Dynamic mixed frequency density pooling

Audronė Virbickaitė [*†] Hedibert F. Lopes [‡] Martina Danielova Zaharieva [§]

Abstract

This article investigates the benefits of combining information available from daily and intraday data to model and predict the dependence structure of equity returns. The two data sources are combined via a density pooling approach, wherein the individual joint densities are represented as a copula function, and the pooling weights are potentially time-varying. The dependence structure in the daily frequency case is extracted from a standard multivariate volatility model, and the high-frequency counterpart is based on the additive inverse Wishart model (AIW). We find that incorporating both high and low frequency information via density pooling provides significant gains in predictive model performance over any individual model and any model combination within the same data frequency. Finally, a portfolio allocation exercise quantifies the economic gains by producing investment portfolios with the smallest variance.

Keywords: Density combination; High Frequency; Mixture; Realized volatility.

JEL: C58, C11.

[*]Corresponding author

[†]Department of Quantitative Methods, CUNEF, Calle Pirineos 55, Madrid 28040, audrone.virbickaite@cunef.edu, Phone: 0034 914 480 892

[‡]School of Mathematical and Statistical Sciences at Arizona State University, USA, and Insper Institute of Education and Research, Sao Paulo, Brazil, Hedibert.Lopes@asu.edu / hedibertfl@insper.edu.br

[§]Department of Quantitative Methods, CUNEF, Calle Pirineos 55, Madrid 28040, martina.zaharieva@cunef.edu

1 Introduction

Since the advent of the availability of high frequency financial data, research on how to use, model and predict measures (such as volatility and co-volatility, for example) extracted from such data has surged (comprehensive review in McAleer & Medeiros 2008). As a result, high frequency data based models have proven to be powerful competitors to the standard modeling approaches, which are based on daily data. Some studies have shown that (co-)volatility models based on high-frequency data perform better at forecasting than models based solely on daily data (Andersen et al. 2003, Koopman et al. 2005, Horpestad et al. 2019, Lyócsa et al. 2021).

As an alternative to choosing a single modeling approach, some authors have combined the best of both worlds by augmenting low frequency models with high frequency information, see Engle 2002, Ghysels et al. 2004, 2005, Shephard & Sheppard 2010, Noureldin et al. 2012, Hansen et al. 2012, 2014 for univariate modeling, and Bauwens & Xu 2022 for a multivariate approach. Such combinations rely on suitable parametrization, in which the high frequency measure enters the model as an exogenous covariate, thus creating a new class of models.

In contrast to previous research, we combine low and high frequency information not through parameters but through the combination of densities. In particular, we model and predict the dependence structure of multiple financial returns as a weighted sum of two predictive densities, the first arising from low frequency data and the second arising from high frequency data. Such combinations are also known as opinion pools (the name was first proposed by Stone 1961). In principle, the proposed approach could be related to the System for Averaging Models (SAM) procedure of Norges Bank (Bjørnland et al. 2008, Aastveit et al. 2011) or the Bayesian predictive synthesis of McAlinn (2021) for combining macroeconomic forecasts. However, we pool models that arise from competing theoretical perspectives of understanding volatility, i.e. volatility as an unobserved process estimated from daily data vs the volatility as an observable quantity extracted from the high frequency data. This is in contrast to pooling several alternative models that differ exclusively in their parametric specifications.

The combination of predictive densities is a recent topic of increasing interest in the financial and macroeconomic literature (Hall & Mitchell 2007, Jore et al. 2010, Geweke & Amisano 2011, Billio et al. 2013, Aastveit et al. 2014, Del Negro et al. 2016, Aastveit et al. 2018, Bassetti et al. 2018,

McAlinn & West 2019, Casarin et al. 2023). Hall & Mitchell (2007) and Geweke & Amisano (2011) have relied on log predictive scores to calculate recursive combination weights, which, in the long term, reach a stable equilibrium. Del Negro et al. (2016) have also used the log scoring rule for modeling dynamic combination weights. Alternatively, Billio et al. (2013) have considered dynamic weights based on the model residuals and not on the log scores; Bassetti et al. (2018) have assumed random combination weights; and Casarin et al. (2023) have modeled the dynamics of the weighting process via nonlinear dynamic factor model. To control for the effects of the particular weighting scheme, we consider four options for density pooling, all based on the log predictive scores (LPS): equally weighted, static (Geweke & Amisano 2011), naïve dynamic (Jore et al. 2010) and dynamic (Del Negro et al. 2016). In the static weighting scheme, the weights are re-balanced daily as a function of the expanding set of the past LPS, converging to a stable equilibrium, hence the name static. In the naïve dynamic scheme, the set of the past LPS is smaller, and only the most recent observations are considered. Finally, in the "fully" dynamic scheme, the weights are latent and are updated as a function of the past weights via an AR-type process. For a general introduction and comprehensive reviews of aggregating probability distributions, readers are referred to Clemen & Winkler (2007) and Aastveit et al. (2018), among others. Timmermann (2018) has briefly reviewed the forecast combinations in the financial econometrics context.

Each of the individual densities in the pool, namely that modeled with low frequency and that modeled with high frequency data, are constructed by a copula function. Using a copula instead of the complete high-dimensional density is a convenient solution when the focus of the modeling is explicitly on the dependence structure rather than on the individual series dynamics. Modeling the dependence via a copula also has practical advantages. It allows to simplify the assessment of the marginal distributions and avoids dealing with highly parametrized and possibly nonstandard multivariate density functions. Models in which dynamic copula parameters are obtained from daily data are considered a standard approach in the financial times series literature (multivariate GARCH models in Dias & Embrechts 2004, Patton 2006*b*, Ausín & Lopes 2010; score-driven models in Koopman et al. 2018, Nguyen & Javed 2021; and factor models in Opschoor et al. 2021). In contrast, models in which the copula dependence structure is obtained from high frequency data are rather sparse (Salvatierra & Patton 2015, Fengler & Okhrin 2016, Okhrin & Tetereva 2017). Salvatierra & Patton (2015) have

modeled the dynamics of the copula parameter as a function of past realized correlations via an autoregressive score-type model, whereas Fengler & Okhrin (2016) and Okhrin & Tetereva (2017) have used multiple univariate AR-type processes for the realized covariances. In this work, we use several standard multivariate GARCH-type models for the copula parameter arising from daily data, and use an additive inverse Wishart model (Jin & Maheu 2013, 2016) for the copula parameter arising from the intra-day data. Flexible functional forms for the evolution of the copula parameter affect the joint density in similar ways to standard multivariate volatility modeling: dynamic specifications are able to model the conditional heteroscedasticity and produce more precise density forecasts than the static models (Patton 2006*a*, Tsay 2014).

In this article, we rely on a Bayesian estimation approach in three stages. In the first stage, the marginal distributions are estimated. In particular, the daily data are transformed to the unit interval by standardization of the de-meaned equity returns and application of the probability integral transform. This approach has the advantage of producing approximately standard normal marginals, which are very easy to handle. In the second stage, we fit a copula model to the resulting uniformly distributed data to obtain the joint predictive density for the returns. Finally, the density pooling weights assigned to a low and to a high frequency model are obtained in order to produce pooled predictive densities for the daily log returns.

We use daily and intraday US equity return data from 2001 to 2009 to construct multiple variants of competing density pools differing in (i) the respective underlying low and high frequency modeling strategy and (ii) the density pooling approach. Empirical results show that pooled models outperform the best individual model in terms of the entire density forecast as well in the left tail. In addition, the density pool shows improvement in predictive performance with respect to other mixed frequency models, such as the natural competitor — the DCC-HEAVY model of Bauwens & Xu (2022). Furthermore, the superiority of the pooled models is supported by different financial market data. In a second empirical data application consisting of exchange rates, market index and a commodity for the period from 2012 to 2021, we confirm the results in terms of model ranking. Finally, we perform a global minimum variance (GMV) portfolio allocation exercise to quantify the economic gains in using the proposed approach. The results illustrate the benefits of pooling by delivering investment portfolios with the smallest variance. Overall, we demonstrate that combining information from daily and intra-

day data sources not only produces superior joint density forecasts but also leads to tangible economic benefits for the investor.

The article is organized as follows. Sections 2 and 3 present the pooled copula model, estimation approach and model evaluation. Sections 4 and 5 contain two empirical applications: one for a ten-variate and another for a five-variate dataset. Finally, Section 6 presents the conclusion.

2 Methodology

In this section, we describe the main idea of the article: combining information arising from high and low frequency data in the copula modeling framework in order to model and predict the dependence structure of equity returns. The choice of modeling the joint distribution via copulas, next to being a flexible way of constructing multivariate densities, is also convenient from computational and methodological perspectives. As noted in Opschoor et al. (2021), when the cross-section dimension d is large, specifying and estimating the marginals separately might considerably ease the computational burden. In addition, such an approach enables a focus on modeling the dependence structure explicitly, independently from the marginals. In particular, we are interested in estimating a copula density $c(u_t|\mathcal{M}_{HF}, \mathcal{M}_{LF})$ for uniformly distributed data. The approach consists of three major tasks:

- Modeling the dynamics of the covariance matrices arising from low frequency data via a model called $\mathcal{M}_{LF}$. In this step, we consider standard specifications for multivariate co-volatilities, such as the Static model, the RiskMetrics© (RM) model of J.P. Morgan, or the dynamic conditional correlation (DCC) model of Tse & Tsui (2002), Engle (2002).

- Modeling the dynamics of the covariance matrices arising from high frequency data via a model called $\mathcal{M}_{HF}$. We concentrate on the additive inverse Wishart (AIW) approach of Jin & Maheu (2013, 2016).

- Modeling the dynamics of the combination weights ω_t. Here, we consider four options, covering a large part of the variety of linear combination strategies. The schemes are equally weighted, static (Geweke & Amisano 2011), naïve dynamic (Jore et al. 2010) and dynamic (Del Negro et al. 2016).

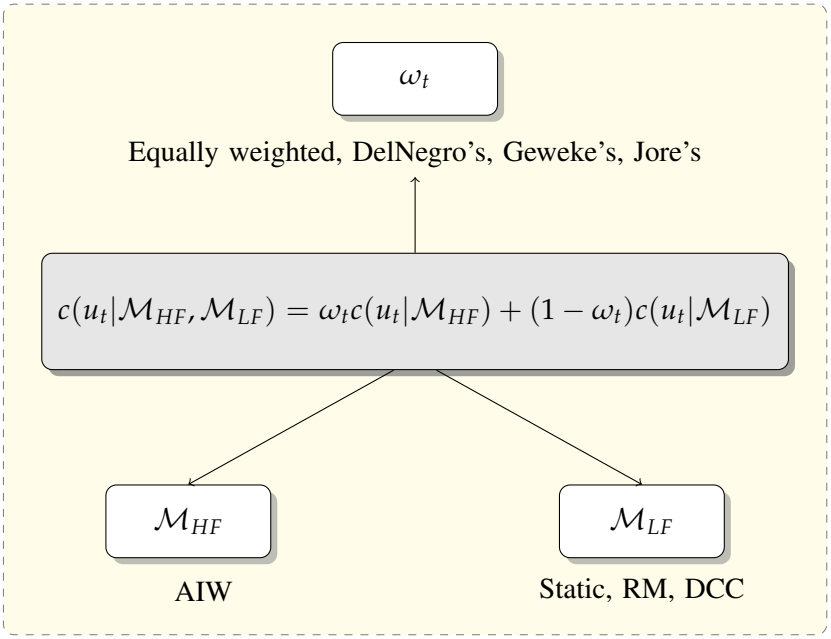

Figure 1: Model components combining high ($\mathcal{M}_{HF}$) and low ($\mathcal{M}_{LF}$) frequency models via density pooling with time varying weights ω_t.

Figure 1 summarizes the main model components and the approach used at each phase.

Even though both models, based on both high- and low- frequency data, are essentially aimed at capturing the dependence structure between the standardized returns, they exhibit very different properties. The low frequency data based models consider the entire series of historical daily data, and the estimated co-volatility processes are usually smooth. In contrast, the high frequency data based models can capture instantaneous changes in co-variation and make predictions accordingly; however, they are more likely to be "contaminated" by the market micro-structure noise. As noted in Kapetanios et al. (2015), Timmermann (2018), some models might be useful while the markets are in decline, whereas other models might be more informative when the markets are booming. Therefore, the time-varying pooling weights might also indicate whether the preference for one model or another is correlated with the overall market conditions. Such a correlation might affect an investor's decisions: if an investor anticipates bull/bear market conditions, they might choose to rely on one model or another to produce density forecasts.

We start by defining $r_{i,t}$ as the de-meaned log returns (in %) for day t and asset i, such that $t =$

$1, \ldots, T$ and $i = 1, \ldots, d$:

$$r_{i,t} = 100 \times \left(\log \frac{P_{i,t}}{P_{i,t-1}} - \mathrm{E}\left[\log \frac{P_{i,t}}{P_{i,t-1}} \right] \right),$$

where $P_{i,t-1}$ and $P_{i,t}$ are the prices at the beginning and at the end of the period, respectively, and $\mathrm{E}\left[\log \frac{P_{i,t}}{P_{i,t-1}}\right]$ is replaced with its sample equivalent.

Next, we present an approach to combine information arising from high and low frequency data for dependence modeling between daily financial returns, by relying on a density combination approach. As noted in Clemen & Winkler (2007), the two major approaches are linear and logarithmic pools. The linear opinion pool is a weighted linear combination of predictive probabilities, whereas multiplicative averaging results in a logarithmic opinion pool. Differently from linear pools, the logarithmic combinations have been shown to result in unimodal, less dispersed (Rufo et al. 2012) and symmetric (Kascha & Ravazzolo 2010) densities, in marked contrast to the empirically observed features of financial returns. Therefore, in this article, we rely on linear pools only because they appear to be more appropriate for financial time series.

The linear combination of individual densities obtained from models $\mathcal{M}$ is given by:

$$p(r_t) = \sum_{j=1}^{N} \omega_j p(r_t | \mathcal{M}_j), \quad t = 1, \ldots, T,$$

where $r_t = (r_{1,t}, \ldots, r_{d,t})'$ is the d-variate return vector, N is the number of alternative models, ω_j is the combination weight, and $p(r_t|\mathcal{M}_j)$ is the candidate density, originating from different models. The dependence structure between low frequency returns r_t can be modeled by combining (i) a model estimated from daily returns $p(r_t|\mathcal{M}_{LF})$ with (ii) a model estimated from high frequency returns $p(r_t|\mathcal{M}_{HF})$.

A convenient way to model the potentially high dimensional joint density $p(r_t|\mathcal{M}_j)$ involves separating the dependence structure from the dynamics of the marginals by using copula functions. Furthermore, the treatment of the marginal densities can be substantially simplified by taking advantage of the available *ex post* realized volatility measure defined by $RV_{i,t} = \sum_{j=1}^{J} \tilde{r}_{i,t,j}^2$. Here, $\tilde{r}_{i,t,j}$ is an l-minute log-return for day t, and J is the number of l-minute intervals in a trading day (Barndorff-Nielsen & Shephard 2002, Andersen et al. 2003, Barndorff-Nielsen & Shephard 2004). For a review of real-

ized volatility, readers are referred to McAleer & Medeiros (2008). To account for the time-varying volatility, the de-meaned log returns are standardized by the realized volatility measure, and some unconditional standard deviation[1] $z_{i,t} = r_{i,t}/(\sqrt{RV_{i,t}} \cdot \sigma_i)$. As seen in Andersen et al. (2000, 2001), is it safe to assume that $z_{i,t} \sim N(0,1)$.[2] Finally, call $u_{i,t} = \Phi_1(z_{i,t})$ the probability integral transform of the $z_{i,t}$, where $\Phi_1(\cdot)$ is a cumulative distribution function for the univariate standard normal distribution, and the resulting variables are uniformly distributed $u_{i,t} \overset{iid}{\sim} \mathcal{U}(0,1) \forall i = 1, \ldots, d$ (serially uncorrelated). This procedure helps reduce the number of parameters and the computational burden of the estimation procedure. Moreover, using the realized volatility for the standardization step circumvents the inclusion of an additional potential source of estimation error.

The dependence structure of the resulting probability integral transforms can be easily modeled by using copulas. To define a copula, we consider a collection of random variables $Y_1, \ldots, Y_d$ with corresponding distribution functions $F_i(y_i) = P[Y_i \leq y_i]$ for $i = 1 \ldots, d$ and a joint distribution function $H(y_1, \ldots, y_d) = P[Y_1 \leq y_1, \ldots, Y_d \leq y_d]$. Then, according to a theorem by Sklar (1959), a copula C exists such that

$$H(y_1, \ldots, y_d) = C(F_1(y_1), \cdots, F_d(y_d)).$$

That is, the dependence structure can be separated from the marginals. The joint density $h(y_1, \ldots, y_d)$ is then a product of individual marginal densities $f_i(y_i)$ and a copula density:
$h(y_1, \ldots, y_d) = c(F_1(y_1), \cdots, F_d(y_d)) \cdot \prod_{i=1}^{d} f_i(y_i)$. Copulas are defined in the unit hypercube $[0,1]^d$, where d is the dimension of the data, and all univariate marginals are uniformly distributed. For a detailed treatment of copulas and areas of applications, readers are referred to McNeil et al. (2005), Nelsen (2006), Patton (2012), Joe (2015).

In this article, we use Gaussian and t copulas, because they are available in high dimensions ($d > 2$), and their implementation is straightforward. Gaussian copulas, although widely used, do not allow for fat-tailed co-dependence — an assumption that can be relaxed by using the t copula. Nonetheless, one could also consider even more flexible vine copulas (Brechmann & Czado 2015, Loaiza-Maya &

[1] σ_i is a scaling factor that allows the standard deviation of the returns to deviate from the RV measure; see Jin & Maheu (2013, 2016).

[2] Andersen et al. (2000, 2001) have found that the distributions of the returns scaled by realized standard deviations are approximately Gaussian.

Smith 2018) or inversion copulas (Demarta & McNeil 2005, Smith et al. 2012, Loaiza-Maya & Smith 2020), for example, which are also available in higher dimensions and can capture stylized features observed in financial time series, such as heteroscedasticity, pair-wise fat tails and asymmetry.

Call $u_t = (u_{1,t}, \ldots, u_{d,t})'$ the collection of uniformly distributed data at time t. The d-variate Gaussian copula has the following distribution and density functions (Joe 2015):

$$C(u_t|\boldsymbol{R}) = \Phi_d(\Phi_1^{-1}(u_{1,t}), \ldots, \Phi_1^{-1}(u_{d,t})|\boldsymbol{R}),$$
$$c(u_t|\boldsymbol{R}) = \frac{\phi_d(\Phi_1^{-1}(u_{1,t}), \ldots, \Phi_1^{-1}(u_{d,t})|\boldsymbol{R})}{\prod_{i=1}^{d} \phi_1(\Phi_1^{-1}(u_{i,t}))}.$$

Here, $\Phi_d(\cdot|\boldsymbol{R})$ and $\phi_d(\cdot|\boldsymbol{R})$ are a d-variate standard normal distribution and density functions with a correlation matrix $\boldsymbol{R}$. The d-variate t copula has the following distribution and density functions (Joe 2015):

$$C(u_t|\boldsymbol{R}, \eta) = T_{d,\eta}(T_{1,\eta}^{-1}(u_{1,t}), \ldots, T_{1,\eta}^{-1}(u_{d,t})|\boldsymbol{R}),$$
$$c(u_t|\boldsymbol{R}, \eta) = \frac{t_{d,\eta}(T_{1,\eta}^{-1}(u_{1,t}), \ldots, T_{1,\eta}^{-1}(u_{d,t})|\boldsymbol{R})}{\prod_{i=1}^{d} t_{1,\eta}(T_{1,\eta}^{-1}(u_{i,t}))}.$$

Here, $T_{1,\eta}$, $T_{d,\eta}(\cdot|\boldsymbol{R})$, $t_{1,\eta}$ and $t_{d,\eta}(\cdot|\boldsymbol{R})$ are the univariate and d-variate t distribution and density functions with degrees of freedom parameter $\eta > 0$ and correlation matrix $\boldsymbol{R}$. When $\eta \to \infty$, the t copula becomes a Gaussian copula.

Another important reason exists to focus on Gaussian copula, at least for the high-frequency model. We make use of the fact that the variance-covariance (or correlation) matrix, estimated from the de-meaned and standardized log returns $z_{i,t}$ (given that they are approximately normally distributed), is equivalent to the copula parameter $\boldsymbol{R}$. This result is valid for only a Gaussian copula with standard normal marginals and is a result of Hoeffding's lemma and Sklar's theorem; details are described in Fengler & Okhrin (2016). This aspect can be easily seen in the bi-variate case. If we consider two random variables Y_1 and Y_2 with standard normal marginal distributions $F_1(y_1)$ and $F_2(y_2)$ and a joint distribution function $H(y_1, y_2)$, we can find the covariance $Cov(y_1, y_2) = \sigma_{12}$ (also a correlation for

the standardized data) by inserting the copula function in Hoeffding's lemma:

$$\begin{aligned}\sigma_{12} &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} [H(y_1, y_2) - F_1(y_1)F_2(y_2)]\, dy_1 dy_2 \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} [C(F_1(y_1), F_2(y_2)|\boldsymbol{R}) - F_1(y_1)F_2(y_2)]\, dy_1 dy_2 \\ &= \kappa.\end{aligned}$$

As a special case, if the two random variables Y_1 and Y_2 are independent, we obtain $H(y_1, y_2) = F_1(y_1)F_2(y_2)$ and zero covariance. That is, the difference between those two integrals is a measure of the linear dependence between the two variables. When $C(F_1(y_1), F_2(y_2)|\boldsymbol{R})$ corresponds to the Gaussian copula, defined above, the resulting scalar κ is just the off-diagonal element of the matrix $\boldsymbol{R}$. Therefore, the Gaussian copula function enables one-to-one mapping between the copula dependence parameter and the linear dependence measure. This result is relevant to our work because it enables use of the realized correlation, obtained from high frequency data, as a copula parameter $\boldsymbol{R}$.

Finally, given our proposed framework, the resulting density to describe daily dependence structure at time t can be written in terms of a copula density pool:

$$c(u_t|\mathcal{M}_{LF}, \mathcal{M}_{HF}) = \omega_t c(u_t|\mathcal{M}_{HF}) + (1 - \omega_t)c(u_t|\mathcal{M}_{LF}), \tag{1}$$

where $\mathcal{M}_{HF}$ and $\mathcal{M}_{LF}$ present the models estimated by using high and low frequency data. For example, consider JP Morgan (JPM) and Bank of America (BAC) log return series, used later in the empirical application. The left panel of Figure 2 shows the realized correlations (in gray), sample correlation (thick black horizontal line) and rolling-window correlations (a window of 50, black line); the middle and left panels show different dependence structures for calm and more volatile periods between standardized BAC/JPM returns and pairwise estimated t copula degrees of freedom. As shown, the dependence structure can be described by using multiple alternative measures, both static and dynamic. Moreover, the dependence structure changes not only in strength but also in tail thickness. The proposed mixed frequency pooled copula would allow for all these features observed in real data.

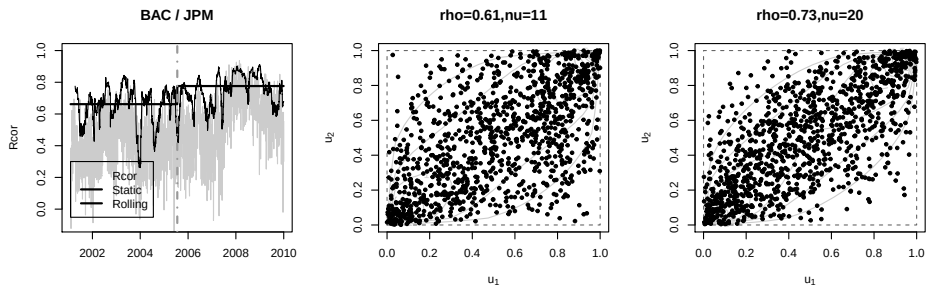

Figure 2: Left: BAC and JPM realized correlations, sample correlation (for two sub-samples) and rolling-window correlation. Middle and right: different dependence structures for the first half of the sample (2001/02/01 to 2005/07/20) and the second half of the sample (2005/07/21 to 2009/12/31) between the standardized BAC/JPM returns and pairwise estimated t copula degrees of freedom.

2.1 Low frequency covariance modeling

Next, we consider several standard approaches to model the dynamics of the correlation matrix that arises from the low frequency (daily) data. Call Ω a variance-covariance matrix of the observed standard normally distributed standardized returns $z_t = (z_{1,t}, \ldots, z_{d,t})'$. Then the corresponding correlation matrix is $\boldsymbol{R} = (\text{diag}\Omega)^{-1/2}\Omega(\text{diag}\Omega)^{-1/2}$. We start with the most straightforward way to measure dependence, by using a sample correlation matrix. The dependence between u_t is modeled either by fitting a Gaussian or t copula with a static correlation matrix $\boldsymbol{R}$, estimated given the daily data up to time t. Another possible model, which is dynamic, is the exponentially weighted moving average (EWMA) specification, popularized by the RiskMetrics© model. The persistence parameter for the RM model can be either fixed ($\lambda = 0.94$, as recommended for daily data) or estimated:

$$\Omega_t = (1-\lambda) z_{t-1} z_{t-1}' + \lambda \Omega_{t-1}.$$

This model is very simple and easy to justify: the covariance at time t depends on the previous period's covariance adjusted by the most recent shock. Finally, the EWMA model can be generalized to include an intercept term, thus resulting in the DCC model (Tse & Tsui 2002, Engle 2002):

$$\Omega_t = \overline{\Omega} \odot (\iota\iota' - A - B) + A \odot z_{t-1} z_{t-1}' + B' \odot \Omega_{t-1}, \tag{2}$$

where $\odot$ is the Hadamard product of two equally sized matrices (element-by-element multiplication); ι is a vector of ones; parameter matrices A, B can be replaced with scalars a, b; and $\overline{\Omega}$ is a sample variance covariance matrix. Naturally, the model choice for daily variance-covariance matrix is not limited to the models outlined above. For extensive reviews of existing multivariate volatility models, readers are referred to Asai et al. (2006), Bauwens et al. (2006), Silvennoinen & Teräsvirta (2009), among others.

2.2 High frequency covariance modeling

As mentioned before, Fengler & Okhrin (2016) have shown that the Gaussian copula's parameter $\boldsymbol{R}_t$ can be estimated by using the correlation matrix of the original data (log returns in our case). In the high frequency data setting, the correlation matrix of the returns can be estimated via Rcor_t, a realized correlation measure, obtained from intraday data (Noureldin et al. 2012):

$$\text{Rcor}_t = (\text{diag}\,\text{Rcov}_t)^{-1/2}\text{Rcov}_t(\text{diag}\,\text{Rcov}_t)^{-1/2},$$

where Rcov_t is a realized covariance measure. Modeling the dynamics of the realized covariance matrices is a notoriously difficult task because of the high dimensions and positive-definite restrictions on the matrices. One method is to decompose/transform the variance-covariance matrix and use standard time-series techniques to model the transformed series. This approach has been used by Bauer & Vorkink (2011) and Chiriac & Voev (2011), among others. Another method is to model the dynamics of the realized variance-covariance matrices directly by using Wishart distributions (Gourieroux et al. 2009, Jin & Maheu 2013, 2016). In Jin & Maheu (2013, 2016), the scale matrix in the Wishart distribution follows either an additive or a multiplicative component structure, and the authors have found that the additive structure performs better. Such additive models capture strong persistence in the covariances and fat-tailed distributions of the returns. They have compared their proposed model with multiple other models, such as Cholesky-VARFIMA from Chiriac & Voev (2011), the Wishart autoregressive model from Gourieroux et al. (2009), vec-MGARCH from Ding & Engle (2001) and DCC from Engle (2002). The additive Wishart model has been found to produce superior density forecasts for all forecast horizons.

Next, we present the additive component model, introduced in Jin & Maheu (2013, 2016). Consider a sequence of realized covariance matrices Rcov_t of dimension $d \times d$, $t = 1, \ldots, T$. The additive component inverse Wishart AIW(L) model is given by:

$$\begin{aligned} \text{Rcov}_t &\sim \mathcal{IW}((\nu - d - 1)V_t, \nu), \\ V_t &= B_0 + \sum_{j=1}^{L} B_j \odot \Gamma_{t-1,l_j}, \\ B_j &= b_j b_j', \quad j = 1, \ldots, L, \\ \Gamma_{t-1,l_j} &= 1/l_j \sum_{i=1}^{l_j} \text{Rcov}_{t-i}. \end{aligned} \tag{3}$$

Here, $\mathcal{IW}(A, b)$ is the inverse-Wishart distribution with scale matrix A and degrees of freedom b. We set $l_1 = 1$; furthermore, l_js indicates how many past observations are used to form a component Γ_{t-1,l_j} and L is the number of autoregressive components. B_0 is a symmetric positive definite matrix and is set to $B_0 = (\iota\iota' - B_1 - \ldots - B_K) \odot \overline{\text{Rcov}}$ so that the long-term mean of the covariances is equal to the sample mean.

2.3 Choosing the weights

In this article, we attempt to cover a large part of the types of linear pooling schemes by focusing on four different approaches: equally weighted, static (Geweke & Amisano 2011), naïve dynamic (Jore et al. 2010) and dynamic (Del Negro et al. 2016).

Geweke & Amisano (2011) have proposed to maximize the log predictive score function at each point in time:

$$\begin{aligned} \omega_{T+k+1}^{Gew} &= \arg\max_{\omega} f(\omega), \quad \text{such that} \\ f(\omega) &= \sum_{t=1}^{T+k} \log[\omega c(u_t|\mathcal{M}_{HF}) + (1-\omega)c(u_t|\mathcal{M}_{LF})], \end{aligned} \tag{4}$$

where $c(u_t|\mathcal{M}_{HF})$ and $c(u_t|\mathcal{M}_{LF})$ are predictive copula densities for u_t, and $k = 1, \ldots, K$ is the out of sample evaluation period. Even though the weights are recalculated at each time point, this weighting scheme is considered static because, for a large K, the weights reach a stable equilibrium

(Del Negro et al. 2016). Another approach involves using the log-score rolling weights calculated at each time t by using $\tilde{m}$ lags, as defined in Jore et al. (2010):

$$\omega^{Jore}_{T+k+1,\tilde{m}} = \frac{\exp[\sum_{\tau=T+k+1-\tilde{m}}^{T+k} \log c(u_\tau|\mathcal{M}_{HF})]}{\sum_{r=\{HF,LF\}} \exp[\sum_{\tau=T+k+1-\tilde{m}}^{T+k} \log c(u_\tau|\mathcal{M}_r)]}. \tag{5}$$

We call this a naïve time-varying weighting approach. The main difference between the weights in Eqs. (4) and (5) is that Geweke's approach considers the predictive densities from the entire sample, whereas Jore's weighting scheme places importance on only the last $\tilde{m}$ observations. Finally, as in Del Negro et al. (2016), we allow for persistence in weights by introducing a latent variable s_t, thus giving rise to a dynamic weighting scheme:

$$\begin{aligned} s_t &= \beta s_{t-1} + \sqrt{1-\beta^2}\xi_t, \quad \xi_t \sim \mathcal{N}(0,1) \\ \omega_t^{DN} &= \Phi(s_t). \end{aligned} \tag{6}$$

The unconditional mean of s_t is 0, and the unconditional variance is 1. Parameter β controls the persistence of the weight dynamics: when $\beta = 1$, the process reduces to a random walk; when $\beta = 0$, at each time t, the weights ω_t^{DN} will be uniformly distributed *a priori*.

For the sake of simplicity, in this article, we consider only a few of the available linear pools, because the main goal is to investigate the potential benefits of high and low frequency data combinations. By no means do we wish to present a horse-race among the pooling methods, given that a more flexible/advanced density combination method would probably result in better performing models. For example, a similar approach to ours has been proposed by McAlinn (2021), wherein macroeconomic data from different frequencies are synthesized by using Bayesian predictive synthesis (McAlinn & West 2019). This approach is particularly powerful and beneficial when models are dependent, as shown by Takanashi & McAlinn (2021). Finally, because the data are multivariate, one might also consider weighing each series separately (McAlinn et al. 2020). This route would potentially be beneficial if the marginal distributions form a part of the overall model. Therefore, the use of Bayesian predictive synthesis in our proposed modeling framework should definitely be pursued in a follow-up research agenda.

2.4 Competitor model

For comparison purposes, we also include a model that augments the low frequency DCC model with the high frequency information via the high-frequency-based volatility (HEAVY) approach (Noureldin et al. 2012), thus resulting in a scalar DCC-HEAVY specification of Bauwens & Xu (2022). In this model, the lagged outer product of standardized returns in Eq.(2) is replaced by the lagged realized correlations. More specifically, we can model the correlation matrix directly, without using the variance-covariance matrix, as follows:

$$\boldsymbol{R}_t = \bar{\boldsymbol{R}} + a(\text{Rcor}_{t-1} - \overline{\text{Rcor}}) + b(\boldsymbol{R}_{t-1} - \bar{\boldsymbol{R}}).$$

Here, $a, b \geq 0$, $b = 0$ if $a = 0$, $b < 1$; $\bar{\boldsymbol{R}}$ is the $d \times d$ sample correlation matrix of z_t; and $\overline{\text{Rcor}}$ is the sample mean of the realized correlation matrices. We consider the DCC-HEAVY model with Gaussian and t copulas. For more details, readers are referred to Bauwens & Xu (2022).

3 Posterior Inference and Model Selection

3.1 Posterior inference

For posterior inference and prediction, we rely on Bayesian computation, particularly Markov chain Monte Carlo (MCMC) methods. To estimate the density pool in Eq. (1), we first sample from the posterior of the individual models $\mathcal{M}_{HF}$ and $\mathcal{M}_{LF}$. Conditional on those samples, the density pooling weights in Eqs. (4)-(6) can be obtained. Next, we briefly describe the posterior sampling details for each of the models presented in Sections 2.1 and 2.2.

Static model. Consider an inverse-Wishart prior on the unconditional variance covariance matrix $\Omega \sim \mathcal{IW}(I_d(\nu_0 - d - 1), \nu_0)$, $\nu_0 \geq d + 1$, so that $E[\Omega] = I_d$ and I_d id the d-dimensional unit matrix. Given the observed standardized approximately normally distributed data $z_t = (z_{1,t}, \ldots, z_{d,t})'$, where $z_{1:T} = (z_1', \ldots, z_T')'$, the parameter Ω can be sampled directly from the posterior $\Omega | z_{1:T} \sim \mathcal{IW}(z_{1:T} z_{1:T}' + I_d(\nu_0 - d - 1), \nu_0 + T)$; derivations are provided in Appendix A in the Online Supple-

mentary Material[3]. The correlation matrix used as a copula parameter is obtained as $\boldsymbol{R} = (\text{diag}\Omega)^{-1/2}$ $\Omega(\text{diag}\Omega)^{-1/2}$.

RiskMetrics. The estimated RiskMetrics (RMe) model contains only one parameter, $\lambda \in (0,1)$. We assume a beta prior $\lambda \sim \mathcal{B}(a_\lambda, b_\lambda)$ so that $0 < \lambda < 1$. Given that the data are normally distributed, the likelihood can be easily written as a function of Ω_t (or $\boldsymbol{R}_t$). Our target density is the posterior $p(\lambda|u_t) \propto \prod_{t=1}^{T} c(u_t|\boldsymbol{R}_t)\pi(\lambda)$, which is of a non-standard form. Therefore, we can sample from $p(\lambda|u_t)$ via a random walk metropolis Hastings (RWMH) step, for $m = 1,\ldots,M$, where M is the length of the MCMC chain, and given some starting value $\lambda^{(0)}$:

1. At iteration m, draw a new value of $\tilde{\lambda}$ from a normal proposal distribution $\mathcal{N}(\lambda^{(m-1)}, V_\lambda)$.
2. Accept the new draw with probability $\alpha = \max\left\{1, p(\tilde{\lambda}|u_t)/p(\lambda|u_t)\right\}$.
3. If the draw is accepted, set $\lambda^{(m)} = \tilde{\lambda}$, if not, set $\lambda^{(m)} = \lambda^{(m-1)}$.

Tuning the parameter V_λ allows to control the acceptance ratio.

DCC, DCC-t and DCC-HEAVY. Similarly to the RMe model above, the parameters for the scalar DCC, DCC-t and DCC-HEAVY models (a,b,η) can be sampled via RWMH. The priors for the parameters (a,b) are assumed beta so that $0 < a,b < 1$, and the prior for the degrees of freedom of the Student t distribution, η, is exponential. We sample (a,b,η) jointly in one step from a trivariate normal proposal distribution given some starting values $(a,b,\eta)^{(0)}$. The algorithm iterates via MH steps, and we always reject the draws for which $a+b > 1$ (except in the DCC-HEAVY model, in which this restriction is not necessary), to ensure that the process is mean-reverting. For the DCC and DCC-HEAVY models with Gaussian copulas, we have only parameters (a,b).

AIW. For estimation of the AIW model, as in Jin & Maheu (2013, 2016), we use MH within Gibbs. We assume $L = 2$ and call $\boldsymbol{b} = (b_1', b_2')$. The priors for the model parameters are $\nu \sim \mathcal{E}_{\nu>d+1}(\xi_\nu)$, $\boldsymbol{b} \sim \mathcal{N}_{2d}(0, V_b \cdot I_{2d})$, $l_2 \sim \mathcal{U}_Z(a_l, b_l)$. Here $\mathcal{E}(\cdot)$ is an exponential distribution, $\mathcal{N}_{2d}(\cdot)$ is a $2d$-variate normal distribution, and $\mathcal{U}_Z(\cdot)$ is a discrete uniform distribution. Given some starting values $(l_2, \nu, \boldsymbol{b})^{(0)}$, the algorithm iterates through the following for $m = 1,\ldots,M$:

[3]Available at `https://sites.google.com/view/audravirbickaitephd`.

1. Sample ν via RWMH from the conditional posterior:
   $p(\nu|l_2, \boldsymbol{b}, \text{Rcov}_{1:T}) \propto \pi(\nu) \prod_t g_{IW}(\text{Rcov}_t|l_2, \nu, \boldsymbol{b})$, where g_{IW} is the density function of the inverse-Wishart distribution.

2. Sample $\boldsymbol{b} = (b_1', b_2')$ via RWMH jointly from the $2d$-variate normal proposal, where the first elements of each vector are truncated to be positive, for identification purposes. As in Jin & Maheu (2013, 2016) we reject such draws of $\boldsymbol{b}$ where B_0 is not positive definite, or the absolute value of any element of $\sum_{i=1}^{2} B_i$ is not less than 1.

3. Sample l_2 via RWMH by using Poisson increments that can be either positive or negative with equal probability.

Pooling weights. Estimation of the static and naïve time-varying weights is straightforward and can be performed by applying the formulas in Eqs. (4) and (5) on the log predictive scores at each MCMC iteration after the estimation is performed for all models individually. For the time-varying persistent weights ω_t^{DN}, we use a variant of particle MCMC called particle marginal Metropolis-Hastings sampler (Andrieu et al. 2010). In particular, we use a bootstrap filter of Gordon et al. (1993) for the latent state s_t filtering and a standard MH step with normal prior truncated at (-1,1) $\beta \sim \mathcal{TN}_{(-1,1)}(m_\beta, V_\beta)$ with a random walk proposal for the persistence parameter β.

3.2 Model selection

To compare model performance, we consider one-step-ahead density prediction. One-step-ahead horizon has also been considered by Billio et al. (2013), for example. For that purpose, we calculate the correlation matrices for $t+1$. For the static and fixed-parameter RiskMetrics (RMf) model, the marginal predictive is available analytically:

$$p^{\text{static}}(u_{t+1}|z_{1:t}) = x^{-1} t_{d,\nu_0+t-d+1}\left(z_{t+1}\left|\frac{I_d(\nu_0-d-1)+z_{1:t}'z_{1:t}}{\nu_0+t-d+1}\right.\right),$$
$$p^{\text{RMf}}(u_{t+1}|z_{1:t}) = x^{-1}\phi_d(z_{t+1}|\boldsymbol{R}_{t+1}(\lambda)),$$

where $z_{t+1} = (\Phi^{-1}(u_{1,t+1}), \ldots, \Phi^{-1}(u_{d,t+1}))'$, $x = \prod_{i=1}^{d} \phi_1(z_{i,t+1})$, and $\boldsymbol{R}_{t+1}(\lambda)$ is a correlation matrix from the RMf model with known parameter λ. Of note, the marginal predictive for the static

model is the Student t density, which is a result of the normal-inverse Wishart conjugacy.

For the DCC (DCC-HEAVY) and RMe models described in Section 2.1, the posterior predictive distributions are given by

$$
\begin{aligned}
p^{\mathrm{DCC}}(u_{t+1}|z_{1:t},\theta_{DCC}) &= x^{-1}\phi_d(z_{t+1}|\boldsymbol{R}_{t+1}(\theta_{DCC})),\\
p^{\mathrm{DCCt}}(u_{t+1}|z_{1:t},\theta_{DCCt}) &= \left(\prod_{i=1}^{d} t_{1,\eta}(T_{1,\eta}^{-1}(u_{i,t+1}))\right)^{-1} t_{d,\eta}(u_{t+1}|\boldsymbol{R}_{t+1}(\theta_{DCCt})),\\
p^{\mathrm{RMe}}(u_{t+1}|z_{1:t},\theta_{RMe}) &= x^{-1}\phi_d(z_{t+1}|\boldsymbol{R}_{t+1}(\theta_{RMe})).
\end{aligned}
$$

Here, $\theta_{DCC} = (\overline{\Omega}, a, b)$; $\theta_{DCCt} = (\overline{\Omega}, a, b, \eta)$; and $\theta_{RMe} = \lambda$ are the estimated parameters for the DCC (DCC-HEAVY), DCC-t (DCC-HEAVY-t) and RMe models.

Finally, the posterior predictive density for the AIW model is:

$$
p^{\mathrm{AIW}}(u_{t+1}|z_{1:t},\theta_{AIW}) = x^{-1} t_{d,v-d+1}\left(z_{t+1}\left|\frac{v-d-1}{v-d+1}V_{t+1}\right.\right),
$$

where θ_{AIW} is a vector of the parameters in the AIW model.

The marginal predictive densities $p(u_{t+1}|z_{1:t})$ that account for parameter uncertainty for the DCC, DCC-t, DCC-HEAVY, RMe and AIW models can be obtained by using the MCMC output:

$$
p(u_{t+1}|z_{1:t}) = \int p(u_{t+1}|z_{1:t},\theta)p(\theta|z_{1:t})d\theta \approx \frac{1}{M}\sum_{m=1}^{M} p(u_{t+1}|z_{1:t},\theta^{(m)}),
$$

where $(\theta^{(1)},\ldots,\theta^{(M)})$ are the M posterior samples obtained from the MCMC.

The model comparison is carried out via predictive Bayes factors (BF) given K out of sample observations. The BF between model 0 ($\mathcal{M}_0$) and model 1 ($\mathcal{M}_1$) is defined as (West 1986, Kass & Raftery 1995):

$$
BF_{T:T+K} = \frac{p(u_{T:T+K}|z_{1:T},\mathcal{M}_0)}{p(u_{T:T+K}|z_{1:T},\mathcal{M}_1)},
$$

where $p(u_{T:T+K}|z_{1:T},\mathcal{M}_r) = \prod_{k=1}^{K} p(u_{T+k}|z_{1:T+k-1},\mathcal{M}_r)$. The exact calculation of $p(u_{T:T+K}|z_{1:T},\mathcal{M}_r)$ is time consuming because of an expanding time horizon, i.e., the model must be re-estimated K times.

For notational convenience, we do not condition on the model $\mathcal{M}_r$ and instead of condition on $z_{1:T}$, because $u_{i,t} = \Phi(z_{i,t})$. Then we can write:

$$\begin{aligned} p(u_{T:T+K}|u_{1:T}) &= \prod_{k=1}^{K} p(u_{T+k}|u_{1:T+k-1}) \\ &= \prod_{k=1}^{K} \int p(u_{T+k}|\theta) p(\theta|u_{1:T+k-1}) d\theta \\ &\overset{\text{T large}}{\approx} \prod_{k=1}^{K} \int p(u_{T+k}|\theta) \hat{p}(\theta) d\theta, \quad \text{where} \quad \theta^{(1)}, \ldots, \theta^{(M)} \sim \hat{p}(\theta), \\ &\approx \prod_{k=1}^{K} \frac{1}{M} \sum_{m=1}^{M} p(u_{T+k}|\theta^{(m)}). \end{aligned}$$

The marginal predictive distribution of $u_{T:T+K}$ can be approximated by using a posterior sample of estimated model parameters $\hat{p}(\theta)$ until time T (instead of re-estimating the model K times).

Another necessary measure used for calculating the pooling weights is the log predictive score (LPS):

$$LPS = \sum_{t=T}^{T+K-1} \log p(u_{t+1}|z_{1:t}). \tag{7}$$

Finally, we also compare the predictive model performance for the lower q^* percentile. Similar metrics have also been considered by Delatola & Griffin (2011) and Opschoor et al. (2021), among others. We define the log predictive tail score ($LPTS$) measure as follows:

$$LPTS_{q^*} = \sum_{t=T}^{T+K-1} I[u_{t+1} < q] \times \log p(u_{t+1}|z_{1:t}),$$

where q is a $d \times 1$ vector, and $I[u_{t+1} < q] = \prod_{i=1}^{d} I[u_{i,t+1} < q_i]$ with $q_i \in [0,1]$. Here $I[a]$ denotes the indicator function, which equals 1 if condition a is fulfilled and 0 otherwise. We select $q = [q_1, \ldots, q_d]$ such that $K^{-1} \sum_{t=T}^{T+K-1} I[u_{t+1} < q] = q^*$, for $q^* = 0.5, 0.25, 0.10$ (Opschoor et al. 2021). That is, we examine the LPS from Eq.(7) only when the d-variate data are jointly in the lower region $[0, q_1] \times \ldots \times [0, q_d]$.

4 Empirical Application I

4.1 Data description

The daily and intraday equity return data as well as the realized variance and covariance data are from the Multivariate HEAVY article of Noureldin et al. (2012), available at Herber et al. (2009). The data are from 2001/02/01 to 2009/12/31, and contain 2242 data points in total. High frequency returns and the realized covariance measures are extracted by using 5-minute returns with subsampling, as described in Noureldin et al. (2012). The dataset contains some of the most liquid stocks in the Dow Jones Industrial Average (DJIA) index. These are Alcoa (AA), American Express (AXP), Bank of America (BAC), Coca Cola (KO), DuPont (DD), General Electric (GE), International Business Machines (IBM), JP Morgan (JPM), Microsoft (MSFT) and Exxon Mobil (XOM). The online Supplementary Material contains descriptive statistics for all assets. To preserve space, the first set of descriptive plots is for BAC and JPM returns. Figure 3 draws the log returns together with the realized standard deviations, QQ-plots for the standardized returns $(z_{1,t}, z_{2,t})$ against the normal distribution and histograms for the probability integral transforms against the uniform distribution. The corresponding plots for all ten assets can be found in the Online Supplementary Material. As seen from the plots, the time series data includes calm and volatile episodes. The QQ-plot indicates that the data, standardized by the *RV* measure, are approximately normally distributed, as shown by Andersen et al. (2000, 2001). This is also confirmed by the probability integral transforms of the standardized returns $u_{i,t} = \Phi(z_{i,t})$, which are uniformly distributed.

4.2 Prior specification and estimation

The prior hyperparameters for the variance-covariance matrix in the static model are set to $\Omega \sim \mathcal{IW}(I_d, 10)$; for the RMe model, the prior is $\lambda \sim \mathcal{B}(10,3)$ and for the DCC (DCC-HEAVY)-t, the priors are $a \sim \mathcal{B}(3,10)$, $b \sim \mathcal{B}(10,3)$, $\eta \sim \mathcal{E}(0.1)$. The priors for the AIW model are $\nu \sim \mathcal{E}_{\nu>d+1}(0.1)$, $\boldsymbol{b} \sim \mathcal{N}_{2d}(0, 10 \cdot I_{2d})$ and $l_2 \sim \mathcal{U}_Z(2,100)$. In general, all priors are somewhat uninformative but proper. The size of the MCMC chain is $M = 50k$ for all models; the first half is retained as burn-in, and thinning is performed every 25th observation from the second half, thus resulting in posterior samples of 1000 observations. For the RWMH steps, the proposal variances are adjusted such that the

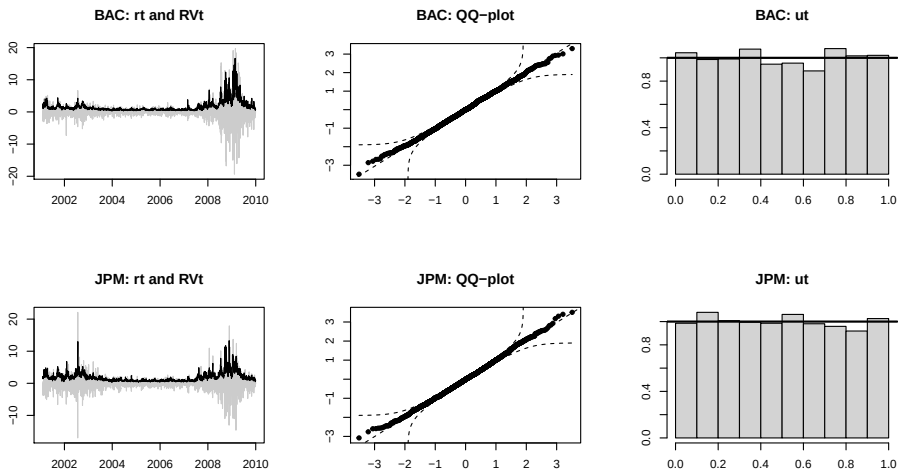

Figure 3: First column: log returns (open-to-close, in gray) and square root of realized volatility (in black). Second column: QQ-plots of the standardized returns against the normal distribution. Third column: histograms of the probability integral transforms against the uniform density for BAC (top row) and JPM (bottom row) assets.

acceptance rate is approximately 0.5 for univariate parameter vectors and 0.10 to 0.30 for multivariate parameter vectors. For sampling l_2 Poisson increments have a rate parameter equal to either 1.5 or 2, depending on the acceptance probability. All MCMC chains have converged after $50k$ iterations. Appendix B in the Online Supplementary Material contains parameter estimation results and trace plots for the parameters for all models. Appendix B also contains the robustness check study, wherein all models are re-estimated by using different hyper-parameter values resulting in more vague priors. The results for all the models using different hyperparameter values remained virtually identical.

4.3 Full model results

For estimation, we used almost all available data, retaining the last year for the out-of-sample performance evaluation. In particular, the data used for estimation are from 2001/02/01 to 2008/12/31 (1990 data points), and the out-of-sample evaluation period is from 2009/01/02 to 2009/12/31 (252 data points). Table 1 presents the average LPS for the $K = 252$ out of sample observations for five low frequency data based models, a high frequency model and a competitor DCC-HEAVY model with t copula (results for the Gaussian copula are not included, because of considerably poorer performance). According to the LPS, the DCC-t model performs best among the low frequency data based models,

and AIW performs best overall. Figure 4 draws expanding-window log predictive BFs for each of the models, wherein the static model is the benchmark. Positive BFs indicate that the model outperforms the static specification. AIW, DCC and DCC-t provide superior out of sample density forecasts, whereas the more restrictive RiskMetrics models are comparable to the static model.

Table 1: 1-step-ahead log predictive scores (LPS) for all individual models: Static, Dynamic conditional correlation with Gaussian and t copulas (DCC and DCC-t), Additive Inverse Wishart (AIW) and DCC-HEAVY model with t copula for 2009/01/02-2009/12/31 out-of-sample period (K = 252 observations).

Static	DCC	DCC-t	AIW	DCC-HEAVY-t
-3134.07	-3125.25	-3119.98	-3111.74	-3127.40

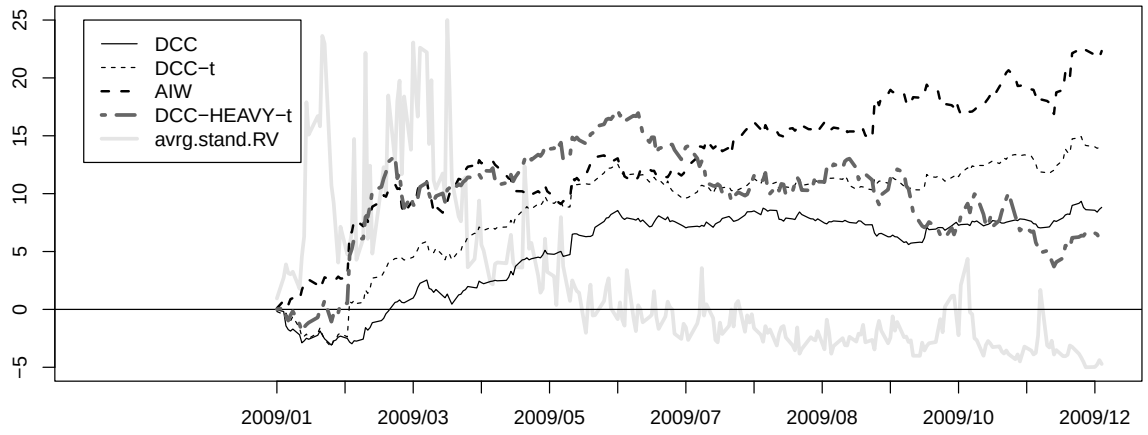

Figure 4: Expanding-window predictive log Bayes factors with a static model as a benchmark for all individual models: RiskMetrics fixed (RMf), RiskMetrics estimated (RMe), Dynamic conditional correlation with Gaussian and t copulas (DCC and DCC-t), Additive inverse Wishart (AIW) and DCC-HEAVY model with t copula for 2009/01/02 to 2009/12/31 out-of-sample period ($K = 252$ observations). Average standardized realized volatility (in gray) in the background.

Next, we perform the predictive density combination exercise, as described in Section 2.3. As seen from Figure 4, model preference is non-constant, and some models that might appear "universally" the best are outperformed by others in certain periods (e.g., AIW vs DCC-t). Therefore, instead of choosing a single model for density prediction, we combine predictive densities by using several alternative weighting schemes. We combine the DCC-t and AIW models, which are the best models in the LF and HF model classes. In fact, the results reported in Appendix D show that combining any other LF model that does not perform as well as DCC-t with the AIW still yields superior predictions to

the best individual model, the AIW. This superiority is not the case for model pools within the same frequency, which perform systematically more poorly than the benchmark. According to those results, we conclude that the predictive gains are primarily from the incorporation of low and high frequency information and not from the model pooling itself.

Figure 5 shows the posterior average of the weights for the HF component (AIW model) for the four weighting schemes: equally weighted, Geweke's as in Eq. (4), Jore's with $\tilde{m} = \{1, 5, 10\}$ as in Eq. (5) and Del Negro (DN) as in Eq. (6). Jore's weights are more volatile because they take into consideration only the last $\tilde{m}$ observations, whereas Geweke's weight takes into consideration the entire out of sample period until the time when the weights are calculated, and reaches a seemingly stable level of approximately 0.6. Drawing the 95% credible intervals around Geweke's weight indicates that the HF component weight is almost always different from 0.5. DN weights are not as volatile as Jore's; however, both follow similar patterns. In particular, because the estimated persistence parameter β in the DN weighting scheme is close to zero, DN and Jore's weights with $\tilde{m} = 1$ are nearly identical, and Jore1 is slightly less smooth. Overall, DN and Jore1 weights fluctuate around 0.5, as expected, given that both DCC-t and AIW perform similarly, particularly during calm periods.

The bottom plot of Figure 5 draws expanding-window predictive log BFs for density combinations and individual models, with the AIW model as the benchmark. All four combination schemes outperform the best individual AIW model. Geweke's and equal weights shows the poorest performance, mainly because the weights are not re-balanced to adjust to a rapidly changing environment. Del Negro's scheme performs better than Geweke's and equal weights, but not as well as Jore1. Even though DN and Jore1 move in very similar patterns, the more extreme movements of Jore1's weights appear to be the source of the superior performance. Overall, all four weighting schemes produce significant improvements in one-step-ahead density prediction over individual LF and HF models. The results also hold for five-step-ahead prediction horizon; Online Supplementary Material Appendix D presents the five-step ahead counterparts of Table 1 and Figures 4 and 5.

Next, Figure 6 draws the posterior densities for the *average per observation* out-of-sample LPS and $LPTS_{q^*}$ for lower 10, 25 and 50% quantiles, only for some individual and some pooled models (for ease of readability of the graph). Of note, 5% or even 1% quantiles would ideally be examined; however, because of the short out-of-sample period ($K = 252$), the resulting sample size would be

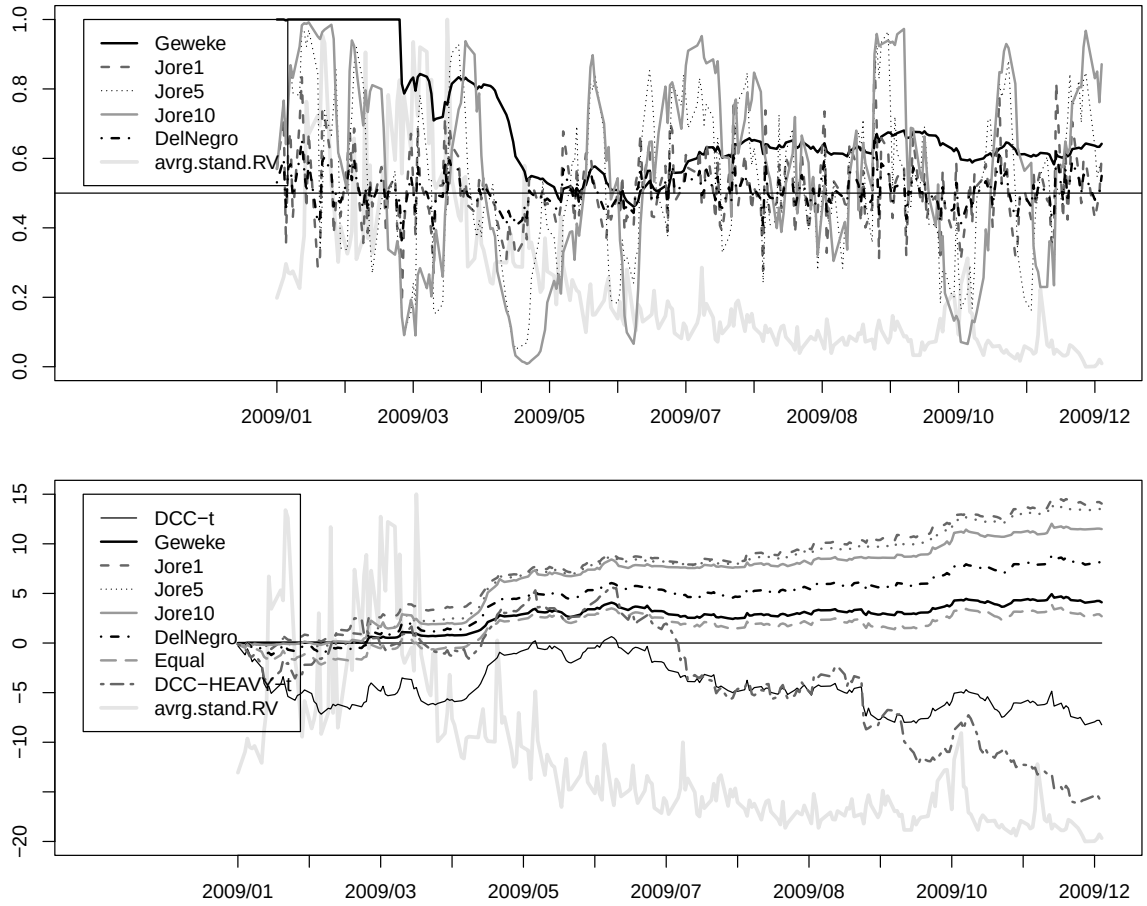

Figure 5: Results for pooling the high-frequency AIW-Gaussian and low-frequency DCC-t copulas. Top plot: posterior mean of the high-frequency component weight for the five different weighting schemes. Bottom plot: expanding-window predictive log Bayes factor for density combinations and individual models, with additive inverse Wishart (AIW) as the benchmark, for 2009/01/02 to 2009/12/31 out-of-sample period ($K = 252$ observations). Average standardized realized volatility (in gray) in the background.

very small.[4] The posterior densities also indicate whether the differences in these average LPS and $LPTS$ are statistically significant. The top left plot indicates that the difference in average overall LPS is statistically significant, pooled models provide the best predictive out of sample performance, and the DCC-t model has the poorest performance. The results change somewhat within the 50% lower quantile (top right plot). Here, the preference for the high-frequency based model is less clear, because the DCC-t and AIW intervals overlap. Pooled models continue to perform best, particularly Jore's 1-period mixing weights. Examination of the first quartile (bottom left) instead of the lower half indicates

[4]There are 12.6 observations in the 5% quantile and only 2.52 in the 1% quantile.

that the model ordering remains the same, but the differences between the models further decrease. Finally, in the 10% lower quantile (bottom right plot), both single component models and two of the pooled models perform virtually identically, and only Jore's pooling scheme performs significantly better than the rest. These results show that different models perform differently depending on the metric used (whole distribution vs the tail of the distribution). Therefore, finding a universally best model is conceptually impossible. In this article, the best model is characterized as that providing the highest log predictive score, because we are interested in the entire predictive distribution of the returns. Nonetheless, if one is interested exclusively in the tails, for example, the log predictive tail score would be a more appropriate metric for calculating the pooling weights. For example, Kapetanios et al. (2015) have proposed to model weights dependent on some variable of interest, which could be some measure related to the lower region of a predictive density.

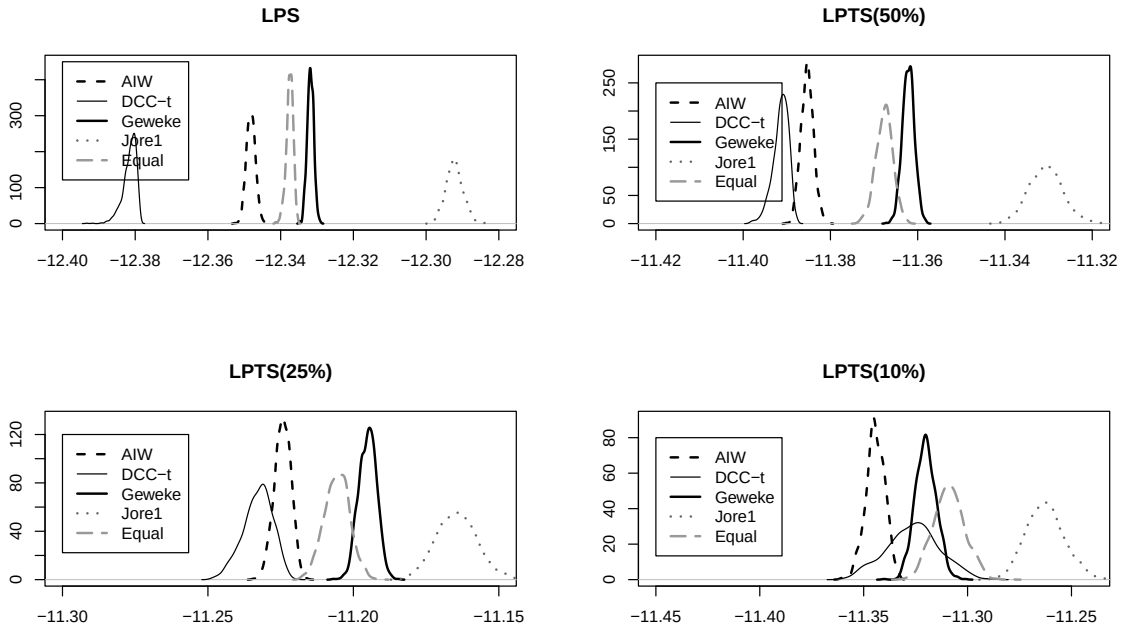

Figure 6: Posterior densities for the average per observation one-step-ahead log predictive score and log predictive tail scores for the lower 50%, 25% and 10% quantiles for 2009/01/02 to 2009/12/31 out-of-sample period ($K = 252$ observations). Additive inverse Wishart (AIW) and dynamic conditional correlation with t copula (DCC-t) are the high and low frequency models, and the pooled models are according to Geweke's, Jore's and equally weighted schemes.

Finally, we sought to determine whether the preference for the high-frequency model might cor-

relate with overall market conditions, proxied by the market volatility. The preference for the high-frequency model is measured as the pooling weight of the high-frequency component in various pooling schemes. We also consider the difference between the predictive log likelihoods between the two best models: AIW and DCC-t. Positive values would indicate that the AIW model is preferred, whereas negative values would indicate that the AIW model is outperformed by the DCC-t. As a proxy for the market volatility, we take the average standardized realized volatility (obtained by using the 5-minute returns with subsampling; see Noureldin et al. 2012) over the ten assets. As alternative proxies, we also consider the equally weighted market portfolio realized volatility, the MCap[5] weighted market portfolio volatility and daily VIX index[6] for the corresponding period. Table 2 reports the posterior medians of the sample correlation coefficients between the preference for the high-frequency weight and the market volatility proxies. Except for Geweke's weights, the rest are negatively correlated, meaning that the preference for the HF model is negatively correlated with the market volatility. The correlations are of relatively small magnitude, but mostly with 95% posterior credible intervals excluding the zero (except for Del Negro's weights). We argue that, as the market volatility increases, the realized variance-covariance measure, used in the AIW model, becomes more contaminated by the market microstructure noise. Hansen & Lunde (2006), for example, have found that the market microstructure noise is negatively correlated with returns. Therefore, the preference for the high-frequency component might decrease because of less reliable estimators during volatile times. Such results have important implications from the investor's perspective: investors anticipating bull (or bear) market conditions, might choose to rely on high-frequency (or low-frequency) models to produce density forecasts.

4.4 Portfolio allocation exercise

Next, we are interested in quantifying how the use of one model versus another translates to a better performing portfolio in terms of economic gains. To this end, we consider the Global Minimum Variance (GMV) portfolio. We note that, even though we have the explicit form of the K one-step-ahead predictive densities for the ten-variate return series at each MCMC iteration, the closed-form expressions for the variance-covariance matrix for the pooled models are analytically unavailable. Therefore,

[5]Market capitalization (MCap) data are from October 31, 2022.

[6]VIX is the Chicago Board Options Exchange's volatility index, based on S&P500 index options.

Table 2: Posterior medians of sample correlations between the preference for high-frequency model and four proxies for the market volatility for 2009/01/02-2009/12/31 out-of-sample period ($K = 252$ observations). The preference for the high-frequency model is measured as a high-frequency component weight in various pooling schemes as well as the difference between the daily log likelihood (diff:logLik) between the AIW and DCC-t models. The proxies for the market volatility are: average standardized realized volatility (avrg RV), equally weighted market portfolio realized volatility (Mkt:eql), MCap weighted market portfolio realized volaltity (Mkt:MCap) and VIX index.

	Geweke	Jore1	Jore5	Jore10	DelNegro	diff:logLik
avrg RV	0.624	-0.028	-0.076	-0.108	-0.060	-0.031
Mkt:eql	0.604	-0.024	-0.069	-0.099	-0.056	-0.027
Mkt:Mcap	0.628	-0.015	-0.067	-0.095	-0.047	-0.019
VIX	0.722	-0.011	-0.026	-0.035	-0.006	-0.015

we use a similar approach to those in Ausín & Lopes (2010) and Opschoor et al. (2021), wherein at each MCMC iteration, and for each out-of-sample point, we draw N replications from the ten-variate predictive distribution, where N is a large number.[7] Given this simulated data, we then can calculate the one-step-ahead variance covariance matrix and perform the GMV portfolio weight calculation. The procedure can be summarized as follows. For each $m = 1, \dots, M$ and for each $k = 1, \dots, K$:

1. Simulate N replications of $u_{T+k}^{(m)}$ from $p(u_{T+k}|z_{1:T+k-1})$ and transform the uniformly distributed data to predictive returns $r_{T+k}^{*(m)}$ via the corresponding quantile function. Because in the copula setting, the modeling of the marginals is performed separately from the dependence structure, all predictive returns have the same marginals (across models, not across assets). To obtain the realized volatility forecasts, we use the log-HAR(1,5,22) model of Corsi (2009); this model produces the best out-of-sample results among some alternative specifications according to six criteria outlined in Hansen & Lunde (2005) (details in Online Supplementary Material Appendix C).

2. Calculate the empirical one-step-ahead variance-covariance matrix $\Sigma_{T+k}^{(m)}$ of the predictive returns and obtain the solution to the quadratic programming problem analytically:

$$w_{T+k}^{(m)} = \min \quad w_{T+k}^{\prime(m)} \Sigma_{T+k}^{(m)} w_{T+k}^{(m)} \quad \text{s.t.} w_{T+k}^{\prime(m)} \iota = 1, \quad \text{with} \quad w_{T+k}^{(m)} = \frac{\Sigma_{T+k}^{-1(m)} \iota}{\iota' \Sigma_{T+k}^{-1(m)} \iota}.$$

[7] $N = 10,000$ in our case.

3. Given the estimated predictive optimal portfolio weights $w_{T+k}^{(m)}$ and the actual *ex post* observed returns r_{T+k}, we can calculate various *ex post* portfolio metrics of interest. Because we have M of such weight vectors for each time period, we can also have the entire posterior distributions of these quantities.

In particular, we calculate the $K = 252$ sequence of the realized portfolio returns. Given these realizations, we can obtain the overall portfolio variance, Sharpe ratio (the ratio between the expected return and the standard deviation), the empirical 5% and 10% quantiles or the expected value in these quantiles, which is the value-at-risk (VaR) and the expected shortfall (ES). Next, as in Opschoor et al. (2021), we also calculate portfolio turnover (TO), concentration (CO) and short position (SP):

$$\begin{aligned}
TO_{T+k}^{(m)} &= \sum_{i=1}^{d} \left| w_{i,T+k+1}^{(m)} - w_{i,T+k}^{(m)} \frac{1 + r_{i,T+k}}{1 + w_{T+k}^{'(m)} r_{T+k}} \right|, \\
CO_{T+k}^{(m)} &= \sum_{i=1}^{d} \left(w_{i,T+k}^{(m)2} \right)^{1/2}, \\
SP_{T+k}^{(m)} &= \sum_{i=1}^{d} w_{i,T+k}^{(m)} \times I[w_{i,T+k}^{(m)} < 0].
\end{aligned}$$

Here, $w_{i,T+k}^{(m)}$ is the ith element of the GMV portfolio weight vector at iteration $m = 1, \ldots, M$ for out-of-sample period $k = 1, \ldots, K$. The portfolio turnover measures the value of the portfolio that is bought/sold from time $T + k$ to $T + k + 1$. An investor would prefer smaller values of TO_t, which imply lower transaction costs. We also calculate the turnover-adjusted realized portfolio returns and Sharpe ratios for a fairer economic comparison (using 1% transaction costs). Portfolio concentration and portfolio short position measure how extreme the portfolio weights are. An investor will prefer a model that provides the smallest concentration and the largest short position measures. Finally, we calculate the gain/loss (G/L) ratio, as in Conrad & Stürmer (2017). The G/L ratio shows the percentage of the expected return gain if we use some other model over the benchmark model to form our portfolio:

$$G/L^{(m)} = 100\%(\sigma_{AIW}^{(m)} - \sigma_P^{(m)})/\sigma_P^{(m)},$$

where $\sigma_{AIW}^{2(m)}$ and $\sigma_P^{2(m)}$ are the realized portfolio variances at iteration m, based on the predicted variance-covariance matrices of the best performing individual model, which is AIW, and some other

competing model, respectively. The preferred portfolio will have a positive and higher G/L ratio.

Table 3 reports the posterior medians of various GMV portfolio metrics based on the one-step-ahead predictions for the competing models. The realized portfolio standard deviations, Sharpe ratios and G/L criteria are annualized. In the upper panel of the table, we focus on the most relevant measures for an investor minimizing the global variance. Two of the pooled models, Geweke's and Jore's, reduce the portfolio variance in comparison to the AIW model, which provides the smallest variance among the individual models. The adjusted Sharpe ratio, calculated by using the turnover-adjusted returns, is the highest for Jore's pooled model, and the DCC-HEAVY-t model is second best. The G/L criteria favor Jore's and Geweke's pooled models, with a 3.5% and 2.6% relative increase in the required access return, respectively, when moving away from the benchmark AIW. In the lower panel of Table 3, we provide the results for additional risk measures, which do not necessarily reflect the investor's objective of minimal portfolio variance. In most cases, Jore's pooled model provides the best results in terms of realized 5% and 10% VaR and ES of the one-step-ahead portfolio returns. In contrast, the SP, CO and TO metrics favor the AIW model. In conclusion, the portfolio allocation results show favorable economic outcomes for density pooling, as compared with the individual models, when the variance-targeting measures are considered.

Table 3: GMV portfolio results based on 1-step-ahead forecasts for 2009/01/02 to 2009/12/31 out-of-sample period ($K = 252$ observations) for 10-variate dataset. The table reports the posterior 5, 50 and 95 percentiles of G/L criteria as well as portfolio standard deviation (in %) for the pooled models (Geweke's, Jore's and equally weighted), two best individual models (Additive Inverse Wishart and Dynamic Conditional Correlation with t copula) and a competitor model (DCC-HEAVY-t).

	G/L			Portfolio stdev.		
	P05	Median	P95	P05	Median	P95
Geweke	-12.606	-9.521	-5.483	91.182	91.611	92.009
Jore1	**0.288**	**12.388**	**19.471**	**89.976**	**90.362**	**91.051**
Equal	-17.907	-11.755	-5.542	91.400	91.747	92.094
AIW	0.000	0.000	0.000	90.692	91.067	91.448
DCC-t	-21.173	-12.276	-3.331	91.410	91.777	92.147
DCC-HEAVY-t	-48.392	-37.908	-28.554	92.846	93.306	93.786

5 Empirical Application II

To verify whether the previous results hold for different samples/periods, we perform a second empirical exercise, which is briefly summarized here. We consider a five-variate dataset of three exchange rates (EUR/USD, USD/JPY and EUR/GBP), the SP500 market index and WTI oil spot prices for the period 2012/01/04 to 2021/12/31 (10 years of data, with 2508 daily observations in total). We use the first 9 years for estimation (2256 observations) and the last 1 year (252 observations) for one-step-ahead forecast evaluation. Table 4 presents the differences between the one-step-ahead LPS between the AIW model and the rest of the models. As in the previous empirical application, among the single frequency models the high-frequency AIW model performs best overall, and the low-frequency DCC model with t copula is the second best. Interestingly, the competitor DCC-HEAVY-t outperforms all single frequency models but not the pools: almost all HF-LF pooling schemes outperform all other models (except for the equally weighted pool), and the Jore1 model performs best. This is the same conclusion that we made according to the first data set considered in Section 4. Finally, Table 5 presents the GMV portfolio allocation results. Jore's pool produces the highest G/L ratio, as compared with the individual AIW model, with a remarkable 17.6% gain. The results in the upper panel of the table are close to those seen in the Empirical Application I. However, for the rest of the measures, we a find less clear preference for one model over another.

Table 4: Differences in the 1-step-ahead log predictive scores (LPS) between the best-fitting individual model (AIW) and the rest of the models: Static, Dynamic conditional correlation with Gaussian and t copulas (DCC and DCC-t), DCC-HEAVY with t copula and various AIW-DCC-t pools for 2021/01/04 - 2021/12/31 out-of-sample period (K = 252 observations).

Static	DCC	DCC-t	AIW	DCC-HEAVY-t	Gew	Jore1	Jore5	Jore10	DN	Equal
-18.12	-3.62	-2.55	0.00	4.84	5.09	10.32	8.79	6.81	6.41	3.06

All the details containing data description, estimation and prediction results, with corresponding tables and figures, can be found in the Appendix E in the Online Supplementary Material.

6 Discussion and Conclusion

In this paper, we propose a mixed frequency copula-based approach that enables to model the dependence between financial returns by using information arising from data sampled at different frequen-

Table 5: GMV portfolio results based on 1-step-ahead predictions for 2021/01/04-2021/12/31 out-of-sample period ($K = 252$ observations) for 5-variate dataset. The table reports the posterior 5, 50 and 95 percentiles of G/L criteria as well as portfolio standard deviation (in %) for the pooled models (Geweke's, Jore's and equally weighted), two best individual models (Additive Inverse Wishart and Dynamic Conditional Correlation with t copula) and a competitor model (DCC-HEAVY-t).

	G/L			Portfolio stdev.		
	P05	Median	P95	P05	Median	P95
Geweke	7.412	13.036	19.115	14.086	14.141	14.193
Jore1	**24.647**	**30.459**	**35.342**	**13.950**	**13.990**	**14.035**
Equal	-1.321	3.772	8.738	14.183	14.224	14.269
AIW	0.000	0.000	0.000	14.221	14.258	14.299
DCC-t	2.027	8.521	14.480	14.143	14.182	14.225
DCC-HEAVY-t	0.068	6.925	13.198	14.148	14.198	14.242

cies. We rely on a density pooling approach to combine alternative copula models to describe the daily dependence structure.

In particular, we pool two copula densities, wherein the parameters are obtained from low and high frequency data. For the high frequency copula parameter, we use a realized correlation measure. We model the dynamics of the realized variance-covariance matrices via an additive inverse Wishart model with a Gaussian copula; meanwhile, for the low-frequency dependence structure, we consider five standard models: static, RMf, RMe and DCC, all with Gaussian copula, and DCC with t copula. The DCC-t model always performs best among the low frequency data based models. In both empirical applications, even though the overall log predictive scores favor the AIW model, incorporating information arising from the low frequency data improves the predictive model's performance. In addition, the density pool shows an improvement in predictive performance over those of other mixed frequency models, such as the natural competitor the DCC-HEAVY model. Finally, we show that the gains arise not from density pooling itself, but from pooling different frequencies, and that the results also hold for longer prediction horizons.

For future research, an infinite component mixture could be considered for high-frequency data based models (Jin & Maheu 2016). In addition, a more flexible pooling scheme, such as Bayesian predictive synthesis, would result in overall better models. Finally, the use of more flexible copulas, such as inversion copulas, should also be considered.

Acknowledgments

This work was partially supported by grant PID2020-113192GB-I00 (Mathematical Visualization: Foundations, Algorithms and Applications) from the Spanish MICINN.

References

Aastveit, K. A., Gerdrup, K. R. & Jore, A. S. (2011), 'Short-term forecasting of gdp and inflation in real-time: Norges bank's system for averaging models', *Norges Bank, Staff Memo* **9**, 1–56.

Aastveit, K. A., Gerdrup, K. R., Jore, A. S. & Thorsrud, L. A. (2014), 'Nowcasting GDP in real time: A density combination approach', *Journal of Business and Economic Statistics* **32**(1), 48–68.

Aastveit, K. A., Mitchell, J., Ravazzolo, F. & Dijk, H. K. V. (2018), 'The evolution of forecast density combinations in economics', *Tinbergen Institute Discussion Paper* (TI 2018-069/III), 1–48.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Ebens, H. (2001), 'The distribution of realized stock return volatility', *Journal of Financial Economics* **61**(1), 43–67.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2000), 'Exchange rate returns standardized by Realized Volatility are (Nearly) Gaussian', *NBER Working Paper Series* **7488**, 1–21.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **71**(2), 579–625.

Andrieu, C., Doucet, A. & Holenstein, R. (2010), 'Particle Markov chain Monte Carlo methods', *Journal of the Royal Statistical Society. Series B: Statistical Methodology* **72**(3), 269–342.

Asai, M., McAleer, M. & Yu, J. (2006), 'Multivariate Stochastic Volatility: A Review', *Econometric Reviews* **25**(2-3), 145–175.

Ausín, M. C. & Lopes, H. F. (2010), 'Time-Varying Joint Distribution Through Copulas', *Computational Statistics and Data Analysis* **54**(11), 2383–2399.

Barndorff-Nielsen, O. E. & Shephard, N. (2002), 'Econometric analysis of realized volatility and its use in estimating stochastic volatility models', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66**(2), 253–280.

Barndorff-Nielsen, O. E. & Shephard, N. (2004), 'Econometric analysis of realized covariation: high frequency covariance, regression and correlation in financial economics', *Econometrica* **72**(3), 885–925.

Bassetti, F., Casarin, R. & Ravazzolo, F. (2018), 'Bayesian nonparametric calibration and combination of predictive distributions', *Journal of the American Statistical Association* **113**(522), 675–685.

Bauer, G. H. & Vorkink, K. (2011), 'Forecasting multivariate realized stock market volatility', *Journal of Econometrics* **160**(1), 93–101.

Bauwens, L., Laurent, S. & Rombouts, J. V. K. (2006), 'Multivariate GARCH Models: a Survey', *Journal of Applied Econometrics* **21**(1), 79–109.

Bauwens, L. & Xu, Y. (2022), 'Dcc- and deco-heavy: Multivariate garch models based on realized variances and correlations', *International Journal of Forecasting* **In Press**.

Billio, M., Casarin, R., Ravazzolo, F. & Van Dijk, H. K. (2013), 'Time-varying combinations of predictive densities using nonlinear filtering', *Journal of Econometrics* **177**(2), 213–232.

Bjørnland, H. C., Jore, A. S., Smith, C. & Thorsrud, L. A. (2008), 'Improving and evaluating short term forecasts at the norges bank', *Staff Memo, Norges Bank* **4**, 1–60.

Brechmann, E. C. & Czado, C. (2015), 'COPAR - Multivariate time series modeling using the copula autoregressive model', *Applied Stochastic Models in Business and Industry* **31**(4), 495–514.

Casarin, R., Grassi, S., Ravazzolo, F. & van Dijk, H. K. (2023), 'A flexible predictive density combination for large financial data sets in regular and crisis periods', *Journal of Econometrics* **In Press**, 1–12.

Chiriac, R. & Voev, V. (2011), 'Modelling and Forecasting Multivariate Realized Volatility', *Journal of Applied Econometrics* **26**(6), 922–947.

Clemen, R. T. & Winkler, R. L. (2007), Aggregating probability distributions, *in* E. Ward, R. Miles & D. von Winterfeldt, eds, 'Advances in Decision Analysis: From Foundations to Applications', Cambridge University Press, pp. 154–176.

Conrad, C. & Stürmer, K. (2017), 'On the economic determinants of optimal stock-bond portfolios: international evidence', *Updated version of: University of Heidelberg, Department of Economics, Discussion Paper Series* **636**.

Corsi, F. (2009), 'A Simple Approximate Long-Memory Model of Realized Volatility', *Journal of Financial Econometrics* **7**(2), 1–23.

Del Negro, M., Hasegawa, R. B. & Schorfheide, F. (2016), 'Dynamic prediction pools: An investigation of financial frictions and forecasting performance', *Journal of Econometrics* **192**(2), 391–405.

Delatola, E.-I. & Griffin, J. E. (2011), 'Bayesian Nonparametric Modelling of the Return Distribution with Stochastic Volatility', *Bayesian Analysis* **6**(4), 901–926.

Demarta, S. & McNeil, A. J. (2005), 'The t copula and related copulas', *International Statistical Review* **73**(1), 111–129.

Dias, A. & Embrechts, P. (2004), 'Dynamic copula models for multivariate high-frequency data in finance', *Manuscript, ETH Zurich* **81**(November), 1–42.

Ding, Z. & Engle, R. F. (2001), 'Large Scale Conditional Covariance Matrix Modeling , Estimation and Testing', *NYU Working Paper No. FIN-01-029* .

Engle, R. F. (2002), 'New Frontiers for ARCH Models', *Journal of Applied Econometrics* **17**(5), 425–446.

Fengler, M. R. & Okhrin, O. (2016), 'Managing risk with a realized copula parameter', *Computational Statistics and Data Analysis* **100**, 131–152.

Geweke, J. & Amisano, G. (2011), 'Optimal prediction pools', *Journal of Econometrics* **164**(1), 130–141.

Ghysels, E., Santa-Clara, P. & Valkanov, R. (2004), 'The MIDAS Touch: Mixed Data Sampling Regression Models', *UCLA: Finance* pp. 1–33.

Ghysels, E., Santa-Clara, P. & Valkanov, R. (2005), 'There is a risk-return trade-off after all', *Journal of Financial Economics* **76**(3), 509–548.

Gordon, N., Salmond, D. & Smith, A. (1993), 'Novel approach to nonlinear/non-Gaussian Bayesian state estimation', *IEE Proceedings F (Radar and Signal Processing)* **140**(2), 107–113.

Gourieroux, C., Jasiak, J. & Sufana, R. (2009), 'The Wishart Autoregressive process of multivariate stochastic volatility', *Journal of Econometrics* **150**(2), 167–181.

Hall, S. G. & Mitchell, J. (2007), 'Combining density forecasts', *International Journal of Forecasting* **23**(1), 1–13.

Hansen, P. R., Huang, Z. & Shek, H. H. (2012), 'Realized GARCH: a Joint Model for Returns and Realized Measures of Volatility', *Journal of Applied Econometrics* **27**(6), 877–906.

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

Hansen, P. R. & Lunde, A. (2006), 'Realized variance and market microstructure noise', *Journal of Business and Economic Statistics* **24**(2), 127–161.

Hansen, P. R., Lunde, A. & Voev, V. (2014), 'Realized beta GARCH: a Multivariate GARCH Model with Realized Measures of Volatility', *Journal of Applied Econometrics* **29**(5), 774–799.

Herber, G., Lunde, A., Shephard, N. & Sheppard, K. (2009), 'Oxford-man institute of quantitative finance realized library', *Oxford-Man Institute, University of Oxford* .

Horpestad, J. B., Štefan Lyócsa, Molnár, P. & Olsen, T. B. (2019), 'Asymmetric volatility in equity markets around the world', *North American Journal of Economics and Finance* **48**(4), 540–554.

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

Jin, X. & Maheu, J. M. (2016), 'Bayesian semiparametric modeling of realized covariance matrices', *Journal of Econometrics* **192**(1), 19–39.

Joe, H. (2015), *Dependence Modeling with Copulas*, CRC Press, Boca Raton, FL.

Jore, A. S., Mitchell, J. & Vahey, S. P. (2010), 'Combining forecast densities from VARs with uncertain instabilities', *Journal of Applied Econometrics* **25**(4), 621–634.

Kapetanios, G., Mitchell, J., Price, S. & Fawcett, N. (2015), 'Generalised density forecast combinations', *Journal of Econometrics* **188**(1), 150–165.

Kascha, C. & Ravazzolo, F. (2010), 'Combining Inflation Density Forecasts', *Journal of Forecasting* **29**(1-2), 231–250.

Kass, R. E. & Raftery, A. E. (1995), 'Bayes Factors', *Journal of the American Statistical Association* **90**(430), 773–795.

Koopman, S. J., Jungbacker, B. & Hol, E. (2005), 'Forecasting daily variability of the s&p 100 stock index using historical, realised and implied volatility measurements', *Journal of Empirical Finance* **12**(3), 445–475.

Koopman, S. J., Lit, R., Lucas, A. & Opschoor, A. (2018), 'Dynamic discrete copula models for high-frequency stock price changes', *Journal of Applied Econometrics* **33**(7), 966–985.

Loaiza-Maya, R. & Smith, M. S. (2018), 'Time series copulas for heteroskedastic data', *Journal of Applied Econometrics* **33**(3), 332–354.

Loaiza-Maya, R. & Smith, M. S. (2020), 'Real-Time Macroeconomic Forecasting With a Heteroscedastic Inversion Copula', *Journal of Business and Economic Statistics* **28**(2), 470–486.

Lyócsa, Š., Molnár, P. & Výrost, T. (2021), 'Stock market volatility forecasting: Do we need high-frequency data?', *International Journal of Forecasting* **37**(3), 1092–1110.

McAleer, M. & Medeiros, M. C. (2008), 'Realized volatility: A review', *Econometric Reviews* **27**(1-3), 10–45.

McAlinn, K. (2021), 'Mixed-frequency bayesian predictive synthesis for economic nowcasting', *Journal of the Royal Statistical Society, Series C* **70**(5), 1143–1163.

McAlinn, K., Aastveit, K. A., Nakajima, J. & West, M. (2020), 'Multivariate Bayesian Predictive Synthesis in Macroeconomic Forecasting', *Journal of the American Statistical Association* **115**(531), 1092–1110.

McAlinn, K. & West, M. (2019), 'Dynamic Bayesian predictive synthesis in time series forecasting', *Journal of Econometrics* **210**(1), 155–169.

McNeil, A. J., Frey, R. & Embrechts, P. (2005), *Quantitative risk management: Concepts, techniques, and tools*, Princeton University Press.

Nelsen, R. B. (2006), *An Introduction to Copulas, 2nd Edition*, Springer.

Nguyen, H. & Javed, F. (2021), 'Dynamic relationship between stock market and bond market: A gas midas copula approach', *Working Paper, Örebro University School of Business, Örebro* **15**, 1–45.

Noureldin, D., Shephard, N. & Sheppard, K. (2012), 'Multivariate High-Frequency-Based Volatility (HEAVY) Models', *Journal of Applied Econometrics* **27**(6), 907–933.

Okhrin, O. & Tetereva, A. (2017), 'The Realized Hierarchical Archimedean Copula in Risk Modelling', *Econometrics* **5**(2), 26.

Opschoor, A., Lucas, A., Barra, I. & van Dijk, D. (2021), 'Closed-Form Multi-Factor Copula Models With Observation-Driven Dynamic Factor Loadings', *Journal of Business and Economic Statistics* **39**(4), 1066–1079.

Patton, A. J. (2006*a*), 'Estimation of multivariate models for time series of possibly different lengths', *Journal of Applied Econometrics* **21**(2), 147–173.

Patton, A. J. (2006*b*), 'Modelling Asymmetric Exchange Rate Dependence', *International Economic Review* **47**(2), 527–556.

Patton, A. J. (2012), 'A review of copula models for economic time series', *Journal of Multivariate Analysis* **110**, 4–18.

Rufo, M. J., Martín, J. & Pérez, C. J. (2012), 'Log-linear pool to combine prior distributions: A suggestion for a calibration-based approach', *Bayesian Analysis* **7**(2), 411–438.

Salvatierra, I. D. L. & Patton, A. J. (2015), 'Dynamic copula models and high frequency data', *Journal of Empirical Finance* **30**, 120–135.

Shephard, N. & Sheppard, K. (2010), 'Realising the Future: Forecasting with High-Frequency-Based Voality (HEAVY) Models', *Journal of Applied Econometrics* **25**(2), 197–231.

Silvennoinen, A. & Teräsvirta, T. (2009), Multivariate GARCH models, *in* T. G. Andersen, R. A. Davis, J. Kreiss & T. Mikosch, eds, 'Handbook of Financial Time Series', Springer-Verlag, pp. 201–226.

Sklar, A. (1959), 'Fonctions de répartition à n dimensions et leurs marges', *Publications Inst. Statis., Univ. Paris-VIII* pp. 229–231.

Smith, M. S., Gan, Q. & Kohn, R. J. (2012), 'Modelling dependence using skew t copulas: Bayesian inference and applications', *Journal of Applied Econometrics* **27**(3), 500–522.

Stone, M. (1961), 'The opion pool.', *Annals of Mathematical Statistics* **32**, 1339–1342.

Takanashi, K. & McAlinn, K. (2021), 'Predictions with dynamic Bayesian predictive synthesis are exact minimax', *arXiv preprint arXiv:1911.08662* pp. 1–27.

Timmermann, A. (2018), 'Forecasting methods in finance', *Annual Review of Financial Economics* **10**(11), 449–479.

Tsay, R. S. (2014), *Multivariate Time Series Analysis*, John Wiley & Sons.

Tse, Y. K. & Tsui, A. K. C. (2002), 'A Multivariate Generalized Autoregressive Conditional Heteroscedasticity Model With Time-Varying Correlations', *Journal of Business and Economic Statistics Economic Statistics* **20**(3), 351–362.

West, M. (1986), 'Bayesian Model Monitoring', *Journal of the Royal Statistical Society. Series B (Methodological)* **48**(1), 70–78.