_{2,i}\right), \tag{10}$$

where $F^{-1}_{\mathrm{CDF}}(\,\cdot\,; \nu_1, \nu_2)$ denotes the inverse CDF of the standard F distribution with ν_1 and ν_2 degrees of freedom.

Given the k simulated eigenvalues, the simulated realized covariance matrix can be constructed using the spectral decomposition, after which the draw of the realized portfolio variance follows immediately via Eq. (8). We summarize the entire procedure in the following Algorithm 2.

From this we obtain the simulated distribution of the portfolio variances, which allows us to compute the 1-step ahead (conditional) *VolaR* for the relevant quantile levels.

4.2. Backtesting Procedures

To assess the accuracy of the portfolio variance forecast distributions, we backtest the *VolaR* forecasts by comparing them with the realized portfolio variances rc_{t+1}. Define the hit sequence $h_t(q)$ as

$$h_t(q) = \mathbb{1}\{rc_t \geq VolaR_t^q\}. \tag{11}$$

From the definition of the *VolaR* we expect rc_{t+1} to exceed $VolaR_{t+1}^q$ around $q \cdot T$ times. Therefore, we can use the unconditional and conditional coverage (UC and CC) tests of Kupiec (1995) and Christoffersen (1998) to backtest the adequacy of the *VolaR* across different models.

We also compare relative performance between models using the tick-loss measure (Giacomini and White (2006)). This statistic is given by

$$L_{t+1}^q = \left(q - \mathbb{1}\{\, VolaR_{t+1}^q < rc_{t+1}\}\right) \cdot \left(VolaR_{t+1}^q - rc_{t+1}\right). \tag{12}$$

This is the appropriate loss function to identify the *VolaR* as the $1-q$-quantile of the forecast distribution of rc_{t+1}. We can compare the performance of different models in terms of their Tick-loss value using the Model Confidence Set (MCS) of Hansen et al. (2011).

To accommodate the conditional VolaR, Du and Escanciano (2017) propose the Unconditional- (UES) and Conditional (CES) Expected Shortfall tests. The UES and CES provide tests for the ES forecast using the cumulative violation process. Consider the integral

$$H_t(q) = q^{-1} \cdot \int_0^q h_t(u)du. \tag{13}$$

The UES test then naturally tests $\mathbb{E}[H_t(q)] = q/2$. For forecasting window size n, a test statistic of the UES test is then given by

$$U_{ES} = \frac{\sqrt{n}(\bar{H} - q/2)}{\sqrt{q \cdot \left(\frac{1}{3} - \frac{q}{4}\right)}}. \tag{14}$$

We use the standard normal distribution to assess significance given the asymptotic results in Du and Escanciano (2017) and our use of a large estimation sample. Furthermore, defining the lag-j autocovariance as

$$\gamma_{nj} = \frac{1}{n-j} \sum_{t=1+j}^{n} (H_t(q) - q/2)(H_{t-j}(q) - q/2), \tag{15}$$

a CES test can be performed using the Box-Pierce statistic

$$C_{ES}(m) = n \sum_{i=1}^{m} \hat{\rho}_{nj}^2, \quad \hat{\rho}_{nj} = \frac{\hat{\gamma}_{nj}}{\hat{\gamma}_{n0}}, \tag{16}$$

where we take $m = 10$. These two Expected Shortfall tests further assess the forecasting quality of the different models over the entire extreme tail of the conditional portfolio variance distribution.

4.3. Data

We apply the ReλHAR model to daily realized covariance matrices of 5 large US financial stocks: AXP, JPM, AIG, BAC, C. Picking large blue chip stocks ensures that the variances are (roughly speaking) similarly sized, and that the eigenvalue-eigenvector decomposition we use is not dominated by a single security. The day t realized covariance matrix is obtained from high-frequency data using 5-minute intraday returns. We follow the guidelines of Barndorff-Nielsen et al. (2009), Brownlees and Gallo (2006) and Shephard and Sheppard (2010). For $k = 5$ assets, we denote the returns $\boldsymbol{r}_{t,n} \in \mathbb{R}^{k \times 1}$ on day t and 5-minute interval n, for $n = 1, \ldots, N$. We then estimate a realized covariance matrix as

$$\boldsymbol{RC}_t = \sum_{n=1}^{N} \boldsymbol{r}_{t,n} \cdot \boldsymbol{r}_{t,n}^{\top}. \tag{17}$$

The $\boldsymbol{RC}_t$ are then decomposed using Eq. (1). Table 2 shows the summary statistics[1] of the eigenvalue series of the five equities.

The complete dataset spans the period from the 2nd of January 2002 to the 22nd of February 2023, a total of 5250 trading days. The model is estimated using the first 1250 observations until the 10th of January 2007. The (conditional) *VolaR* application is then carried out using predictions spanning from 2007 until 2023, covering the financial crisis, the COVID-19 period, and beyond. This implies the model is initially estimated based on a stable financial landscape: it is therefore challenging for the model to adapt itself to the heat of the subsequent financial crisis, and it is interesting to see whether the

[1] As an intermediate step, we bound the influence of AIG in the scale of the realized covariance matrices. We do so by capping the realized variances in the variance-correlation decomposition at 1000, which is exceeded 3 times during the 2008 financial crisis. In this way, we ensure positive definite $\boldsymbol{RC}_t$ and preserve the correlation structure. This step has no notable influence on the empirical application of the ReλHAR model and its benchmarks.

Table 2: Summary Statistics of the Eigenvalues.

This table shows the summary statistics of the daily realized eigenvalues of 5 financial U.S. equities; AXP, JPM, AIG, BAC, C. The sample runs from the 2nd of January 2002 to the 22nd of February 2023 (5250 trading days).

	Mean	St.dev.	Min	Max	Skew.	Kurt.
$\lambda_{1,t}$	16.068	56.293	0.415	1149.403	11.390	174.687
$\lambda_{2,t}$	4.131	15.538	0.185	320.334	10.687	148.813
$\lambda_{3,t}$	2.130	6.833	0.131	176.593	11.351	188.936
$\lambda_{4,t}$	1.331	3.641	0.093	88.196	8.931	123.496
$\lambda_{5,t}$	0.861	2.316	0.035	67.728	10.613	195.398

model succeeds in accurately forecasting high levels of volatility even if such levels were not observed during the estimation period. The model is re-estimated every 20 days (or approximately every month). We also perform a follow-up analysis that focuses on the periods of financial turmoil by testing forecasting performance during the sub-periods of the 2008 financial crisis and the COVID-19 crisis. This additional test comprises 1750 observations in total and spans from January 2007 until January 2013, and from December 2019 until December 2020.

4.4. Benchmark Models

We compare the new model to several models from the literature. Our first benchmark is a restricted version of the Matrix-F model of Opschoor et al. (2018), which we subsequently label as 'Matrix-F'. It assumes that $\boldsymbol{RC}_t$ is conditionally Matrix-F distributed with mean $\boldsymbol{V}_t$ and scalar tail parameters ν_1 and ν_2. The conditional mean matrix $\boldsymbol{V}_t$ is modeled using the matrix generalization of the HAR model of Corsi (2009) as

$$\boldsymbol{V}_{t+1} = (1-\alpha_d-\alpha_w-\alpha_m-\beta)\cdot\boldsymbol{\Omega}+\alpha_d\cdot\boldsymbol{RC}_t+\alpha_w\cdot\overline{\boldsymbol{RC}}_{w,t}+\alpha_m\cdot\overline{\boldsymbol{RC}}_{m,t}+\beta\cdot\boldsymbol{V}_t, \quad (18)$$

where $\boldsymbol{\Omega}$ is estimated by the unconditional mean of $\boldsymbol{RC}_t$, and $\overline{\boldsymbol{RC}}_{w,t}$ and $\overline{\boldsymbol{RC}}_{m,t}$ are averages of the observed realized covariance matrices over the past week and month, respectively. In contrast to the new ReλHAR model, the Matrix-F model describes the full realized covariance matrix and its dynamics directly, rather than through its eigenvalue-eigenvector decomposition and the copula link.

Our second benchmark is the adapted DPC-CAW model of Gribisch and Stollenwerk (2020). It assumes a conditional Wishart distribution for $\boldsymbol{RC}_t$ with a time-varying conditional mean $\boldsymbol{V}_t$, which is decomposed into eigenvectors and eigenvalues

$$\begin{aligned}\boldsymbol{RC}_t \mid \mathcal{F}_{t-1} &\sim \mathcal{W}(\boldsymbol{V}_t, \nu_W) \\ \boldsymbol{V}_t &= \boldsymbol{L}_t\boldsymbol{\Lambda}_t\boldsymbol{L}_t^\top,\end{aligned} \quad (19)$$

with $\boldsymbol{L}_t$ a matrix of latent eigenvectors and $\boldsymbol{\Lambda}_t$ a diagonal matrix of ascending eigenvalues. The eigenvectors $\boldsymbol{L}_t$ are modeled by a matrix-variate auxiliary process while the eigenvalues $\gamma_{i,t}$ of the diagonal matrix $\boldsymbol{\Lambda}_t$ are modeled by a GARCH type process. Note that in this model, the latent conditional eigenvalues of $\boldsymbol{V}_t$ are modeled, whereas the ReλHAR model directly focuses on the first moment of the observed realized eigenvalue.

We adapt the DPC-CAW model by including HAR type recursion as in (4). This empirical improvement to the DPC-CAW benchmark model allows us to discriminate whether the improved fit of the ReλHAR model stems from the specific distributional assumptions and decompositions that are used, rather than from any differences in short versus (quasi) long memory dynamics. Indeed, not applying the HAR recursion within the DPC-CAW framework yields less accurate *VolaR* forecasts for the DPC-CAW compared to the new ReλHAR model. We therefore omit this specification from our comparison, providing an additional advantage to the DPC-CAW benchmark model.

The third benchmark model concentrates on the effect of the decomposition and uses the variance-correlation decomposition of the covariance matrix (see also Vassallo et al., 2021)

$$\boldsymbol{RC}_t = \boldsymbol{D}_t^{1/2} \boldsymbol{P}_t \boldsymbol{D}_t^{1/2}, \tag{20}$$

rather than the eigenvalue-eigenvector decomposition, where $\boldsymbol{D}_t$ is a diagonal matrix containing the asset specific realized variances $(rv_{1,t}, \ldots, rv_{k,t})$, and $\boldsymbol{P}_t$ denotes the realized correlation matrix corresponding to $\boldsymbol{RC}_t$. Similar to the DCC model in traditional covariance matrix modeling (Engle, 2002a), the DCC-CAF benchmark model separately models the variance and correlation components in (20), similar to the DCC-HEAVY and DECO-HEAVY models of Bauwens et al. (2012) and Bauwens and Xu (2023). The asset realized variances $(rv_{1,t}, \ldots, rv_{k,t})$ are assumed to be conditionally independent and F distributed, each element having potentially different degrees of freedom parameters. Unlike the original specification, we endow the models for the variances with univariate HAR dynamics. This improves the benchmark's empirical fit, and provides an even stronger competitor for the ReλHAR model than the original benchmark model specification. Similar to the estimation of the marginal models for the eigenvalues in Algorithm 1, the models for the individual volatilities are estimated using MLE and yield estimates of the matrices $\hat{\boldsymbol{D}}_t$. The correlation matrix dynamics are subsequently estimated in a second stage by maximizing the likelihood of the matrix-valued model conditional on $\hat{\boldsymbol{D}}_t$:

$$\begin{aligned} \boldsymbol{RC}_t &\sim \text{Matrix-}F\left(\hat{\boldsymbol{D}}_t^{1/2} \boldsymbol{R}_t \hat{\boldsymbol{D}}_t^{1/2}, \nu_1, \nu_2\right), \\ \boldsymbol{R}_{t+1} &= (1 - \alpha_d - \alpha_w - \alpha_m - \beta)\boldsymbol{\Omega} + \beta\, \boldsymbol{R}_t + \alpha_d\, \boldsymbol{P}_t + \alpha_w\, \overline{\boldsymbol{P}}_{w,t} + \alpha_m\, \overline{\boldsymbol{P}}_{m,t}. \end{aligned} \tag{21}$$

The DCC-CAF model thus differs from the ReλHAR model in the decomposition method that is used, but uses a similar two-stage estimation approach and similar assumptions for the dynamics (HAR) and tails ((matrix)-F). We again stress that these additions provide an even stronger benchmark for the new ReλHAR model than some of the thin-tailed, short-memory specifications from the literature. Alternatively, they could even be seen as new, improved benchmark models.

In addition to the three matrix-valued benchmark models, we also consider univariate benchmark models to predict the *VolaR*. This can be done by directly modeling the realized volatility rc_t of (8). To remain congruent with the multivariate models, we assume rc_t has a conditional F distribution, similar to the dynamics for a single eigenvalue λ_t with a time-varying conditional mean with HAR dynamics (i.e. the univariate version of (18)). Directly concentrating on the univariate realized portfolio measure rc_t may help to avoid noise and parameter uncertainty that can be the result of a matrix modeling perspective based on the full $\boldsymbol{RC}_t$. The mapping to a univariate construct collapses all the variance and correlation dynamics into a single measure, which can lead to parsimony and efficiency gains. The flip side of this, of course, is that we can no longer distinguish between the different drivers of increases and decreases in the univariate realized variance measure rc_t.

5. Empirical Results

We first analyze the ReλHAR model in-sample based on the full dataset for five large financial stocks. Next, we investigate the forecasting performance of the ReλHAR model in an extensive conditional Volatility-at-Risk application, comparing it to the various benchmark models of Section 4.4.

5.1. Full-Sample Results

We estimate the ReλHAR model using the 2-step estimation procedure for the marginals and the copula, targeting the intercepts ω as $(1 - \alpha_d - \alpha_w - \alpha_m - \beta)\hat{\omega}^*$, where $\hat{\omega}^*$ is the in-sample mean of the realized eigenvalues. Table 3 shows the parameter estimates. To save space, but still give an impression of the model's performance, we present the results for $k = 5$ assets using the full in-sample period of $T = 5250$ daily observations.

The table yields five main take-aways. First, the parameters α_d, α_w, α_m, and β, are quite comparable across the different realized eigenvalue series. Their values are in line with a typical (quasi) long-memory type pattern (Corsi, 2009). Second, the first degrees of freedom parameter for the marginal models is more difficult to pin down based on the data: the values appear to differ somewhat across the different eigenvalue series, but it is difficult to see whether these differences are significant given the considerable

Table 3: ReλHAR Full Sample Estimation Results.

This table shows the estimation results of the two-stage MLE for the full sample-period 2 Jan 2002–22 Feb 2023 ($T = 5250$) for $k = 5$ large US financial stocks: AXP, JPM, AIG, BAC, C. In Panel A, the parameter estimates of $\boldsymbol{\theta} = (\alpha_d, \alpha_w, \alpha_m, \beta, \nu_1, \nu_2)$, $i = 1, \ldots, k$, for each of the realized eigenvalue series are shown with MLE standard errors in parentheses. Panel B shows the t-copula parameter estimates. Panel C presents the loglikelihood values of the ReλHAR model for a t, Gaussian, and an independence copula specification.

	λ_1	λ_2	λ_3	λ_4	λ_5
Panel A: Marginal model parameter estimates					
α_d	0.527	0.476	0.506	0.531	0.513
	(0.016)	(0.016)	(0.015)	(0.015)	(0.016)
α_w	0.052	0.067	0.138	0.113	0.097
	(0.031)	(0.035)	(0.035)	(0.033)	(0.034)
α_m	0.157	0.183	0.145	0.134	0.136
	(0.015)	(0.017)	(0.015)	(0.014)	(0.015)
β	0.250	0.262	0.198	0.205	0.235
	(0.038)	(0.044)	(0.039)	(0.036)	(0.038)
ν_1	60.087	126.850	213.210	146.524	35.384
	(11.477)	(29.575)	(41.524)	(36.929)	(3.365)
ν_2	10.900	13.379	17.290	15.764	15.955
	(0.408)	(0.408)	(0.419)	(0.515)	(0.729)
Panel B: t-copula correlation matrix, $\hat{\nu} = 15.282$ (0.985)					
λ_2	0.599				
	(0.008)				
λ_3	0.655	0.684			
	(0.007)	(0.007)			
λ_4	0.644	0.548	0.784		
	(0.007)	(0.009)	(0.005)		
λ_5	0.601	0.472	0.685	0.821	
	(0.008)	(0.010)	(0.006)	(0.004)	
Panel C: Log-likelihood values					
	ReλHAR-t	ReλHAR-Gaussian		ReλHAR-independence	
Loglik	-7588	-7773		-16448	

standard errors. Third, there appear to be considerable differences in fat-tailedness (ν_2) between the five series. As standard errors for ν_2 are small, these differences are also likely significant. In particular, the tails appear to be fattest for the first realized eigenvalue series. Fourth, Panel B shows that there are substantial positive conditional copula correlations between the eigenvalues. This is in line with all eigenvalues following the increases/decreases in the scale of the realized covariance matrix. It is also interesting to see that the copula correlations between the largest eigenvalues (nearing 60%) are lower than the correlations between smaller eigenvalues (nearing 80%). This suggests that even

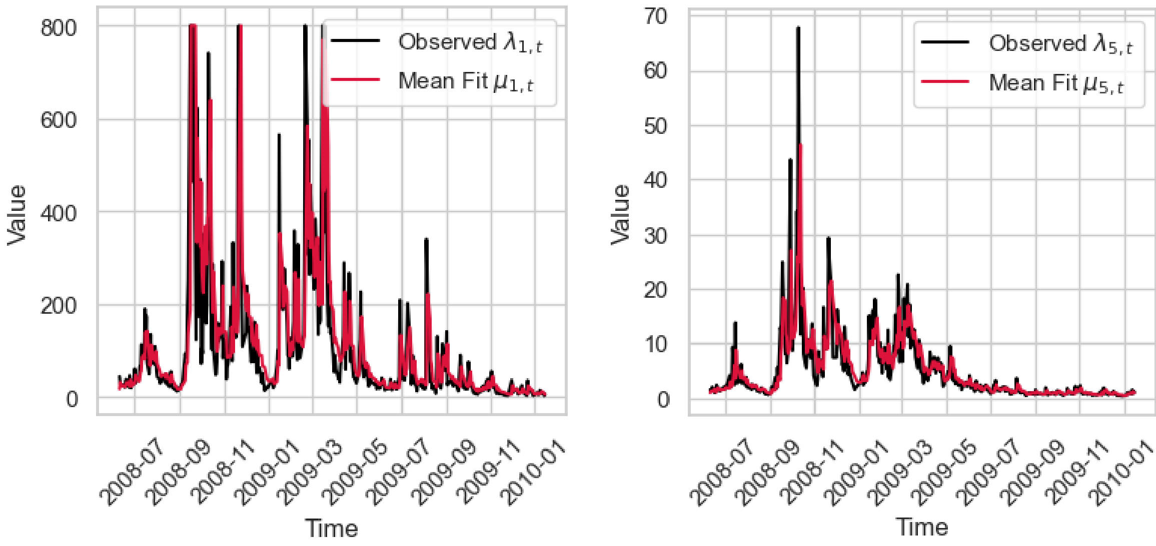

(a) largest eigenvalue $\lambda_{1,t}$ **(b)** smallest eigenvalue $\lambda_{5,t}$

Figure 2: Eigenvalue Series.

This figure shows the time series of realized eigenvalues (λ_t) and their fitted values (μ_t) derived for a setting with $k = 5$ large US financial stocks: AXP, JPM, AIG, BAC, C. Forecasts are based on the ReλHAR model estimated using MLE. The series is shown for the financial crisis sub-period, ranging from June 2007 until December 2009. Panel (a) shows the primary eigenvalue $\lambda_{1,t}$ of $\boldsymbol{RC}_t$, while Panel (b) shows the smallest eigenvalue $\lambda_{5,t}$ over time. In both panels, the black line represents the observed series $\{\lambda_{i,t}\}$, while the red line represents the forecasts $(\mu_{i,t})$. Values in Panel (a) are capped at 800 for visibility.

though there is considerable conditional co-movement between all eigenvalues, it may be important to allow for differences in these correlations to better match the patterns present in the empirical data. The log-likelihood values finally confirm the empirical evidence of conditional dependence: the maximized log-likelihood values of the ReλHAR models are much higher than that of the independence copula specification. Moreover, there is even some degree of tail-dependence as the t copula specification improves upon the Gaussian copula by roughly 200 log-likelihood points.

To illustrate the empirical behavior of the ReλHAR model, Figure 2 shows the largest $(i = 1)$ and smallest $(i = 5)$ realized eigenvalue series and their fitted values during the 2008 financial crisis. Both panels of the figure confirm that the model succeeds in following the secular dynamics of the realized (large and small) eigenvalues well, even during periods of stress.

5.2. Forecasting Performance

We now turn to the out-of-sample performance of the ReλHAR model using the (conditional) *VolaR* application introduced in Section 4 and the benchmark models listed in the same section. Table 4 shows the *VolaR* results over the out-of-sample period from January 2007 to January 2023 using an equally weighted test portfolio for all models. The table shows that the ReλHAR models outperform all of the benchmark models, all of which have been improved empirically compared to their original specification in the literature to become more competitive (see Section 4.4). The empirical exceedance probability $\hat{q}$ for the ReλHAR model is consistently very close to the nominal q. The ReλHAR model also attains the smallest Tick-loss for all levels of q and has the highest MCS p-value, such that the model (with t or independence copula) is always in the Model Confidence Set. The univariate model that directly models the dynamics of the univariate realized portfolio variance is also in the model confidence set, but with a lower p-value than the full ReλHAR model. By contrast, the MCS p-values are low for the three models that use matrix distributions like the Wishart or matrix-F to capture the realized covariance matrix dynamics. The latter three benchmark models also have too many *VolaR* violations ($\hat{q}$) and obtain a substantially higher tick-loss. They thus effectively underestimate volatility risk and fail to accurately model the tails of the portfolio variance distribution. This is underlined by the unconditional and conditional coverage tests. The matrix-valued benchmark models have significant UC test results for all values of q, and also their CC test results generally have low p-values. The ReλHAR model, by contrast, safely passes both the UC and CC tests at all levels of q.

To isolate the effect of the eigenvalue-eigenvector decomposition on the empirical fit, we compare the ReλHAR model to the DCC-CAF benchmark model. The main difference between these models lies in the decomposition used: eigenvalues-eigenvectors versus variances-correlations. In addition, the ReλHAR model allows for more flexibility via the copula approach. We see that the DCC-CAF model performs much worse than the ReλHAR model in terms of VolaR, tick-loss, UC, and CC. Concentrating on the eigenvalues and providing their dynamics with more flexibility thus increases the model's usefulness for risk measurement. To understand this, we note that the first eigenvalue generally explains around 55% of the variance in the realized covariance matrix. Capturing its dynamics well thus immediately has an important impact on the quality of the realized covariance matrix predictions. By contrast, the DCC-CAF does not have a similarly important single time-varying parameter, as all realized variances (and correlations) of all assets have a much more symmetric impact on the total realized covariance matrix predictions and associated *VolaR*s.

It is also interesting to see that the simple univariate F_{pfv} performs reasonably well in several dimensions. Its tick-loss is worse than the ReλHAR model, but better than that of the other three benchmark models. The CC results of the model are typically fine, illustrating that a simple univariate realized volatility model for a fixed portfolio is generally able to avoid clustering of *VolaR* violations. The UC results, however, illustrate that the univariate model has difficulty in capturing the unconditional tail probabilities of realized portfolio volatilities, particularly in the far tails. This contrasts with the ReλHAR model, which always passes the UC test for all tail levels q. Compared to the matrix

Table 4: *VolaR* Performance 2007-2023

This table shows the *VolaR* application results for five large US financial stocks: AXP, JPM, AIG, BAC, C, using 4000 out-of-sample observations ranging from January 2007 to January 2023. The restricted ReλHAR model uses the independence rather than the t copula. The benchmark models were introduced in Section 4.4. Panels A, B, and C show the *VolaR* measures for the 5%, 2.5% and 1% exceedance levels, respectively. For each nominal tail level q , the empirical exceedance fraction $\hat{q}$ and Tick-loss measure L^q are shown. Furthermore, the test statistic and p-value for the Unconditional (UC) and Conditional (CC) Coverage tests are shown. For the Tick-loss measure, the Model Confidence Set (MCS) p-value is provided. The portfolio variance distributions are simulated using 100000 draws. The best model for each measure is bolded.

Metric	ReλHAR	ReλHAR$^{\text{res}}$	F_{pfv}	Matrix-F	DPC-CAW	DCC-CAF
Panel A: 5% level						
$\hat{q}$	0.046	**0.048**	0.055	0.066	0.097	0.152
Tick-loss L^q	**0.390**	**0.390**	0.433	0.461	0.541	0.522
MCS p-value	0.791	**1.000**	0.682	0.112	0.072	0.008
UC stat	1.213	**0.261**	2.042	18.511	147.704	579.034
UC p-value	0.271	**0.610**	0.153	0.000	0.000	0.000
CC stat	**0.025**	0.320	0.319	2.817	57.850	1.161
CC p-value	**0.875**	0.572	0.572	0.093	0.000	0.281
Panel B: 2.5% level						
$\hat{q}$	**0.026**	**0.026**	0.033	0.046	0.077	0.116
Tick-loss L^q	**0.264**	0.265	0.313	0.357	0.451	0.431
MCS p-value	**1.000**	0.690	0.611	0.105	0.074	0.013
UC stat	**0.091**	0.162	8.446	59.485	288.253	724.776
UC p-value	**0.762**	0.687	0.004	0.000	0.000	0.000
CC stat	0.628	**0.574**	0.709	2.217	44.117	3.077
CC p-value	0.428	**0.449**	0.400	0.137	0.000	0.079
Panel C: 1% level						
$\hat{q}$	**0.013**	**0.013**	0.017	0.031	0.058	0.085
Tick-loss L^q	**0.155**	0.156	0.200	0.258	0.368	0.339
MCS p-value	**1.000**	0.404	0.404	0.112	0.066	0.024
UC stat	**2.811**	**2.811**	15.303	112.085	433.897	869.706
UC p-value	**0.094**	**0.094**	0.000	0.000	0.000	0.000
CC stat	**0.000**	**0.000**	0.014	3.855	39.363	4.159
CC p-value	**1.000**	**1.000**	0.905	0.050	0.000	0.041

distribution based models, the univariate F_{pfv} thus seems able to avoid much of the noise in the realized covariance matrix estimate by only considering its projection along the fixed portfolio dimension. However, it lacks the flexibility to accommodate different tail behavior of the realized covariance matrix elements along the different spectral dimensions, as well as the possibly changing conditional correlations between the realized eigenvalues. The ReλHAR model, by contrast, picks up both of these effects in a numerically efficient way.

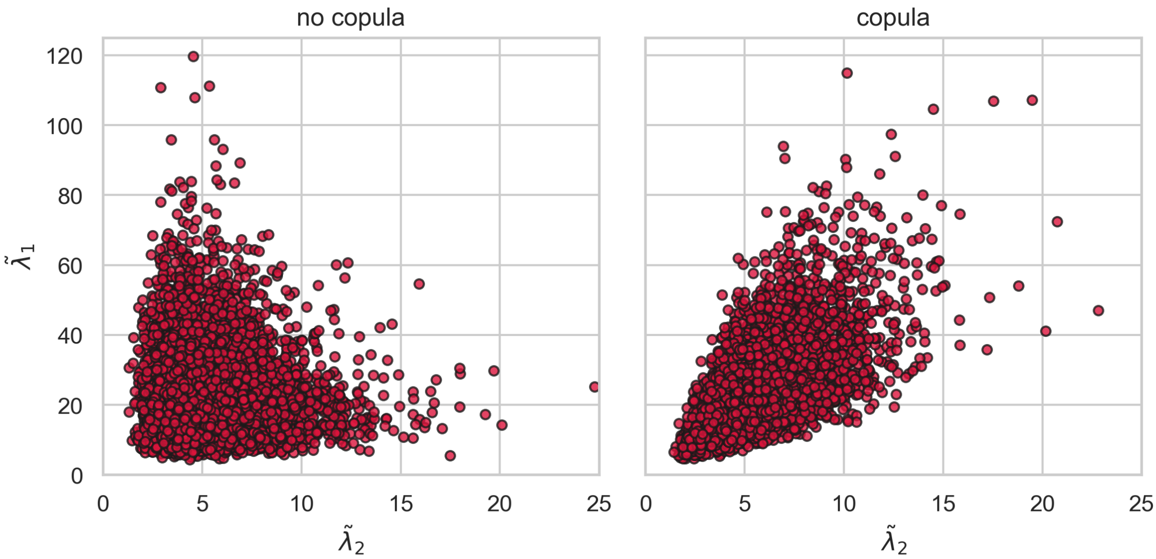

Figure 3: t versus Independence Copula

This figure shows the distribution of the largest and 2nd largest eigenvalue. The left panel shows the bivariate distribution for an independence copula. The right panel shows the bivariate distribution under a Student's t conditional copula. The figure shows 10000 simulations from the bivariate distribution. Estimating the t copula on the inverse PITs yields correlation $\rho_{12} \approx 0.7$ and $\nu \approx 15$.

Comparing the ReλHAR model with a t or an independence copula specification, we find no clear ordering of the two versions of the ReλHAR: both models are typically in the MCS, and their UC and CC performance is very similar, sometimes favoring the t copula, and sometimes the independence copula. We conclude that the addition of conditional dependence to the eigenvalue model substantially increases the likelihood fit, but has little effect on the *VolaR* performance.

Figure 3 shows scatter diagrams of 10000 simulated pairs of the first two eigenvalues of the ReλHAR specification based on a t copula with copula correlation $\rho_{12} = 0.7$ versus an independence copula ($\rho_{12} = 0.0$). There is a clear difference in the bivariate distributions obtained by the two specifications. In particular, the positive copula correlation leads to many more incidences of increased variance along *both* the first and the second spectral dimension. As the equally weighted test portfolio typically lies close to the first spectral dimension, the realized portfolio variance lies close to the first eigenvalue. This explains why the copula inclusion or exclusion has little effect on the *VolaR* quality. The variations do, however, obtain different forecasts as is illustrated in Figure B6 in the appendix.[2]

Figure 4 plots the time series of the 5% *VolaR*s over the out-of-sample period for a selection of three models. Panel (a) shows that the *VolaR* estimates follow the secular

[2] One could alternatively construct test portfolios by taking linear combinations of eigenvectors to imply the portfolio weightings. This can assess how the copula dynamics can accurately capture the overall matrix dynamics. Since such linear combinations need to be normalized before being interpreted as portfolio weights, they amount to specific choices of weight vectors in the asset space, so that our random-weight portfolio robustness check in B.1 provides a more general version of this exercise and shows that the ReλHAR model also performs adequately for other portfolio compositions.

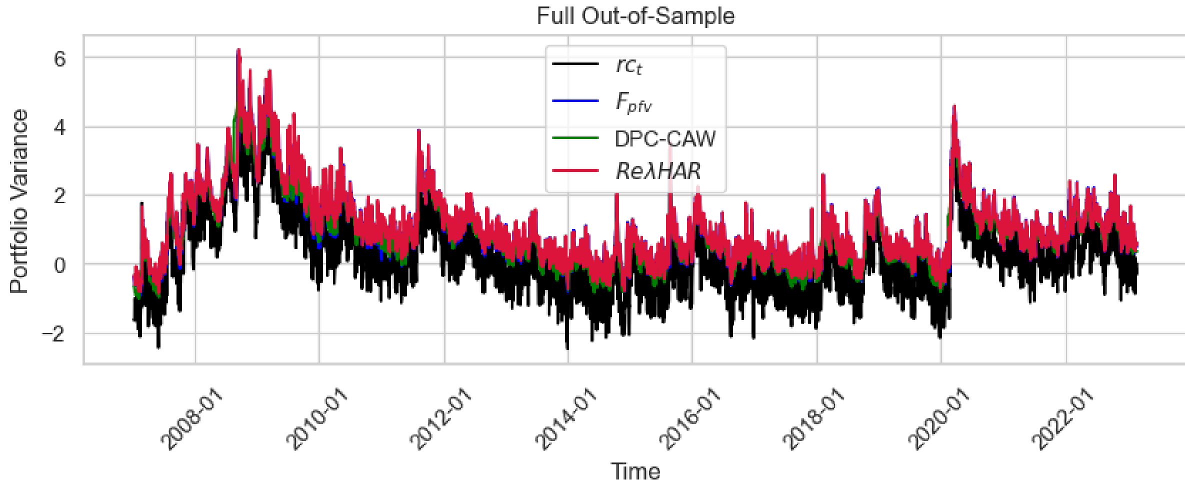

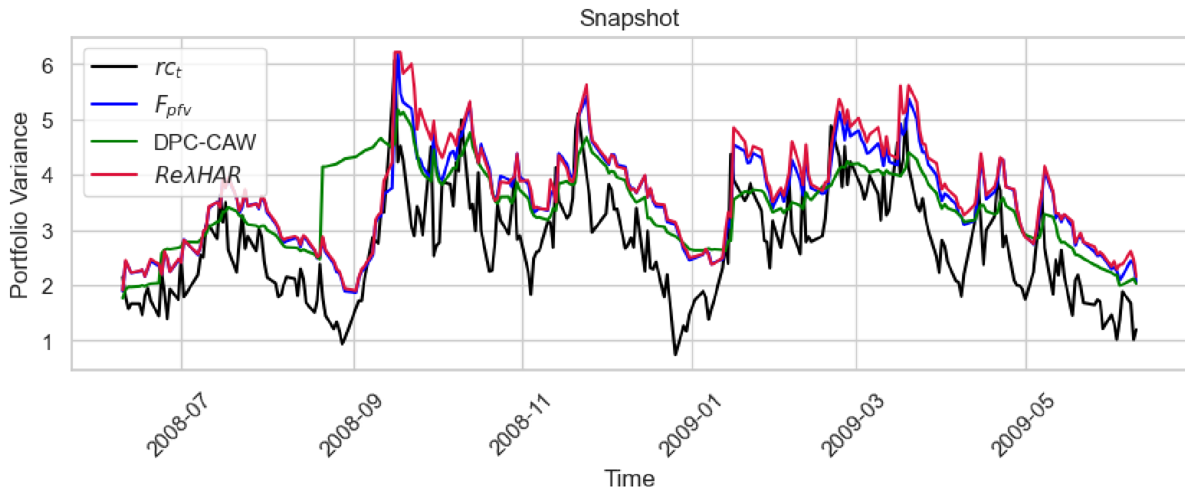

Figure 4: Predicted *VolaR* during 2007-2023 for 5 financial assets.

This figure shows the time series of the predicted portfolio variance *VolaR* estimates for 5 financial assets. The top panel shows the full forecast period from January 2007 to January 2023. The bottom panel shows a snapshot of the series during the peak of financial turmoil from June 2008 to July 2009. In particular, the observed portfolio variances are compared to the estimated 95% quantiles of the portfolio variances. The quantiles are shown for the ReλHAR, F, and DPC-CAW model. The series are based on 5 financial assets. The black line is the observed equally-weighted portfolio variance. The blue, orange, green, and red lines are the 95% portfolio variance quantiles of the F, Matrix-F, DPC-CAW and ReλHAR model, respectively. The natural logarithm is taken for visibility. The quantiles are again taken based on 100000 simulated density values for each model.

dynamics of the portfolio variance rc_t, despite the noisiness of the latter. Zooming in on a snapshot of the financial crisis, the pattern is even clearer. The *VolaR* based on the thin-tailed Wishart (DPC-CAW) clearly underestimates the *VolaR*. The fat-tailed HAR type models do a much better job, with the ReλHAR model being slightly more conservative than the univariate F_{pfv} model, resulting in the improved *VolaR* documented earlier in Table 4.

The overall tail shape accuracy of the models is analyzed by the Conditional VolaR. Table 5 shows the test statistics and p-values of the Unconditional and Conditional Expected Shortfall tests for 5 financial assets, and for the 5%, 2.5%, and 1% quantile.

Table 5: Conditional *VolaR* Performance.

This table shows the Conditional *VolaR* results for 5 financial assets with 4000 observations ranging from January 2007 to January 2023. The proposed ReλHAR model is depicted in Column 1. The remaining columns represent the benchmark models of Section 4.4. Panels A, B, and C show the conditional *VolaR* measures for the 5%, 2.5% and 1% quantile levels, respectively. For each quantile, the test statistic and p-value for the Unconditional Expected Shortfall and Conditional Expected Shortfall are shown. The portfolio variance distributions are simulated using 100000 density values. The best model for each measure is bolded.

Metric	ReλHAR	F_{pfv}	Matrix-F	DPC-CAW	DCC-CAF
Panel A: 5% level					
UES stat	**0.033**	2.751	9.552	23.785	42.762
UES p-value	**0.487**	0.003	0.000	0.000	0.000
CES stat	**11.375**	21.222	61.767	200.890	425.279
CES p-value	**0.329**	0.020	0.000	0.000	0.000
Panel B: 2.5% level					
UES stat	**0.980**	3.699	13.654	31.584	51.282
UES p-value	**0.164**	0.000	0.000	0.000	0.000
CES stat	**5.805**	13.887	60.757	195.173	380.340
CES p-value	**0.831**	0.178	0.000	0.000	0.000
Panel C: 1% level					
UES stat	**0.947**	3.652	18.240	40.808	63.032
UES p-value	**0.172**	0.000	0.000	0.000	0.000
CES stat	**5.650**	9.651	55.616	140.383	309.240
CES p-value	**0.844**	0.472	0.000	0.000	0.000

The table confirms the results obtained in the *VolaR* application. The expected shortfall tests indicate that the ReλHAR model outperforms the benchmark models. The univariate F_{pfv} model is again able to avoid rejections conditionally, but performs considerably worse than the ReλHAR model. This indicates that the flexibility of the F_{pfv} model makes it able to adjust to recent information quickly, but overall fails to capture the shape of the distribution due to filtering much of the data. Again we find that the ReλHAR model is able to avoid rejections in most tests, underlining the usefulness of its modular set-up in terms of separate marginals plus a copula structure.

To illustrate this difference in distribution shape further, Figure 5 shows an instance of the simulated portfolio variance densities. In particular, the kernel density estimates (KDE) of the portfolio variance are shown for January 14th 2009 and provide an example of the different distributional shape that can be obtained by the different models. The DPC-CAW KDE shows how the Wishart assumption in the model leads to thin tails in the portfolio density estimate. This follows directly from the Wishart distributional assumption, which collapses to a scaled χ^2 distribution for the univariate portfolio variance. The F distributions obtained by the alternative models lead to fat-tailed behavior. Differences between models, however, are still considerable. In particular, the

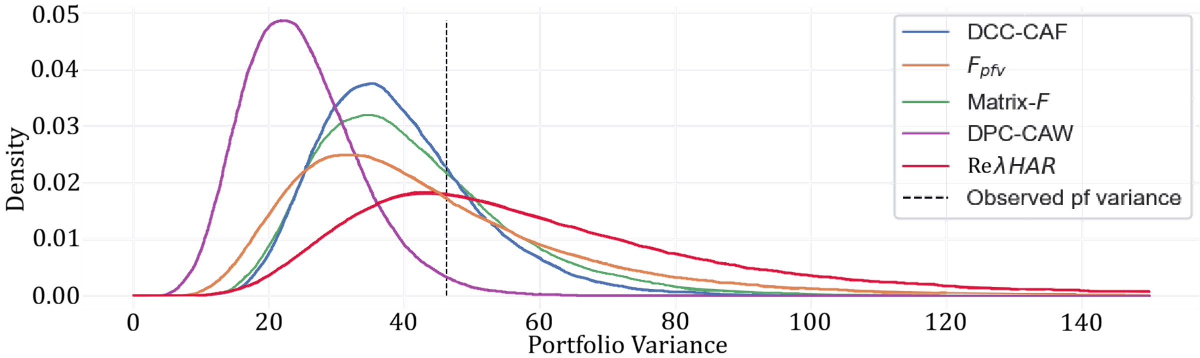

Figure 5: Portfolio Variance Distribution.

This figure shows kernel density estimates (KDE) of forecasted portfolio variance distribution of the ReλHAR, F, DCC-CAF, DPC-CAW and Matrix-F model. The KDE is shown for January 14, 2009. The densities are based on 5 financial assets. The blue, orange, green, and red lines are the KDE values of the F, Matrix-F, DPC-CAW and ReλHAR model, respectively. The KDEs are based on 100000 simulated density values for each model.

ReλHAR model has a much larger variance on this example day than the F and Matrix-F model, leading to larger *VolaR* estimates and better behavior in terms of risk measures.

The *VolaR* results for the period 2007 to 2023 indicate that the ReλHAR performs best out of all considered models. To see where the gain comes from, Table 6 shows the *VolaR* and ES results for the different models during the Financial crisis and COVID-19 period, i.e., January 2007 to January 2013 and November 2019 to November 2020, or in total 1750 observations.

The table shows that if we focus on periods of financial turmoil, the ReλHAR model also outperforms the benchmarks. Although rejections are obtained in UC and UES tests, we find that the benchmark models cannot accurately capture the portfolio variance dynamics. Ultimately, the ReλHAR model seems able to capture most of the dynamics in the series of portfolio volatilities. It consistently outperforms the benchmark models across the various measures and tests, both in-sample, out-of-sample, and during turbulent periods. It is able to achieve highly accurate $\hat{q}$ for the different tail risk levels and often avoids rejections in the coverage and expected shortfall tests. The ReλHAR model also attains the smallest Tick-loss levels. All results thus point to the value-added of using the eigenvalue-eigenvector decomposition methodology and modular marginal-copula modeling set-up of the ReλHAR model for accurate modeling of realized portfolio volatilities.

6. Conclusion

We introduced the ReλHAR model for estimating and forecasting the distribution of realized portfolio variances. The ReλHAR model uses the eigenvalue-eigenvector decomposition of realized covariance matrices to accurately describe their dynamics.

Table 6: (Conditional) *VolaR* Performance During Financial Turmoil.

This table shows the *VolaR* application results for 5 financial assets with 1750 observations during the crisis periods, ranging from January 2007 to January 2013 and November 2019 to November 2020. The proposed ReλHAR model is depicted in Column 1. The remaining columns represent the benchmark models of Section 4.4. Panels A and B show the *VolaR* measures for the 5% and 1% quantile levels, respectively. For each nominal quantile, the empirical exceedance fraction $\hat{q}$ and Tick-loss measure L^q are shown. Furthermore, the test statistic and p-value for the coverage and expected shortfall tests are shown. For the Tick-loss measure, the p-value of the Model Confidence Set (MCS) is shown at a 5% significance level. The portfolio variance distributions are simulated using 100000 density values. The best model value for each measure is balded.

Metric	ReλHAR	F_{pfv}	Matrix-F	DPC-CAW	DCC-CAF
Panel A: 5% level					
$\hat{q}$	**0.056**	0.071	0.079	0.110	0.165
Tick-loss L^q	**0.759**	0.856	0.920	1.097	1.038
MCS p-value	**0.914**	0.652	0.055	0.040	0.011
UC stat	**1.378**	14.422	26.396	102.394	318.549
UC p-value	**0.240**	0.000	0.000	0.000	0.000
CC stat	**1.013**	1.927	2.381	42.515	2.020
CC p-value	**0.314**	0.165	0.123	0.000	0.155
UES stat	**1.742**	4.986	10.896	20.185	32.352
UES p-value	**0.041**	0.000	0.000	0.000	0.000
CES stat	**18.694**	28.088	73.100	166.812	325.681
CES p-value	**0.044**	0.002	0.000	0.000	0.000
Panel B: 1% level					
$\hat{q}$	**0.016**	0.022	0.042	0.071	0.095
Tick-loss L^q	**0.298**	0.398	0.525	0.762	0.682
MCS p-value	**1.000**	0.428	0.101	0.049	0.012
UC stat	**5.027**	19.036	103.221	283.569	471.698
UC p-value	**0.025**	0.000	0.000	0.000	0.000
CC stat	**0.000**	0.024	3.877	26.886	7.868
CC p-value	**1.000**	0.876	0.049	0.000	0.005
UES stat	**1.746**	4.624	17.965	35.060	50.622
UES p-value	**0.040**	0.000	0.000	0.000	0.000
CES stat	**7.391**	12.359	64.010	104.074	261.811
CES p-value	**0.688**	0.262	0.000	0.000	0.000

The eigenvalues capture the dominant direction of time-series variation of the realized covariance matrices. We modeled the eigenvalues separately using a conditional F distribution with time-varying mean with HAR dynamics. The conditional dependence between eigenvalues was captured using a copula. This set-up made the ReλHAR model highly efficient for empirical risk management applications. Due to its modular set-up, the model is also easily scalable if the number of assets increases.

In an extensive (Conditional) Volatility-at-Risk (*VolaR*) application on US blue chip stocks, we found that the new ReλHAR model outperforms a number of recent benchmarks, even after further empirically improving these benchmarks. In particular, compared to recent matrix-valued models, we found that the new composite modeling approach provided better *VolaR* results in terms of coverage, tick-loss, and other portfolio variance risk measures.

In conclusion, the ReλHAR model substantially and consistently outperformed the benchmark models from the literature. It attains the most accurate *VolaR* level estimates over time in combination with the smallest error measured in the tick-loss. Furthermore, it is able to avoid rejections in the unconditional and conditional tests. The coverage test indicated that indeed the ReλHAR model more precisely estimates the *VolaR* level and is able to accurately incorporate new information during financial turmoil. This is further underlined by the Expected Shortfall tests, indicating that the tail dynamics are accurately captured. This makes the ReλHAR model a highly accurate indicator of volatility risk and a useful tool for modeling realized covariance matrix dynamics in a parsimonious and scalable way.

Data availability statement

The raw HF data that support the findings of this study are available at the TAQ database via WRDS. Restrictions apply to the availability of these data, which were used under license for this study. Calculated Realized Covariance matrices are available from the authors.

References

Andersen, T. G., Bollerslev, T., Christoffersen, P. F., and Diebold, F. X. (2007). Practical volatility and correlation modeling for financial market risk management. In *The Risks of Financial Institutions*, pages 513–544. University of Chicago Press, Chicago.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2009). Realized kernels in practice: Trades and quotes. *The Econometrics Journal*, 12(3):C1–C32.

Bauwens, L., Storti, G., Violante, F., et al. (2012). Dynamic conditional correlation models for realized covariance matrices. *CORE DP*, 60:104–108.

Bauwens, L. and Xu, Y. (2023). DCC-and DECO-HEAVY: Multivariate GARCH models based on realized variances and correlations. *International Journal of Forecasting*, 39(2):938–955.

Berkes, I. and Horváth, L. (2004). The efficiency of the estimators of the parameters in garch processes. *The Annals of Statistics*, 32(2):633–655.

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

Brownlees, C. T. and Gallo, G. M. (2006). Financial econometric analysis at ultra-high frequency: Data handling concerns. *Computational statistics & data analysis*, 51(4):2232–2245.

Caporin, M., Kolokolov, A., and Renò, R. (2017a). Systemic co-jumps. *Journal of Financial Economics*, 126(3):563–591.

Caporin, M., Rossi, E., and De Magistris, P. S. (2017b). Chasing volatility: A persistent multiplicative error model with jumps. *Journal of Econometrics*, 198(1):122–145.

Chamberlain, G. and Rothschild, M. (1983). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, 51(5):1281–1304.

Chen, X., Fan, Y., and Tsyrennikov, V. (2006). Efficient estimation of semiparametric multivariate copula models. *Journal of the American Statistical Association*, 101(475):1228–1240.

Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, 39(4):841–862.

Christoffersen, P. F. and Diebold, F. X. (2000). How relevant is volatility forecasting for financial risk management? *Review of Economics and Statistics*, 82(1):12–22.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196.

Creal, D., Koopman, S. J., and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795.

Du, Z. and Escanciano, J. C. (2017). Backtesting expected shortfall: accounting for tail risk. *Management Science*, 63(4):940–958.

Engle, R. (2002a). 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. (2002b). New frontiers for arch models. *Journal of Applied Econometrics*, 17(5):425–446.

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

Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6):1545–1578.

Gouriéroux, C., Jasiak, J., and Sufana, R. (2009). The wishart autoregressive process of multivariate stochastic volatility. *Journal of Econometrics*, 150(2):167–181.

Gribisch, B. and Stollenwerk, M. (2020). Dynamic principal component caw models for high-dimensional realized covariance matrices. *Quantitative Finance*, 20(5):799–821.

Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2):453–497.

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Number 52. Cambridge University Press.

Hediger, S., Näf, J., Paolella, M. S., and Polak, P. (2023). Heterogeneous tail generalized common factor modeling. *Digital Finance*, 5(2):389–420.

Hetland, S., Pedersen, R. S., and Rahbek, A. (2023). Dynamic conditional eigenvalue GARCH. *Journal of Econometrics*, 237(2):105175.

Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge university press.

Huang, D., Schlag, C., Shaliastovich, I., and Thimme, J. (2019). Volatility-of-volatility risk. *Journal of Financial and Quantitative Analysis*, 54(6):2423–2452.

Jin, X. and Maheu, J. M. (2013). Modeling realized covariances and returns. *Journal of Financial Econometrics*, 11(2):335–369.

Joe, H. and Xu, J. (2005). Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of Multivariate Analysis*, 94(2):401–419.

Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2):73–84.

Laloux, L., Cizeau, P., Potters, M., and Bouchaud, J.-P. (2000). Random matrix theory and financial correlations. *International Journal of Theoretical and Applied Finance*, 3(03):391–397.

Näf, J., Paolella, M. S., and Polak, P. (2019). Heterogeneous tail generalized comfort modeling via cholesky decomposition. *Journal of Multivariate Analysis*, 172:84–106.

Noureldin, D., Shephard, N., and Sheppard, K. (2012). Multivariate high-frequency-based volatility (heavy) models. *Journal of Applied Econometrics*, 27(6):907–933.

Opschoor, A., Janus, P., Lucas, A., and Van Dijk, D. (2018). New heavy models for fat-tailed realized covariances and returns. *Journal of Business & Economic Statistics*, 36(4):643–657.

Opschoor, A., Lucas, A., and Rossini, L. (2025). The conditional autoregressive f-riesz model for realized covariance matrices. *Journal of Financial Econometrics*, 23(2):nbae023.

Patton, A. J. (2006). Estimation of multivariate models for time series of possibly different lengths. *Journal of applied econometrics*, 21(2):147–173.

Shephard, N. and Sheppard, K. (2010). Realising the future: forecasting with high-frequency-based volatility (heavy) models. *Journal of Applied Econometrics*, 25(2):197–231.

Sinclair, E. (2013). *Volatility trading*. John Wiley & Sons.

Stollenwerk, M. (2026). Probability distributions for realized covariance measures. *Journal of Econometrics*, 254(Part A):105954.

Straumann, D. and Mikosch, T. (2006). Quasi-maximum-likelihood estimation in conditionally heteroscedastic time series: A stochastic recurrence equations approach. *Annals of Statistics*, 34(5):2449–2495.

Tasche, D. (2002). Expected shortfall and beyond. *Journal of Banking & Finance*, 26(7):1519–1533.

Vassallo, D., Buccheri, G., and Corsi, F. (2021). A DCC-type approach for realized covariance modeling with score-driven dynamics. *International Journal of Forecasting*, 37(2):569–586.

Zhou, J., Jiang, F., Zhu, K., and Li, W. K. (2022). Time series models for realized covariance matrices based on the matrix-f distribution. *Statistica Sinica*, 32(2):755–786.

Zumbach, G. (2011). Empirical properties of large covariance matrices. *Quantitative Finance*, 11(7):1091–1102.

Supplementary Materials for "Composite Univariate Modeling of Realized Covariance Matrix Dynamics and Volatility-at-Risk"

This supplement includes all technical proofs and additional results.

Contents

A. Theoretical Appendix

A.1. Proofs

Proof of Theorem 1: We establish the consistency and asymptotic normality of the first-stage ML Estimators (MLEs) of $\boldsymbol{\theta}_i = (\nu_{1,i}, \nu_{2,i}, \omega_i, \alpha_i^d, \alpha_i^w, \alpha_i^m, \beta_i)^\top$ in Model (2)–(4). For notational simplicity, we drop the subscript i throughout, and use a generic constant c to bound functions in different contexts.

We apply the well-established literature on GARCH models to a MEM model of the form

$$\lambda_t = \mu_t \varepsilon_t, \quad \varepsilon_t \overset{iid}{\sim} F(1, \nu_1, \nu_2)$$

where μ_t is time-varying under the stochastic recurrent equation (SRE)

$$\mu_t = (1 - \alpha^d - \alpha^w - \alpha^m - \beta)\omega + \alpha^d \mu_{t-1}\varepsilon_{t-1} + \alpha^w \sum_{i=1}^{5} \mu_{t-i}\varepsilon_{t-i} + \alpha^m \sum_{i=1}^{22} \mu_{t-i}\varepsilon_{t-i} + \beta\mu_{t-1}.$$

Using $\alpha^d, \alpha^w, \alpha^m, \beta > 0$ and the sufficient condition $\alpha^d + \alpha^w + \alpha^m + \beta < 1$ of Bollerslev (1986) we obtain that the sequence μ_t converges e.a.s. to a unique strictly stationary, ergodic sequence uniformly in θ, i.e., that the filter is uniformly invertible (see, e.g. Francq and Zakoïan, 2019).

The properties of the MLE of $\boldsymbol{\theta}$ are then studied using the ML criterion function

$$Q_T(\boldsymbol{\theta}) = \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{\theta}). \tag{A.1}$$

Here, $\ell_t(\boldsymbol{\theta})$ denotes the loglikelihood contribution evaluated at the stationary solution $\mu_t(\boldsymbol{\theta})$, whereas $\hat{\ell}_t(\boldsymbol{\theta})$ denotes the corresponding contribution evaluated at the initialized filtered path $\hat{\mu}_t(\boldsymbol{\theta})$, obtained with the fixed initial value $\hat{\mu}_1$. We write

$$\hat{Q}_T(\boldsymbol{\theta}) = \frac{1}{T}\sum_{t=1}^{T} \hat{\ell}_t(\boldsymbol{\theta}). \tag{A.2}$$

Uniqueness of the optimizer follows, given that ℓ_t is the loglikelihood contribution of the F density, whose static parameters are identified under correct specification by a standard Gibbs inequality. Identifiability of the remaining parameters is obtained under the sufficient condition of Berkes and Horváth (2004) that the polynomial

$$\alpha^d/\beta + \alpha^w P_5(\beta^{-1}) + \alpha^m P_{22}(\beta^{-1})$$

is nonzero, which holds by imposing the sufficient condition $\alpha^d, \alpha^w, \alpha^m, \beta > 0$.

To establish the convergence of the feasible criterion $\hat{Q}_T$ to its stationary counterpart Q_T, we first establish a uniform bound on $\partial \ell_t / \partial \mu_t$. From the conditional F loglikelihood contribution, the terms depending on μ_t are

$$\ell_t(\boldsymbol{\theta}) = -\frac{\nu_1}{2} \log \mu_t - \frac{\nu_1 + \nu_2}{2} \log \left(1 + \frac{\nu_1}{\nu_2 - 2} \frac{\lambda_t}{\mu_t}\right) + C_t(\nu_1, \nu_2, \lambda_t),$$

where $C_t(\nu_1, \nu_2, \lambda_t)$ does not depend on μ_t. We then get

$$\begin{aligned}
\sup_{\boldsymbol{\theta}\in\Theta} \left| \frac{\partial \ell_t(\boldsymbol{\theta})}{\partial \mu_t} \right| &\leq \sup_{\boldsymbol{\theta}\in\Theta} \left| \frac{\nu_1}{2\mu_t} \right| + \sup_{\boldsymbol{\theta}\in\Theta} \sup_{(\lambda/\mu)\in\mathbb{R}^+} \left| \frac{\nu_1 + \nu_2}{2} \frac{\nu_1(\lambda/\mu)}{(\nu_2 - 2) + \nu_1(\lambda/\mu)} \right| \\
&\leq \sup_{\boldsymbol{\theta}\in\Theta} \left| \frac{\nu_1}{2\underline{\omega}} \right| + \sup_{\boldsymbol{\theta}\in\Theta} \left| \frac{\nu_1 + \nu_2}{2} \right| < \infty,
\end{aligned} \tag{A.3}$$

for all t, where the second inequality follows from $\inf_t \mu_t \geq \omega \geq \underline{\omega} > 0$ for $\underline{\omega} = \inf_{\boldsymbol{\theta}\in\Theta} \omega$, and the final inequality follows from the compactness of the parameter space. The mean-value theorem then yields

$$\begin{aligned}
\sup_{\boldsymbol{\theta}\in\Theta} |\hat{Q}_T(\boldsymbol{\theta}) - Q_T(\boldsymbol{\theta})| &\leq \frac{1}{T} \sum_{t=1}^{T} \sup_{\boldsymbol{\theta}\in\Theta} |\hat{\ell}_t(\boldsymbol{\theta}) - \ell_t(\boldsymbol{\theta})| \\
&\leq \frac{1}{T} \sum_{t=1}^{T} \sup_{\boldsymbol{\theta}\in\Theta} \left| \frac{\partial \ell_t(\tilde{\mu}_t(\boldsymbol{\theta}), \boldsymbol{\nu})}{\partial \mu_t} \right| \cdot \sup_{\boldsymbol{\theta}\in\Theta} |\hat{\mu}_t(\boldsymbol{\theta}) - \mu_t(\boldsymbol{\theta})| \\
&\leq \frac{c}{T} \sum_{t=1}^{T} \sup_{\boldsymbol{\theta}\in\Theta} |\hat{\mu}_t(\boldsymbol{\theta}) - \mu_t(\boldsymbol{\theta})| \overset{a.s.}{\to} 0,
\end{aligned} \tag{A.4}$$

where $\tilde{\mu}_t$ lies between $\hat{\mu}_t(\boldsymbol{\theta})$ and $\mu_t(\boldsymbol{\theta})$, and where the final convergence follows from the initialized filter converging e.a.s. to its stationary counterpart; see, e.g., Lemma 2.1 of Straumann and Mikosch (2006).

Define the population criterion function $Q_\infty(\boldsymbol{\theta}) := \mathbb{E}\left[\ell_t(\boldsymbol{\theta})\right]$. Next, we establish the uniform convergence of the stationary criterion $Q_T(\boldsymbol{\theta})$ to $Q_\infty(\boldsymbol{\theta})$. This follows from the uniform ergodic theorem of Ranga Rao (1962) (see also Theorem 2.7 of Straumann and Mikosch, 2006). Let $\ell_t(\cdot) \in C(\Theta)$. Since Θ is compact and the conditional F loglikelihood is continuous in (μ_t, ν_1, ν_2) on the admissible parameter space, $\ell_t(\boldsymbol{\theta})$ is a.s. continuous in $\boldsymbol{\theta}$. Moreover, $\{\ell_t\}$ is stationary and ergodic because it is a measurable transformation of the stationary and ergodic sequence generated by the SRE. Furthermore,

$$
\begin{aligned}
&\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}|\ell_t(\boldsymbol{\theta})|\right] \\
&=\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}\left|\log\Gamma\left(\frac{\nu_1+\nu_2}{2}\right)-\log\Gamma\left(\frac{\nu_1}{2}\right)-\log\Gamma\left(\frac{\nu_2}{2}\right)+\frac{\nu_1}{2}\log\left(\frac{\nu_1}{(\nu_2-2)\mu_t(\boldsymbol{\theta})}\right)\right.\right. \\
&\qquad\qquad\left.\left.+\frac{\nu_1-2}{2}\log(\lambda_t)-\frac{\nu_1+\nu_2}{2}\log\left(1+\frac{\nu_1}{\nu_2-2}\frac{\lambda_t}{\mu_t(\boldsymbol{\theta})}\right)\right|\right] \\
&\leq\sup_{\boldsymbol{\theta}\in\Theta}\left|\log\Gamma\left(\frac{\nu_1+\nu_2}{2}\right)\right|+\sup_{\boldsymbol{\theta}\in\Theta}\left|\log\Gamma\left(\frac{\nu_1}{2}\right)\right|+\sup_{\boldsymbol{\theta}\in\Theta}\left|\log\Gamma\left(\frac{\nu_2}{2}\right)\right| \\
&\quad+\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}\left|\frac{\nu_1}{2}\log\left(\frac{\nu_1}{\nu_2-2}\right)\right|\right]+\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}\left|\frac{\nu_1}{2}\log\mu_t(\boldsymbol{\theta})\right|\right] \\
&\quad+\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}\left|\frac{\nu_1-2}{2}\log\lambda_t\right|\right]+\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}\left|\frac{\nu_1+\nu_2}{2}\log\left(1+\frac{\nu_1}{\nu_2-2}\frac{\lambda_t}{\mu_t(\boldsymbol{\theta})}\right)\right|\right].
\end{aligned}
\tag{A.5}
$$

The first four terms are finite by compactness of Θ, the imposed bounds $\nu_1 > 0$ and $\nu_2 > 4$, and the continuity of the log-gamma function on the admissible parameter space. For the remaining terms, compactness of Θ implies that ν_1 and ν_2 are uniformly bounded from above. We also have $\inf \nu_2 > 2$, and $\mu_t(\boldsymbol{\theta}) \geq \underline{\omega} > 0$. Moreover, the contraction condition $\alpha^d + \alpha^w + \alpha^m + \beta \leq \rho < 1$ together with the uniform bound in (A.3) implies a uniform bound $\sup_{\boldsymbol{\theta}\in\Theta} \mu_t(\boldsymbol{\theta}) \leq c < \infty$, thus bounding the fifth term and seventh term by also using the assumption that λ_t has a first moment due to the true $\nu_{0,2} > 2$. Finally,

$$
\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}\left|\frac{\nu_1-2}{2}\log\lambda_t\right|\right] \leq c\,\mathbb{E}\,|\log\mu_t(\boldsymbol{\theta}_0)| + c\,\mathbb{E}\,|\log\varepsilon_t| \leq c_1 + c\,\mathbb{E}\,|\log\varepsilon_t|,
$$

where the last expectation is finite for $\varepsilon_t \sim F(1, \nu_{0,1}, \nu_{0,2})$ due to $\nu_{0,1}, \nu_{0,2} > 0$. Combining these bounds gives $\mathbb{E}\left[\sup_{\boldsymbol{\theta}\in\Theta}|\ell_t(\boldsymbol{\theta})|\right] < \infty$, so that the uniform ergodic theorem yields

$$
\sup_{\boldsymbol{\theta}\in\Theta}|Q_T(\boldsymbol{\theta}) - Q_\infty(\boldsymbol{\theta})| \xrightarrow{a.s.} 0.
$$

Combining this result with (A.4), we have $\sup_{\boldsymbol{\theta}\in\Theta}|\hat{Q}_T(\boldsymbol{\theta}) - Q_\infty(\boldsymbol{\theta})| \xrightarrow{a.s.} 0$. This establishes consistency.

Asymptotic normality (AN) of the MLE of $\boldsymbol{\theta}$ is established using the conditions of Straumann (2005). The proof is naturally adjusted to the MEM setting in (2)–(4), in contrast to the GARCH setting in Straumann (2005). The key feature we adopt is that the (expected) information matrix admits a *block decomposition* with respect to the parameters in the recursion for $\mu_t(\vartheta)$. To avoid additional notation and subsequent lengthy arguments, we refer to Straumann (2005) and directly verify the relevant necessary conditions.

First, we partition the full parameter vector as $\boldsymbol{\theta} = (\boldsymbol{\vartheta}^\top, \boldsymbol{\nu}^\top)^\top \in \mathbb{R}^7$, where $\boldsymbol{\vartheta} = (\omega, \alpha^d, \alpha^w, \alpha^m, \beta)^\top$, and $\boldsymbol{\nu} = (\nu_1, \nu_2)^\top$. Following Straumann (2005), we write the parameter space as $\Theta = K \times V$, with $\boldsymbol{\vartheta} \in K$ and $\boldsymbol{\nu} \in V$. The model can be written as

$$\lambda_t = \mu_t(\boldsymbol{\vartheta}_0)\varepsilon_t, \qquad \varepsilon_t \overset{iid}{\sim} k_{\boldsymbol{\nu}_0}(\varepsilon)\, d\varepsilon,$$

where $k_{\boldsymbol{\nu}}$ denotes the unit-mean scaled $F(1, \nu_1, \nu_2)$ innovation density. Thus,

$$k_{\boldsymbol{\nu}}(\varepsilon) = \frac{\Gamma((\nu_1+\nu_2)/2)}{\Gamma(\nu_1/2)\Gamma(\nu_2/2)} \varepsilon^{(\nu_1-2)/2} \left(\frac{\nu_1}{\nu_2-2}\right)^{\nu_1/2} \left(1 + \frac{\nu_1}{\nu_2-2}\varepsilon\right)^{-(\nu_1+\nu_2)/2}, \qquad \varepsilon > 0.$$

Continuing upon the results in the proof of consistency of the MLE, it suffices to verify conditions M.4-M.9 in Chapter 6.3 of Straumann (2005).

Condition M.4 concerns the existence, invertibility, and differentiability of the stationary conditional mean filter. For the present HAR-MEM recursion, the first derivative satisfies

$$\mu_t'(\boldsymbol{\vartheta}) = \begin{pmatrix} 1 \\ \lambda_{t-1} \\ \bar{\lambda}_{t-1}^w \\ \bar{\lambda}_{t-1}^m \\ \mu_{t-1}(\boldsymbol{\vartheta}) \end{pmatrix} + \beta \mu_{t-1}'(\boldsymbol{\vartheta}), \tag{A.6}$$

and, writing $e_\beta = (0, 0, 0, 0, 1)^\top$, its second derivative satisfies

$$\mu_t''(\boldsymbol{\vartheta}) = \beta \mu_{t-1}''(\boldsymbol{\vartheta}) + e_\beta \mu_{t-1}'(\boldsymbol{\vartheta})^\top + \mu_{t-1}'(\boldsymbol{\vartheta}) e_\beta^\top. \tag{A.7}$$

Under the maintained compactness, positivity, contraction, and moment assumptions on K, the standard SRE arguments for GARCH/MEM filters imply that the stationary filter and its first two derivative filters exist, are stationary and ergodic, and that the initialized filters converge to their stationary counterparts exponentially almost surely, uniformly over K (see, e.g., Straumann and Mikosch, 2006; Berkes and Horváth, 2004).

The remaining part of M.4 concerns the filter-dependent population components entering the expected information matrix

$$E\left[\frac{\mu_t'(\boldsymbol{\vartheta}_0)\mu_t'(\boldsymbol{\vartheta}_0)^\top}{\mu_t(\boldsymbol{\vartheta}_0)^2}\right], \quad -E\left[\frac{\mu_t'(\boldsymbol{\vartheta}_0)}{\mu_t(\boldsymbol{\vartheta}_0)}\right],$$

which are finite under the moment conditions imposed on the derivative filter and the uniform lower bound on $\mu_t(\boldsymbol{\vartheta})$. The positive definiteness of the first matrix naturally follows.

Condition M.5 concerns the positivity and smoothness of the innovation density. It is directly satisfied because the unit-mean scaled F density is strictly positive and twice continuously differentiable in the innovation argument and in $\boldsymbol{\nu}$ on every compact subset V of the admissible parameter space.

Condition M.6 concerns the usual score identities and moment requirements for the innovation density. These follow from standard regularity of the scaled F family, with the required moments ensured by the imposed lower bound $\nu_2 > 4$.

Condition M.7 is the corresponding non-degeneracy condition excluding collinearity between the scale direction and the innovation-parameter directions, which is satisfied by the identifiable two-parameter scaled F family on V.

Condition M.8 is important for the application of the Martingale Central Limit Theorem. It requires

$$\mathbb{E}\left\|\frac{\partial^2 \ell_0}{\partial \boldsymbol{\theta}^2}\right\|_\Theta < \infty, \quad \mathbb{E}\left|\frac{\partial \ell_0(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}}\right|^2 < \infty.$$

We verify these two requirements in Appendix A.2 in detail. In particular, the appendix derives the score and Hessian of $\ell_t(\boldsymbol{\theta})$ and shows that the resulting terms admit integrable envelopes uniformly over $\boldsymbol{\theta} \in \Theta$. This combines the regularity and tail restrictions of the scaled F density with the moment bounds for the normalized filter derivatives $\mu_t'(\boldsymbol{\vartheta})/\mu_t(\boldsymbol{\vartheta})$ and $\mu_t''(\boldsymbol{\vartheta})/\mu_t(\boldsymbol{\vartheta})$. Hence the uniform integrability of the Hessian follows. Evaluating the corresponding score expressions at the true parameter $\boldsymbol{\theta}_0$ gives the finite second moment of the score, establishing the second part of M.8. See for further details Appendix A.2.

Condition M.9 regards limit conditions for the log-likelihood. Given the proof of consistency, we additionally prove

$$\sqrt{T}\left\|\frac{\partial \hat{Q}_T}{\partial \boldsymbol{\theta}} - \frac{\partial Q_T}{\partial \boldsymbol{\theta}}\right\|_\Theta \xrightarrow{\text{a.s.}} 0.$$

By the mean-value theorem applied to the score as a function of the filter quantities

(μ_t, μ_t'), there exists a non-negative envelope L_t such that

$$
\begin{aligned}
\sqrt{T} & \left\| \frac{\partial \hat{Q}_T}{\partial \boldsymbol{\theta}} - \frac{\partial Q_T}{\partial \boldsymbol{\theta}} \right\|_{\Theta} \\
& \leq \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \sup_{\boldsymbol{\theta} \in \Theta} \left\| \frac{\partial \hat{\ell}_t(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} - \frac{\partial \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right\| \\
& \leq \frac{1}{\sqrt{T}} \sum_{t=1}^{T} L_t \left(\|\hat{\mu}_t - \mu_t\|_K + \|\hat{\mu}_t' - \mu_t'\|_K \right) \xrightarrow{\text{a.s.}} 0.
\end{aligned}
$$

Here, the envelope L_t collects the relevant mixed derivatives of the score with respect to the filter arguments (μ_t, μ_t'). We show in Eq. (A.15) and its argumentation that this envelope can be chosen stationary and non-negative with $EL_0^q < \infty$ for some $q > 0$. Hence, by Lemma 2.2 of Straumann and Mikosch (2006),

$$
E \log^+ L_0 < \infty.
$$

Moreover, by the exponential forgetting of the initialized filter and derivative filter,

$$
\|\hat{\mu}_t - \mu_t\|_K + \|\hat{\mu}_t' - \mu_t'\|_K \xrightarrow{\text{e.a.s.}} 0.
$$

Applying Lemma 2.1 of Straumann and Mikosch (2006) to the final sum then yields the almost sure convergence, so that also

$$
\sqrt{T} \left\| \frac{\partial \hat{Q}_T}{\partial \boldsymbol{\theta}} - \frac{\partial Q_T}{\partial \boldsymbol{\theta}} \right\|_{\Theta} \xrightarrow{\text{a.s.}} 0,
$$

and hence Condition M.9 holds. This concludes the proof of consistency and asymptotic normality of the estimator of $\boldsymbol{\theta}$. □

A.2. Derivatives and Moments

For a candidate parameter $\boldsymbol{\theta}$, the implied innovation is

$$
\varepsilon_t(\boldsymbol{\vartheta}) = \frac{\lambda_t}{\mu_t(\boldsymbol{\vartheta})},
$$

and the conditional density of λ_t is

$$
f_t(\lambda_t; \boldsymbol{\theta}) = \frac{1}{\mu_t(\boldsymbol{\vartheta})} k_{\boldsymbol{\nu}} \left(\frac{\lambda_t}{\mu_t(\boldsymbol{\vartheta})} \right).
$$

Hence the loglikelihood contribution is

$$\ell_t(\boldsymbol{\theta}) = -\log \mu_t(\boldsymbol{\vartheta}) + \log k_{\boldsymbol{\nu}}\left(\varepsilon_t(\boldsymbol{\vartheta})\right).$$

Let $\ell_0(\boldsymbol{\theta})$ denote the stationary version of this loglikelihood contribution. Define the score and expected Hessian by

$$\dot{\ell}_0(\boldsymbol{\theta}) = \frac{\partial \ell_0(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}, \qquad \mathbf{F}_0 = \mathbb{E}\left[\frac{\partial^2 \ell_0(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}\partial \boldsymbol{\theta}^\top}\right] = -\mathbb{E}\left[\dot{\ell}_0(\boldsymbol{\theta}_0)\dot{\ell}_0(\boldsymbol{\theta}_0)^\top\right],$$

and partition $\mathbf{F}_0$ conformably as

$$\mathbf{F}_0 = \begin{bmatrix} \mathbf{F}_{\boldsymbol{\vartheta}\boldsymbol{\vartheta}} & \mathbf{F}_{\boldsymbol{\vartheta}\boldsymbol{\nu}} \\ \mathbf{F}_{\boldsymbol{\nu}\boldsymbol{\vartheta}} & \mathbf{F}_{\boldsymbol{\nu}\boldsymbol{\nu}} \end{bmatrix}, \qquad \mathbf{F}_{\boldsymbol{\vartheta}\boldsymbol{\vartheta}} \in \mathbb{R}^{d\times d}, \quad \mathbf{F}_{\boldsymbol{\vartheta}\boldsymbol{\nu}} \in \mathbb{R}^{d\times d'}, \quad \mathbf{F}_{\boldsymbol{\nu}\boldsymbol{\nu}} \in \mathbb{R}^{d'\times d'}.$$

We decompose the Hessian into its three blocks:

$$\boldsymbol{\vartheta}\boldsymbol{\vartheta}, \qquad \boldsymbol{\vartheta}\boldsymbol{\nu}, \qquad \boldsymbol{\nu}\boldsymbol{\nu}.$$

First, the derivatives of the scaled F density with respect to the innovation argument are uniformly bounded after multiplication by the appropriate powers of $\varepsilon_t(\boldsymbol{\vartheta})$. Second, the derivatives with respect to $\boldsymbol{\nu}$ are bounded by logarithmic terms, which are integrable under the scaled F distribution. Third, the remaining terms are ratios of $\mu_t'(\boldsymbol{\vartheta})$ and $\mu_t''(\boldsymbol{\vartheta})$ to $\mu_t(\boldsymbol{\vartheta})$, whose required moments follow from the derivative recursions and the uniform lower bound on $\mu_t(\boldsymbol{\vartheta})$.

We first record the relevant kernel expressions. Define

$$A_{\boldsymbol{\nu}}(x) := 1 + x\frac{k_{\boldsymbol{\nu}}'(x)}{k_{\boldsymbol{\nu}}(x)}, \qquad B_{\boldsymbol{\nu}}(x) := \frac{\dot{k}_{\boldsymbol{\nu}}(x)}{k_{\boldsymbol{\nu}}(x)} = \frac{\partial}{\partial \boldsymbol{\nu}} \log k_{\boldsymbol{\nu}}(x).$$

For the scaled $F(1, \nu_1, \nu_2)$ density, let

$$b(\boldsymbol{\nu}) := \frac{\nu_1}{\nu_2 - 2}.$$

Since

$$\log k_{\boldsymbol{\nu}}(x) = C(\boldsymbol{\nu}) + \frac{\nu_1 - 2}{2}\log x - \frac{\nu_1 + \nu_2}{2}\log\left(1 + b(\boldsymbol{\nu})x\right),$$

where $C(\boldsymbol{\nu})$ does not depend on x, we have

$$\frac{k_{\boldsymbol{\nu}}'(x)}{k_{\boldsymbol{\nu}}(x)} = \frac{\nu_1 - 2}{2x} - \frac{\nu_1 + \nu_2}{2}\frac{b(\boldsymbol{\nu})}{1 + b(\boldsymbol{\nu})x}.$$

Therefore,

$$A_{\boldsymbol{\nu}}(x) = 1 + x\frac{k'_{\boldsymbol{\nu}}(x)}{k_{\boldsymbol{\nu}}(x)} = \frac{\nu_1}{2} - \frac{\nu_1+\nu_2}{2}\frac{b(\boldsymbol{\nu})x}{1+b(\boldsymbol{\nu})x}. \tag{A.8}$$

Since $V \subset (0,\infty)\times(4,\infty)$ is compact for the asymptotic normality result,

$$\sup_{\boldsymbol{\nu}\in V}\sup_{x>0}|A_{\boldsymbol{\nu}}(x)| < \infty. \tag{A.9}$$

Moreover,

$$xA'_{\boldsymbol{\nu}}(x) = -\frac{\nu_1+\nu_2}{2}\frac{b(\boldsymbol{\nu})x}{\{1+b(\boldsymbol{\nu})x\}^2},$$

so that

$$\sup_{\boldsymbol{\nu}\in V}\sup_{x>0}|xA'_{\boldsymbol{\nu}}(x)| < \infty. \tag{A.10}$$

Similarly, for

$$C(\boldsymbol{\nu}) = \log\Gamma\left(\frac{\nu_1+\nu_2}{2}\right) - \log\Gamma\left(\frac{\nu_1}{2}\right) - \log\Gamma\left(\frac{\nu_2}{2}\right) + \frac{\nu_1}{2}\log\left(\frac{\nu_1}{\nu_2-2}\right),$$

we have

$$B_{\nu_j}(x) := \frac{\partial}{\partial\nu_j}\log k_{\boldsymbol{\nu}}(x), \qquad j=1,2, \tag{A.11}$$

and

$$B_{\nu_1}(x) = \frac{\partial C(\boldsymbol{\nu})}{\partial\nu_1} + \frac{1}{2}\log x - \frac{1}{2}\log\big(1+b(\boldsymbol{\nu})x\big) - \frac{\nu_1+\nu_2}{2}\frac{x\,\frac{\partial b(\boldsymbol{\nu})}{\partial\nu_1}}{1+b(\boldsymbol{\nu})x},$$

and

$$B_{\nu_2}(x) = \frac{\partial C(\boldsymbol{\nu})}{\partial\nu_2} - \frac{1}{2}\log\big(1+b(\boldsymbol{\nu})x\big) - \frac{\nu_1+\nu_2}{2}\frac{x\,\frac{\partial b(\boldsymbol{\nu})}{\partial\nu_2}}{1+b(\boldsymbol{\nu})x}.$$

Here

$$\frac{\partial b(\boldsymbol{\nu})}{\partial\nu_1} = \frac{1}{\nu_2-2}, \qquad \frac{\partial b(\boldsymbol{\nu})}{\partial\nu_2} = -\frac{\nu_1}{(\nu_2-2)^2}.$$

Since V is compact and bounded away from $\nu_2 = 2$, the derivatives of $C(\boldsymbol{\nu})$ and $b(\boldsymbol{\nu})$ are uniformly bounded on V. Thus $B_{\nu_1}(x)$ and $B_{\nu_2}(x)$ consist only of bounded coefficient terms, logarithmic terms, and fractions of the form $x/(1+b(\boldsymbol{\nu})x)$ for $x>0$ and $b(\boldsymbol{\nu})>0$.

Direct differentiation of the scaled F log-density gives

$$\sup_{\boldsymbol{\nu}\in V}\sup_{x>0}\left\|x\frac{\partial B_{\boldsymbol{\nu}}(x)}{\partial x}\right\|<\infty, \tag{A.12}$$

and

$$\sup_{\boldsymbol{\nu}\in V}\left\|\frac{\partial B_{\boldsymbol{\nu}}(x)}{\partial \boldsymbol{\nu}^\top}\right\|\le C\{1+|\log x|+\log(1+b(\boldsymbol{\nu})x)\}.$$

The logarithmic terms are integrable under the scaled F distribution.

We also record the required moment bounds for the conditional mean derivatives. Since ω is bounded away from zero on the compact parameter space and all terms in the conditional mean recursion are nonnegative, there exists $\underline{\omega}>0$ such that

$$\mu_t(\boldsymbol{\vartheta})\ge\underline{\omega}\qquad\text{uniformly in } t \text{ and } \boldsymbol{\vartheta}\in K. \tag{A.13}$$

Hence,

$$\left\|\frac{\mu_t'(\boldsymbol{\vartheta})}{\mu_t(\boldsymbol{\vartheta})}\right\|_K^2\le\underline{\omega}^{-2}\|\mu_t'\|_K^2,\qquad\left\|\frac{\mu_t''(\boldsymbol{\vartheta})}{\mu_t(\boldsymbol{\vartheta})}\right\|_K\le\underline{\omega}^{-1}\|\mu_t''\|_K.$$

The derivative recursions are stable by the contraction argument used in the verification of Condition M.4. Under the strengthened second-moment condition imposed for asymptotic normality, in particular $\inf_{\boldsymbol{\nu}\in V}\nu_2>4$, the stationary HAR-MEM recursion has the required second moments. Therefore,

$$\mathbb{E}\|\mu_t'\|_K^2<\infty,\qquad\mathbb{E}\|\mu_t''\|_K<\infty, \tag{A.14}$$

and thus

$$\mathbb{E}\left[\left\|\frac{\mu_t'}{\mu_t}\right\|_K^2\right]<\infty,\qquad\mathbb{E}\left[\left\|\frac{\mu_t''}{\mu_t}\right\|_K\right]<\infty.$$

The $\boldsymbol{\vartheta\vartheta}$-block. From

$$\ell_t(\boldsymbol{\theta})=-\log\mu_t(\boldsymbol{\vartheta})+\log k_{\boldsymbol{\nu}}\left(\varepsilon_t(\boldsymbol{\vartheta})\right),\qquad\varepsilon_t(\boldsymbol{\vartheta})=\frac{\lambda_t}{\mu_t(\boldsymbol{\vartheta})},$$

we have

$$\frac{\partial\ell_t(\boldsymbol{\theta})}{\partial\boldsymbol{\vartheta}}=-A_{\boldsymbol{\nu}}\left(\varepsilon_t(\boldsymbol{\vartheta})\right)\frac{\mu_t'(\boldsymbol{\vartheta})}{\mu_t(\boldsymbol{\vartheta})}.$$

Differentiating once more gives

$$\frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\vartheta} \partial \boldsymbol{\vartheta}^\top} = \left[A_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right) + \varepsilon_t(\boldsymbol{\vartheta}) A'_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right) \right] \frac{\mu'_t(\boldsymbol{\vartheta}) \mu'_t(\boldsymbol{\vartheta})^\top}{\mu_t(\boldsymbol{\vartheta})^2} \\ - A_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right) \frac{\mu''_t(\boldsymbol{\vartheta})}{\mu_t(\boldsymbol{\vartheta})}.$$

Using the uniform boundedness of $A_{\boldsymbol{\nu}}$ and $x A'_{\boldsymbol{\nu}}(x)$, together with the moment bounds above,

$$\mathbb{E} \left[\left\| \frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\vartheta} \partial \boldsymbol{\vartheta}^\top} \right\|_\Theta \right] < \infty.$$

The $\boldsymbol{\vartheta\nu}$-block. The score with respect to $\boldsymbol{\nu}$ is

$$\frac{\partial \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\nu}} = B_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right).$$

Therefore,

$$\begin{aligned} \frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\vartheta} \partial \boldsymbol{\nu}^\top} &= \frac{\partial}{\partial \boldsymbol{\vartheta}} B_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right) \\ &= \frac{\partial B_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right)}{\partial \varepsilon} \frac{\partial \varepsilon_t(\boldsymbol{\vartheta})}{\partial \boldsymbol{\vartheta}} \\ &= -\varepsilon_t(\boldsymbol{\vartheta}) \frac{\partial B_{\boldsymbol{\nu}} \left(\varepsilon_t(\boldsymbol{\vartheta}) \right)}{\partial \varepsilon} \frac{\mu'_t(\boldsymbol{\vartheta})}{\mu_t(\boldsymbol{\vartheta})}. \end{aligned}$$

Since

$$\sup_{\boldsymbol{\nu} \in V} \sup_{x > 0} \left\| x \frac{\partial B_{\boldsymbol{\nu}}(x)}{\partial x} \right\| < \infty,$$

and since

$$\mathbb{E} \left[\left\| \frac{\mu'_t}{\mu_t} \right\|_K \right] < \infty,$$

it follows that

$$\mathbb{E} \left[\left\| \frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\vartheta} \partial \boldsymbol{\nu}^\top} \right\|_\Theta \right] < \infty.$$

The $\boldsymbol{\nu\nu}$-block. Finally,

$$\frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\nu} \partial \boldsymbol{\nu}^\top} = \frac{\partial B_{\boldsymbol{\nu}}\left(\varepsilon_t(\boldsymbol{\vartheta})\right)}{\partial \boldsymbol{\nu}^\top}.$$

The bound above gives

$$\left\| \frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\nu} \partial \boldsymbol{\nu}^\top} \right\|_\Theta \leq C \left\{ 1 + \| \log \varepsilon_t(\boldsymbol{\vartheta}) \|_K + \sup_{\boldsymbol{\theta} \in \Theta} \log \left(1 + b(\boldsymbol{\nu}) \varepsilon_t(\boldsymbol{\vartheta})\right) \right\}.$$

Since

$$\varepsilon_t(\boldsymbol{\vartheta}) = \frac{\lambda_t}{\mu_t(\boldsymbol{\vartheta})},$$

and $\mu_t(\boldsymbol{\vartheta})$ is uniformly bounded away from zero, the logarithmic terms are bounded by the logarithmic terms already treated in the consistency proof and by $C\lambda_t$ using $\log(1+x) \leq x$. Hence their expectations are finite, and therefore

$$\mathbb{E}\left[\left\| \frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\nu} \partial \boldsymbol{\nu}^\top} \right\|_\Theta \right] < \infty.$$

Combining the three blockwise bounds yields

$$\mathbb{E}\left[\left\| \frac{\partial^2 \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^\top} \right\|_\Theta \right] < \infty.$$

We also verify the second moment of the score at the true parameter. At $\boldsymbol{\theta}_0$,

$$\frac{\partial \ell_t(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\vartheta}} = -A_{\boldsymbol{\nu}_0}(\varepsilon_t) \frac{\mu_t'(\boldsymbol{\vartheta}_0)}{\mu_t(\boldsymbol{\vartheta}_0)}, \qquad \frac{\partial \ell_t(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\nu}} = B_{\boldsymbol{\nu}_0}(\varepsilon_t).$$

Since $A_{\boldsymbol{\nu}_0}$ is bounded and

$$\mathbb{E}\left[\left\| \frac{\mu_t'(\boldsymbol{\vartheta}_0)}{\mu_t(\boldsymbol{\vartheta}_0)} \right\|^2 \right] < \infty,$$

we have

$$\mathbb{E}\left[\left\| \frac{\partial \ell_t(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\vartheta}} \right\|^2 \right] < \infty.$$

Moreover, the components of $B_{\boldsymbol{\nu}_0}(\varepsilon_t)$ (see Eq. (A.11)) are bounded by logarithmic terms in ε_t and $\log(1 + b(\boldsymbol{\nu}_0)\varepsilon_t)$. These terms have finite second moments under the scaled F

distribution. Hence

$$\mathbb{E}\left[\left\|\frac{\partial \ell_t(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\nu}}\right\|^2\right] < \infty.$$

Consequently,

$$\mathbb{E}\left[\left\|\frac{\partial \ell_t(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}}\right\|^2\right] < \infty.$$

We also record the score-envelope bound used in the verification of Condition M.9. Viewing the score as a function of the filter arguments (μ_t, μ_t'), the mean-value theorem, implies that, for all sufficiently large t, almost surely,

$$\begin{aligned}
&\sup_{\boldsymbol{\theta}\in\Theta}\left\|\frac{\partial \hat{\ell}_t(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} - \frac{\partial \ell_t(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right\| \\
&\quad \le \sup_{\boldsymbol{\theta}\in\Theta}\left\| \frac{\partial}{\partial(\mu,\mu')}\left[\frac{\partial \ell_t(\boldsymbol{\theta};\mu,\mu')}{\partial \boldsymbol{\theta}}\right]\Bigg|_{(\mu,\mu')=(\tilde{\mu}_t,\tilde{\mu}_t')} \begin{pmatrix}\hat{\mu}_t-\mu_t\\ \hat{\mu}_t'-\mu_t'\end{pmatrix}\right\| \\
&\quad \le \sup_{\boldsymbol{\theta}\in\Theta}\left\| \frac{\partial}{\partial(\mu,\mu')}\left[\frac{\partial \ell_t(\boldsymbol{\theta};\mu,\mu')}{\partial \boldsymbol{\theta}}\right]_{(\mu,\mu')=(\tilde{\mu}_t,\tilde{\mu}_t')}\right\| \left(\|\hat{\mu}_t-\mu_t\|_K + \|\hat{\mu}_t'-\mu_t'\|_K\right) \\
&\quad = \sup_{\boldsymbol{\theta}\in\Theta}\left\| \begin{pmatrix} \dfrac{A_{\boldsymbol{\nu}}(\tilde{x}_t)+\tilde{x}_t A_{\boldsymbol{\nu}}'(\tilde{x}_t)}{\tilde{\mu}_t^2}\tilde{\mu}_t' & -\dfrac{A_{\boldsymbol{\nu}}(\tilde{x}_t)}{\tilde{\mu}_t} I \\ -\dfrac{\tilde{x}_t B_{\boldsymbol{\nu}}'(\tilde{x}_t)}{\tilde{\mu}_t} & 0 \end{pmatrix}\right\| \left(\|\hat{\mu}_t-\mu_t\|_K + \|\hat{\mu}_t'-\mu_t'\|_K\right) \\
&\quad \le C\left[\frac{|A_{\boldsymbol{\nu}}(\tilde{x}_t)|}{\tilde{\mu}_t} + \frac{|A_{\boldsymbol{\nu}}(\tilde{x}_t)+\tilde{x}_t A_{\boldsymbol{\nu}}'(\tilde{x}_t)|}{\tilde{\mu}_t^2}\|\tilde{\mu}_t'\| + \frac{\|\tilde{x}_t B_{\boldsymbol{\nu}}'(\tilde{x}_t)\|}{\tilde{\mu}_t}\right]\left(\|\hat{\mu}_t-\mu_t\|_K + \|\hat{\mu}_t'-\mu_t'\|_K\right) \\
&\quad \le C\left(1+\|\mu_t'\|_K + \|\hat{\mu}_t'-\mu_t'\|_K\right)\left(\|\hat{\mu}_t-\mu_t\|_K + \|\hat{\mu}_t'-\mu_t'\|_K\right) \\
&\quad = L_t\left(\|\hat{\mu}_t-\mu_t\|_K + \|\hat{\mu}_t'-\mu_t'\|_K\right), \qquad L_t := C(1+\|\mu_t'\|_K + \|\hat{\mu}_t'-\mu_t'\|_K).
\end{aligned} \tag{A.15}$$

Here $(\tilde{\mu}_t, \tilde{\mu}_t')$ denotes an intermediate value on the line segment between (μ_t, μ_t') and $(\hat{\mu}_t, \hat{\mu}_t')$, and $\tilde{x}_t = \lambda_t/\tilde{\mu}_t$. For each $\boldsymbol{\vartheta} \in K$,

$$\tilde{\mu}_t'(\boldsymbol{\vartheta}) = \mu_t'(\boldsymbol{\vartheta}) + a_t(\boldsymbol{\vartheta})\{\hat{\mu}_t'(\boldsymbol{\vartheta}) - \mu_t'(\boldsymbol{\vartheta})\}, \qquad a_t(\boldsymbol{\vartheta}) \in [0,1],$$

so that

$$\|\tilde{\mu}_t'\|_K \leq \|\mu_t'\|_K + \|\hat{\mu}_t' - \mu_t'\|_K.$$

The bound following the displayed derivative matrix is obtained by taking a finite-dimensional matrix norm and collecting the absolute values of the derivative blocks. The terms involving $A_{\boldsymbol{\nu}}(\tilde{x}_t)$, $A_{\boldsymbol{\nu}}(\tilde{x}_t) + \tilde{x}_t A_{\boldsymbol{\nu}}'(\tilde{x}_t)$, and $\tilde{x}_t B_{\boldsymbol{\nu}}'(\tilde{x}_t)$ are uniformly bounded by (A.9), (A.10), and (A.12), while $\tilde{\mu}_t^{-1}$ and $\tilde{\mu}_t^{-2}$ are uniformly bounded by (A.13). The $B_{\boldsymbol{\nu}}'$-term contains no filter derivative and is therefore absorbed into the constant part of the envelope.

The displayed envelope may be written as

$$L_t = C\left(1 + \|\mu_t'\|_K + \|\hat{\mu}_t' - \mu_t'\|_K\right).$$

Since $\|\hat{\mu}_t' - \mu_t'\|_K \to 0$ exponentially almost surely, this envelope is eventually dominated almost surely by the stationary envelope

$$\bar{L}_t = C(1 + \|\mu_t'\|_K).$$

Moreover, by (A.14),

$$E\bar{L}_0^2 \leq C\left(1 + E\|\mu_0'\|_K^2\right) < \infty.$$

Thus $E\log^+ \bar{L}_0 < \infty$ by Lemma 2.2 of Straumann and Mikosch (2006). Since finite initial observations are asymptotically irrelevant, the score approximation can be bounded using the stationary dominating envelope $\bar{L}_t$, and Lemma 2.1 of Straumann and Mikosch (2006) applies.

B. Empirical Appendix

B.1. Robustness check: Alternative portfolio compositions.

Table B.7: Random weights performance statistics.

This table shows performance metrics of the forecasted realized portfolio variance quantiles for the ReλHAR model and the F_{pfv} model. In particular, summary statistics of each model are shown for the distribution of $\hat{q}$ across different assumed portfolio compositions. The weights are randomly simulated from a dirichlet distribution with $\alpha_i = 10$, so that $w_i > 0$ and $\sum_{i=1}^{5} w_i = 1$ and approximately $0.05 < w_i < 0.5 \; \forall \; i = 1, \ldots, 5$. The results are based on the full out-of-sample period of $T = 4000$ observations.

	5%		2.5%		1%	
Metric	ReλHAR	F_{pfv}	ReλHAR	F_{pfv}	ReλHAR	F_{pfv}
Mean $\hat{q}$	0.050	0.055	0.027	0.032	0.013	0.016
SD $\hat{q}$	0.010	0.002	0.005	0.001	0.003	0.001
Mean tick-loss	0.407	0.441	0.277	0.320	0.164	0.206
Mean VolaR	7.715	6.528	9.223	7.768	11.460	9.596
UC non-rej.	55	78	74	2	64	0
Mean conditional VolaR	10.143	8.511	11.917	9.954	14.572	12.098
UES non-rej.	69	0	57	0	57	0

The table shows that even for other portfolio compositions, the ReλHAR model outperforms the consistently best-performing benchmark. The F_{pfv} model attains more non-rejections for the unconditional coverage test at the 5% level due to the lower estimation variance, but is unable to avoid rejections at 2.5%, 1% and for the conditional VolaR estimates. The larger (conditional) VolaR estimates of the ReλHAR model do, however, not lead to overestimation as the mean tick-loss is consistently significantly smaller. From these observations, we conclude that the ReλHAR model is sufficiently accurate for alternative portfolio compositions.

B.2. Robustness Check: Copula Contribution

The figure shows that the distribution is centered at a positive value of around 0.5%. This implies that on average the t-copula leads to slightly larger estimates of the 5% portfolio variance quantile forecasts. The histogram shows that the percentage differences are approximately normally distributed with variance close to one.

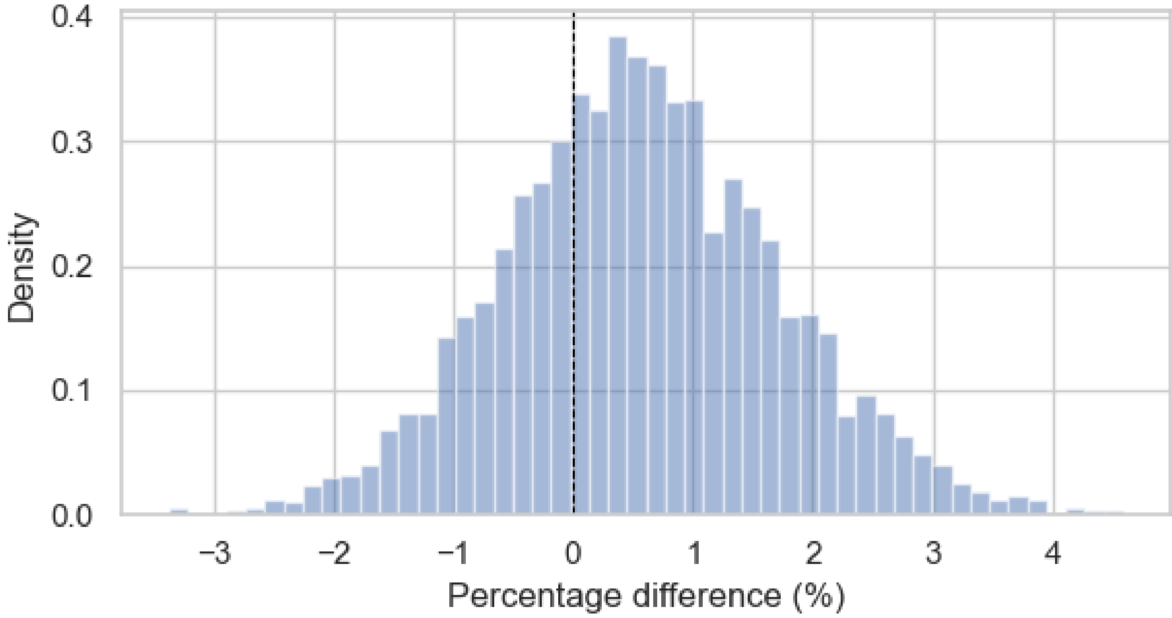

Figure B6: Portfolio Variance differences between copulas.

This figure shows the percentage differences in portfolio variance distribution quantile (q=0.05) forecasts between the two copula specifications. In particular, it shows the distribution of differences between the t-copula and the independence copula. ν is estimated for the full-sample at $\nu = 15.282$. The figure shows that the distribution of differences ranges between -3% and 4%, indicating that the t copula does lead to different portfolio variance forecasts. Furthermore, the distribution is centered around a small positive value of approximately 0.5%, indicating that the t copula on average leads to a larger 95% forecasted portfolio variance quantile.

References in Online Appendix

Berkes, I. and L. Horváth (2004). The efficiency of the estimators of the parameters in garch processes. *The Annals of Statistics 32*(2), 633–655.

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

Francq, C. and J.-M. Zakoïan (2019). *GARCH Models: Structure, Statistical Inference and Financial Applications*. Wiley.

Ranga Rao, R. (1962). Relations between weak and uniform convergence of measures with applications. *The Annals of Mathematical Statistics 33*(2), 659–680.

Straumann, D. (2005). *Estimation in conditionally heteroscedastic time series models*. Springer.

Straumann, D. and T. Mikosch (2006). Quasi-maximum-likelihood estimation in conditionally heteroscedastic time series: A stochastic recurrence equations approach. *Annals of Statistics 34*(5), 2449–2495.