Bayesian Multivariate GARCH Models for Within-Person Variance Forecasting

Philippe Rast, Stephen R. Martin, Josue E. Rodriguez

University of California, Davis

Corresponding author:

Philippe Rast

Department of Psychology

University of California Davis

prast@ucdavis.edu

Author Note

Research reported in this publication was supported by the National Institute On Aging of the National Institutes of Health (NIH) under Award Number R01AG050720 to PR, as well as NIH awards P30 AG024978 to Jeffrey Kaye at ORCATECH and R01 AG024059, P30 AG008017 and Intel Corporation to JK at ORCATECH Living Lab. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

Abstract

We present a multivariate GARCH (MGARCH) model to forecast model averaged (co-)variances, correlations, and means in behavioral data for single individuals. We consider four MGARCH model parameterizations; three classic models (CCC, DCC, and BEKK) as well as a recently introduced model (pdBEKK) optimized for behavioral research. To obtain the averaged forecasts across the four models, we will first need to compute the model weights, obtained from a model stacking method. To do so, we need to approximate the expected log predictive density via a fully Bayesian leave-future-out cross validation technique. This approach has not been described so far in the literature. We provide an illustrative implementation using real data on two individuals from the longitudinal Intelligent Systems for Assessing Aging Change (ISAAC) study covering up to 4 years of daily measurements on individual computer use and walking speed. The individual participants show distinct patterns in the model weights, suggesting that individuals differ in the parameterizations that best capture their behavior. We generate weighted forecasts for up to 5 consecutive two-week periods for both individuals. These forecast are compared to forecasts from the single best model. The resulting predictions are shown to be superior or at least equivalent to the forecasts from the single best model. We close with a discussion on limitations and an outlook. We provide an R package **bmgarch** that implements all the discussed elements and weighted forecasts.

Keywords: Behavioral forecasting, leave-future-out cross validation, Bayesian MGARCH, within-person variability

Bayesian Multivariate GARCH Models for Within-Person Variance Forecasting

Within-person variability has been the focus of psychological research for many years (e.g. [Fiske & Rice, 1955](#); [Woodrow, 1932](#)) but the recent decade has seen a sharp increase in statistical methods geared toward the modeling of such variance ([Hedeker & Mermelstein, 2007](#); [Hedeker, Mermelstein, & Demirtas, 2008](#); [Leckie, French, Charlton, & Browne, 2014](#); [Ram & Gerstorf, 2009](#); [Rast & Ferrer, 2018](#); [Rast, Hofer, & Sparks, 2012](#)). With the rise of ecological momentary assessments, daily diary studies and intensive measurement designs, more and more data are available at the individual level. To date, most models focus on extracting meaningful parameters to explain observed behavior. But with increasing amounts of data, available continuously, one might be able to forecast individual behavior. This is potentially of interest for fields like precision medicine or clinical interventions where one wants to forecast behavior for a particular individual.

These changes in the availability of data have also resulted in a re-evaluation of psychological models in terms of generalizability of findings from current statistical population based models to the individual ([Nesselroade & Molenaar, 2010](#)). In a seminal piece, [Molenaar, Huizenga, and Nesselroade \(2003\)](#) have shown that generalizations from population based models to individuals only hold under very strict ergodicity assumptions that do not typically exist in psychological research ([Molenaar, Sinclair, Rovine, Ram, & Corneal, 2009](#)). As a consequence of both, the availability of person level data and the insight that predictions at the individual level necessitate individual level models, research methods for estimating individual behavior have received a new impetus in recent years. Especially models for the investigation of patterns of variability within individuals, and more generally models and methods on single-case designs, are a growing and active field of research in psychology ([Castro-Schilo & Ferrer, 2013](#); [Ferrer, Gonzales, & Steele, 2013](#); [Natesan & Hedges, 2017](#); [Rindskopf, 2014](#); [Shadish, 2014](#); [Sullivan, Shadish, & Steiner, 2015](#)).

At the same time, the research focus has widened from focusing on the mean

structure over time to also including within-person variation over the same time span. It is not unusual to observe no average changes over time in a person or a group, but the actual behavior in question of an individual might be fluctuating substantially around that mean. While in many applications these fluctuations are seen as residual error and are not of further interest, others seek to identify meaningful patterns in the nature of that variation itself (Almeida, Piazza, & Stawski, 2009; Cleveland, Denby, & Liu, 2002; Hamaker, Asparouhov, Brose, Schmiedek, & Muthén, 2018; Hedeker et al., 2008; Ram & Gerstorf, 2009; Rast, Hofer, & Sparks, 2012; Rush, Rast, Almeida, & Hofer, 2019; Williams, Mulder, Rouder, & Rast, 2020). While these works all include fixed and random effects for the within-person or residual variance and allow that variance to vary across time, given individual specific observations, they are all population models. As such they don't allow to make statements about individuals. This present work combines single-case designs with the focus on forecasting dynamic changes of within-person variability in single individuals

Psychology is not the only field concerned with identifying meaningful patterns of variability in individual time series. In fact, other disciplines have been exploring such models for decades. For example, the field of econometrics has a long tradition of modeling time series data in order to estimate fluctuations in market indices and individual stocks, mostly with the goal of forecasting the volatility for the next day or the near future. One of the standard tools since the 1980s is the autoregressive conditional heteroskedasticity (ARCH) model and its generalized extension (GARCH; Bollerslev, 1986; Engle, 1982; Tsay, 2013).

GARCH models can be readily extended to multivariate models (MGARCH; Engle & Kroner, 1995) and since their introduction a number of different parameterizations have been presented (for an overview see eg. Bauwens, Laurent, & Rombouts, 2006; de Almeida, Hotta, & Ruiz, 2018).

Here, we investigate the usefulness of MGARCH models to forecast within-person behavioral variability for one or few time periods ahead in a limited number of variables

per individual. While the focus in this current paper is on the within-person variance, MGARCH models can be combined with other time-series models for the means, such as vector autoregressive (VAR), moving average models or any other model that accommodates the mean response. In practice, however, MGARCH models tend to be fit on the residuals of detrended data, although that approach is only recommended with gaussian innovations (Carnero & Eratalay, 2014).

The remainder of this paper is structured as follows. First we introduce the MGARCH process along with four parameterizations that will be used later in the illustrative example. Next, we will describe a fully Bayesian approximation, recently introduced by Bürkner, Gabry, and Vehtari (2019), to the log expected predictive density (ELPD) based on leave-future-out cross validation (LFO-CV). On the basis of the $\text{ELPD}_{\text{LFO-CV}}$ we show how composite forecasts for the conditional first and second moments can be obtained by averaging over all candidate models using stacking weights. This approach is then illustrated in data from two individuals from a longitudinal study with daily measurements on walking speed and computer use over up to four years. The forecasts obtained from the averaged model and the single best model will be compared on the basis of Kullback-Leibler divergence (D_{KL}) for the variances and covariances as well as the root mean squared forecasting error (RMSE) for the expected means. Finally, we discuss advantages and short-comings of the presented approach.

Volatility Models

For ease of presentation, we can define a time series with zero mean ($\mu = \mathbf{0}$) as

$$\mathbf{y}_t \sim N(\mu, \mathbf{H}_t), \quad (1)$$

where $\mathbf{H}_t$ is a conditional covariance matrix and $\mathbf{y}_t$ is a $T \times d$ matrix containing the observed d time series of length T . That is, $\mathbf{H}_t$ is time-dependent and its values will be a result of changes that happened in the past. $\mathbf{H}_t$ is the main focus of MGARCH models. While conceptually simple, the crux is to define a model that maintains stationarity of the

variances and covariances among the time-series for all time points as well as positive definiteness for all $\mathbf{H}_t$.

The initial idea behind modeling heteroskedasticity in time-series was derived from the autoregressive and moving average (ARMA) model for the means structure. These mechanisms were developed to partition the residual variance components of the models. First, [Engle \(1982\)](#) introduced autoregressive conditional heteroskedasticity (ARCH) models that were later generalized (GARCH; [Bollerslev, 1986](#)). Essentially, GARCH models follow the logic of ARMA models but the processes of interest are the error variances and not the means. While typical time-series models assume constant error variance (σ^2), ARCH models allow that variance to be time varying and resulting from the sum of a constant and conditional components. To account for the possibility of a serial correlation in the error variance, the ARCH model considers the conditional variance as time dependent, $\mathcal{V}(\varepsilon_t|\varepsilon_{t-1}) = h_t$ with

$$h_t = c + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2, \quad (2)$$

where α governs the effect of previous shocks operationalized as squared residuals (similar to the MA component). In that simple case, the time-series for y can be given by

$$y_t = h_t^{1/2} \eta_t, \quad (3)$$

where $\mathcal{V}(\eta_t) = 1$. Adding the assumption of normality we can express $y_t|\mathcal{F}_{t-1} \sim N(0, h_t)$, conditional on the information set $\mathcal{F}_{t-1}$, that is, all available information through time t .

Model (2) can be extended to a generalized ARCH model (GARCH; [Bollerslev, 1986](#)) of order p, q that includes an AR process on the past variance

$$h_t = c + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^p \beta_i h_{t-i}. \quad (4)$$

Overall, the conditional parameters α and β of the GARCH(p, q) models are scale invariant while c is not.

MGARCH Model Parameterizations

The definition of the GARCH model can be readily extended to multivariate cases with $d = 1, \dots, D$ simultaneous time-series (Bollerslev, Engle, & Wooldridge, 1988). One of the main challenges, however, is the estimation of such multivariate GARCH (MGARCH) models. For one, the number of parameters to be estimated can rise very quickly while the covariance matrix $\mathbf{H}_t$ ($\epsilon_t | \mathcal{F}_{t-1} \sim N(\mathbf{0}, \mathbf{H}_t)$) needs to stay positive-definite across all t time points.

The following sections describe four specifications of MGARCH models: BEKK, pdBEKK, CCC and DCC. BEKK and pdBEKK are attractive due to their possible interpretability in behavioral contexts while CCC and DCC are known to perform well in forecasting tasks and are, thus, popular choices in econometric applications (de Almeida et al., 2018)

Generally speaking, there are two approaches to estimating $\mathbf{H}_t$. One approach is to model $\mathbf{H}_t$ directly while the other approach is to partition $\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$ and estimate the conditional standard deviations $\mathbf{D}_t$ and conditional correlations $\mathbf{R}_t$ separately. We will discuss the parameterizations by first introducing two models that operate on $\mathbf{H}_t$ directly, and two models that operate on $\mathbf{D}_t$ and $\mathbf{R}$.

BEKK. Engle and Kroner (1995) introduced the BEKK specification to address a number of issues arising from an earlier MGARCH model (VECH Bollerslev et al., 1988).

The BEKK(P,Q) can be defined as

$$\mathbf{H}_t = \mathbf{C} + \sum_{q=1}^Q \mathbf{A}'_q (\mathbf{y}_{t-q} \mathbf{y}'_{t-q}) \mathbf{A}_q + \sum_{p=1}^P \mathbf{B}'_p \mathbf{H}_{t-p} \mathbf{B}_p, \quad (5)$$

where $\mathbf{C}$ is a $d \times d$ positive definite symmetric covariance matrix that captures the baseline (co-)variance, while $\mathbf{A}$ and $\mathbf{B}$ are $d \times d$ square parameter matrices that govern the influence of the past residuals, and the past conditional variance, respectively. Q and P govern the lag of the ARCH and GARCH process. For the remainder of this and other model parameterizations we will keep $Q = P = 1$. Engle and Kroner (1995) showed that the

BEKK specification is stationary if and only if all the absolute eigenvalues for the Kronecker product $\mathbf{A} \otimes \mathbf{A} + \mathbf{B} \otimes \mathbf{B}$ are less than one. Note that the cross lagged effects in the off-diagonal elements in both parameter matrices allow the variance terms to spill over from one time-series into the other – and vice versa. Given that these matrices are square, the spill over does not need to be symmetrical.

Conceptually, the GARCH process (captured by $\mathbf{B}$) carries-over past conditional variance. In absence of an ARCH process, the GARCH process alone would recycle a given portion of the conditional variance, resulting in a process with stable variance. Translated to a behavioral outcome, the GARCH process indicates the tendency of an individual to be influenced by previous behavioral (in)consistencies regardless of the events. The ARCH process (captured by $\mathbf{A}$), in contrast, indicates how strongly an individual tends to be perturbed by previous random shocks.

pdBEKK. So far, the BEKK ensures positive definiteness and stationarity by imposing constraints on the eigenvalues of the Kronecker product $\mathbf{A} \otimes \mathbf{A} + \mathbf{B} \otimes \mathbf{B}$. Positive definiteness can be obtained by constraining one diagonal element in $\mathbf{A}$ and $\mathbf{B}$ to be positive, typically $a_{11} \geq 0$ and $b_{11} \geq 0$. This approach can result in negative diagonal elements – which is unproblematic in terms of estimation but poses a problem in terms of interpretation. Hence, [Rast, Martin, Liu, and Williams \(2020\)](#) additionally impose a positivity constraint on the diagonal element of the ARCH and GARCH parameters such that all elements in $\text{diag}(\mathbf{A}) \geq 0$ and $\text{diag}(\mathbf{B}) \geq 0$, to achieve the *positive diagonal* BEKK (pdBEKK) parameterization.

CCC. In the specifications discussed so far, the conditional covariance matrix $\mathbf{H}_t$ is modeled directly. An alternative approach is to first decompose $\mathbf{H}_t$ and then model its respective elements. We may define the conditional covariance as

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R} \mathbf{D}_t. \quad (6)$$

$\mathbf{D}_t$ is the $d \times d$ diagonal matrix with conditional standard deviations while $\mathbf{R}$ is the $d \times d$ constant correlation matrix.

The elements of $\text{diag}(\mathbf{D}_t) = \sqrt{h_{11,t}}, \dots, \sqrt{h_{dd,t}}$ are specified as univariate GARCH(p, q) models, as in (4), so that each conditional variance for row $i = 1, \dots, d$ (and column $j = i$) and $p = q = 1$

$$h_{ii,t} = c_i + \alpha_i \epsilon_{i,t-1}^2 + \beta_i (h_{ii,t-1}). \quad (7)$$

c_i is the constant variance term for a given univariate time-series i while α_i and β_i weigh the squared residuals and the conditional variance from the previous time point. It's important to note that (7) only models the conditional variance for each individual time-series.

DCC. In contrast to CCC, the dynamic conditional correlation (DCC [Engle, 2002](#); [Engle & Sheppard, 2001](#)) model introduces a time-varying and conditional correlations matrix. That is, Equation (6) now includes a subscript on $\mathbf{R}_t$:

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t, \quad (8)$$

with $\mathbf{D}_t$ being the $d \times d$ diagonal matrix with conditional standard deviations. $\mathbf{R}_t$ is the $d \times d$ dynamic conditional correlation matrix. Both matrices are subscripted by t indicating that both the standard deviations and the correlations are time-varying.

The submodel for the conditional standard deviations defined in Equation (7) remains the same as in the CCC parameterization. In addition, the dynamic conditional correlation matrix $\mathbf{R}_t$ is now further broken down into

$$\mathbf{R}_t = \mathbf{Q}_t^{*-1} \mathbf{Q}_t \mathbf{Q}_t^{*-1}, \quad (9)$$

where $\mathbf{Q}_t$ are $d \times d$ symmetric positive-definite matrices given as

$$\mathbf{Q}_t = (1 - a - b) \mathbf{S} + a(\mathbf{u}_{t-1} \mathbf{u}_{t-1}') + b \mathbf{Q}_{t-1}, \quad (10)$$

where a and b are ≥ 0 . Note that the conditional covariance matrix $\mathbf{Q}_t$ is derived from a standard linear GARCH(1,1) model specification (see equations 18 through 23 in [Engle, 2002](#)). $\mathbf{u}_t = \mathbf{D}_t^{-1} \boldsymbol{\epsilon}_t = (u_{1t}, \dots, u_{dt})'$ is a $d \times 1$ vector where $u_{it} = \epsilon_{it} / \sqrt{h_{it}}$ are the standardized residuals. Given that this specification is based on a standard GARCH

model, stationarity is obtained as long as $a + b < 1$. The scalars a and b are again parameters that weigh the past residual variance and the past conditional variance. $\mathbf{S}$ is the unconditional correlation of the standardized residuals.

Notably, [Engle and Sheppard \(2001\)](#) show that the use of a GARCH(p, q) model for the elements in $\mathbf{Q}_t$ ensures positive definiteness of $\mathbf{Q}_t$ and, given (9), positive definiteness of $\mathbf{R}_t$. Hence, $\mathbf{R}_t$ is now simply expressed in terms of the covariance matrix $\mathbf{Q}_t$ multiplied by the inverse of its standard deviations. That is, $\mathbf{Q}_t^*$ is a diagonal matrix with the square roots, $[\sqrt{q_{11}}, \dots, \sqrt{q_{ii}}]$, of the diagonal elements of $\mathbf{Q}_t$. This implies that the conditional covariances are given by $h_{ij,t} = q_{ij,t} \sqrt{h_{ii,t} h_{jj,t}} / \sqrt{q_{ii,t} q_{jj,t}}$.

In essence, the DCC results in two sets of GARCH parameters. Ones for the conditional variances (7) and the others for the conditional correlations (10). The covariances are defined in the separate model (10) where we obtain parameters that govern the relation among the conditional covariance among the time-series. Specifically, we will obtain an unconditional correlation $\mathbf{S}$ over the whole observation period as well as a and b , the ARCH and GARCH parameters for the conditional correlation matrix.

If one seeks to interpret the GARCH parameters, methods that are based on directly estimating $\mathbf{H}_t$ (such as BEKK and pdBEKK) and methods that operate on $\mathbf{D}_t$ and $\mathbf{R}_t$ (such as CCC and DCC) imply a slightly different mechanism under which past variance can spill over from one time series to the other and on how it propagates throughout the same time series. In terms of forecasting volatility, however, both approaches are known to perform well; with DCC slightly outperforming the others in the context of econometric forecasting (see e.g. [de Almeida et al., 2018](#)).

Evaluation of Predictive Accuracy

Before using our models to forecast future data points, we must evaluate their predictive accuracy for the observations at hand. [Eklund and Karlsson \(2007\)](#), for example, proposed to base the forecast weights on the predictive likelihood of each model. That is,

instead of computing the model weights from in-sample measures we approximate the out of sample performance. In a situation where we do not have access to future data, we need to rely on our current observations to give us an estimate of the model’s predictive performance. Once we have that information, we can average across the model space to either obtain averaged estimates or, as in the present case, obtain model averaged forecasts.

The general approach followed here is not a new one. In a seminal paper by Dawid (1984) where he coined the term “Prequential Forecast”, a combination of the two words *probability* and *sequential*, he follows the rationale of de Finetti (1975) by arguing that

... the only concept needed to express uncertainty is Probability, and that this is most meaningful associated with genuine observables. Consequently, the “forecasts” I shall be considering will be probability distributions over future events.

The third main characteristic of the approach to be considered is the sequential nature of the forecasting task. At any time, we can make a forecast for tomorrow. Come tomorrow, we can observe the outcomes of the events which were the subject of today’s forecasts and, with this additional experience to draw on, formulate our new forecasts for the following day. And so on, day after day, drawing on accumulating experience. I call this *prequential forecasting*. (Dawid, 1984, pp. 278)

Further, (Dawid, 1984, p. 279) notes that “Given any statistical model, we can convert it into a prequential forecasting system, in various possible ways – for example, by replacing the unknown parameter in the predictive distribution of the next observation by an estimate based on the data collected so far.”

Following Dawid (1984, 1992), from a Bayesian perspective, predictive accuracy can be understood as a probabilistic prediction of the unknown out-of-sample distribution. In the context of probability distributions, logarithmic scoring rules are known to be proper (Gneiting & Raftery, 2007) and in alignment with the Kullback-Leibler (KL) divergence

measure (Bernardo & Smith, 2000), used to compare true and estimated probability distributions. Briefly, the logarithmic score is the log predictive density (LPD) $\log [p(\tilde{y}|y)]$ for predictions of new data $\tilde{y}$. High expected LPD, coincides with low KL divergence and highest posterior probability (Bernardo & Smith, 2000). For a vector of observations $\mathbf{y}$, $p(\tilde{y}_i|\mathbf{y})$ is the predictive density of $\tilde{y}_i$ induced by the posterior distribution $p(\theta|\mathbf{y})$ (see eg. Gelman & Carlin, 2014). Hence, for a single data-point prediction accuracy can be defined as $\log \int p(\tilde{y}_i|\theta)p(\theta|\mathbf{y})d\theta$. In real applications, the future data $\tilde{y}_i$ are themselves unknown but we can define an expectation E_f under the true data distribution $f(y)$. For N data points we can define a measure of predictive accuracy, the expected out-of-sample log pointwise predictive density (ELPD):

$$\text{ELPD} = \sum_{i=1}^N \int f(\tilde{y}_i) \log [p(\tilde{y}_i|\mathbf{y})] d\tilde{y}_i. \quad (11)$$

Therefore, ELPD quantifies how close the estimated posterior predictive distribution $p(\tilde{y}_i|\mathbf{y})$, is to the true distribution $f(y)$. Accordingly, ELPD is maximized when the estimated distribution equals the true one. In absence of new data $\mathbf{y}$ and the true density of a single data point $f(\tilde{y})$, we need to approximate ELPD with a measure of predictive accuracy. A popular choice to approximate ELPD is to use LOO-CV where predictive accuracy is determined on the left-out data point y_{-i} . This approach requires additional computational steps after the model is fit. That is, for LOO-CV one needs to estimate n training sets in order to compute its statistic. This can be computationally very expensive making exact LOO-CV unfeasible for many real data applications. To alleviate that problem, Vehtari, Simpson, Gelman, Yao, and Gabry (2015) proposed an approximation to LOO-CV via Pareto smoothed importance sampling (PSIS) that circumvents the need to re-estimate models for the left-out data.

LOO-CV, and PSIS-LOO-CV, assume time independence in data points and their use is either restricted to models where data ordering does not play a role, or for the assessment of the non-time-dependent part of a time-series model. If LOO-CV is applied to evaluate predictive accuracy in time-series models the left out data-point is allowed to be

informed by future data-points and the resulting estimate will be too optimistic. Instead, as noted by [Eklund and Karlsson \(2007\)](#), for time-series, one could separate the available sample into a training and test set where the training set consists of the first L observations (or a window of observations of size L). Predictive accuracy would then be obtained by using the trained model on the left out data. However, for time-series of moderate size a leave-one out approach might yield better results as the training set increases over time. In fact, [Bürkner et al. \(2019\)](#) presented an approximation to ELPD based on leave-future-out cross-validation (LFO-CV), where the left out data points maintain their sequential order. With LFO-CV, we can obtain predictive accuracy for M step-ahead predictions (M -SAP) using the existing dataset.

As with LOO-CV, LFO-CV requires refitting the models for each M ahead left out future data-point. Practically, we can compute LFO-CV for a time series $y = (y_1, y_2, \dots, y_N)$ after having defined the minimum number of observations (L) before starting to predict data-points later in that sequence. There is no general rule for the size of the initial data set, but [Eklund and Karlsson \(2007\)](#) and others ([Feldkircher, 2012](#)) propose to set L to the 25th percentile of the time series and use the remaining data to compute predictive accuracy.

If we predict one-step ahead ($M = 1$) we can start with the computation of the predictive densities $p(y_{t+1}|y_{1:t})$ where $t \in \{L, \dots, N - 1\}$. Here we follow the notation used in [Vehtari and Ojanen \(2012\)](#) where $y_{1:t}$ is a short-hand for the vector $y_{1:t} = (y_1, \dots, y_t)$. Accordingly, the expected log predictive density (ELPD) can be defined as

$$\text{ELPD} = \sum_{t=L}^{N-1} \int f(\tilde{y}_{t+1}) \log [p(\tilde{y}_{t+1}|y_{1:t})] d\tilde{y}_{t+1}. \quad (12)$$

The problem encountered in Equation (11) persists in that we do not know the true data generating process $f(\tilde{y}_{t+1})$. Instead we need to approximate ELPD via a leave-future-out method with the observations at hand y_{t+1} by assuming that y_{t+1} is a realization from $f(\tilde{y}_{t+1})$ ([Vehtari & Ojanen, 2012](#)). Accordingly, [Bürkner et al. \(2019\)](#) replace $f(\tilde{y}_{t+1})$ with

its realization y_{t+1} to obtain

$$\text{ELPD}_{\text{LFO-CV}} = \sum_{t=L}^{N-1} \log [p(y_{t+1}|y_{1:t})]. \quad (13)$$

We can multiply $\text{ELPD}_{\text{LFO-CV}}$ by -2 , to yield LFO-CV and place it on a more familiar deviance scale. In that case, as with other information criteria (such as AIC [Akaike, 1973](#)), the goal would be to minimize LFO-CV. In analogy to LOO, we can compute the posterior predictive density $p(y_{t+1}|y_{1:t})$ from the posterior distribution $p(\theta|y_{1:t})$ but based only on the first t observations:

$$p(y_{t+1}|y_{1:t}) = \int p(y_{t+1}|y_{1:t}, \theta) p(\theta|y_{1:t}) d\theta. \quad (14)$$

The integral for the posterior predictive density can not be solved analytically but it can be approximated using S random draws $(\theta_{1:t}^{(1)}, \dots, \theta_{1:t}^{(S)})$ from a Monte-Carlo estimation ([Bürkner et al., 2019](#); [Vehtari et al., 2015](#)) as

$$p(y_{t+1}|y_{1:t}) \approx \frac{1}{S} \sum_{s=1}^S p(y_{t+1}|y_{1:t}, \theta_{1:t}^{(s)}). \quad (15)$$

Unfortunately, this approximation can be computationally costly and unfeasible in many applied settings as one would have to potentially fit a very large number of models.

Instead, expanding on previous work ([Vehtari, Gelman, & Gabry, 2017](#); [Vehtari et al., 2015](#)), [Bürkner et al. \(2019\)](#) introduce an approximation to the posterior predictive density using weights from Pareto smoothed importance sampling (PSIS [Vehtari et al., 2015](#)). In the context of time-series, we can approximate the posterior predictive density as (see also Equation (8) in [Bürkner et al., 2019](#))

$$p(y_{t+1}|y_{1:t}) \approx \frac{\sum_{s=1}^S w_t^{(s)} p(y_{t+1}|\theta^{(s)})}{\sum_{s=1}^S w_t^{(s)}}. \quad (16)$$

The PSIS weights w_t^s are computed from raw importance ratios (see [Vehtari et al., 2015](#)).

The $\theta^{(s)}$ are draws from the model based on the first L observations.

[Bürkner et al. \(2019\)](#) adopt a sampling scheme based on the stability of importance ratios that allows one to only refit a small number of the whole sequence of the time-series

that ranges from $L \geq 1$ to N , referred to as PSIS-LFO-CV. That is, the individual raw importance ratios are computed by the given observations y_t and the parameter θ for increasing t . At some point, the variability of the ratios will increase unduly which requires refitting the model to obtain a new $\theta^{(s)}$ based on the latest sample $1 : t$. This approach is repeated until we reach the end of the time-series at $t = N$. Hence, instead of fitting $N - (L + 1)$ models for one-step-ahead predictions, one only refits a handful of models, which makes LFO-CV feasible also for longer time-series. For example, instead of having to estimate 50 models with increasing sample size, one may only fit 5 models with this method.

At this point, we have defined a valid approach to assess predictive accuracy of time-series models. Based on the PSIS-LFO-CV method, we can now obtain model weights to, eventually, compute composite forecasts.

Ensemble Method

When comparing different models for the same problem at hand, a common approach is to select the best fitting model from a list of candidate models and proceed with the winner. However, in the presence of multiple models that could potentially explain our data, we can guard against possible misspecifications by *not* putting all the weight on one single model. While there are multiple approaches to weigh models, we focus on Bayesian model averaging, as it is the only approach that incorporates all sources of uncertainty present in the posterior.

Generally there are two popular approaches to averaging Bayesian models. One is based on the marginal likelihood (Draper, 1995; Hoeting, Madigan, Raftery, & Volinsky, 1999; Raftery, Madigan, & Volinsky, 1996) and is representative for settings where the true, generative process, is part of the tested models. This situation is referred to as the $\mathcal{M}$ -closed setting. The other approach is based on the predictive likelihood (Feldkircher, 2012; Yao, Vehtari, Simpson, & Gelman, 2018) for $\mathcal{M}$ -open settings where the true model

is not part of the model space.

While most models described in the literature use marginal likelihoods to obtain averaged forecasts (Chen & Gerlach, 2019; Liu & Maheu, 2009; Vrontos, 2012; Vrontos, Dellaportas, & Politis, 2000; Vrontos et al., 2000) these forecasts tend to be less accurate compared to forecasts from predictive likelihood methods (Feldkircher, 2012). As noted by Kapetanios, Labhard, and Price (2008), one would expect that the data-generating “true” model should always do better than an average of a range of “misspecified” models. In real data, however, the observed phenomena can never be fully characterized with the observed data and the resulting models may be best seen as good approximations. As such we would argue that often we will find ourselves in the $\mathcal{M}$ -open setting.

Model Stacking. As noted before, the selection of the ensemble method depends on whether we assume that the true generative model is contained in the list of models to select from or not. While MGARCH models seem to be promising approaches to explaining and forecasting variability in time-series, it is very unlikely that any of the parameterizations represent the true generative process. Hence, we will use a technique called stacking which was originally developed for averaging point estimates but has been shown to be optimal for using with predictive distributions.

Stacking is a model averaging method that results in a composite predictor by combining individual predictors using model weights that optimize a cross-validation criterion. While stacking is defined for minimizing mean squared error of the leave-one-out predictor, Le and Clarke (2017) showed that the stacking solution is asymptotically the Bayes solution if the point estimate is replaced with the predictive distribution for logarithm scoring rules. Based on these findings, Yao et al. (2018) extended the original stacking approach for point estimates and squared loss, to maximizing scoring rules – thus making it the appropriate approach for combining Bayesian predictive distributions, such as $\text{ELPD}_{\text{LFO-CV}}$.

Briefly, stacking has two general steps. Each M_k model, from a list of $k \in \{1, \dots, K\}$

models, is first fitted separately and a predictor criterion (e.g., LOO MSE) is obtained for each k model and data point t . Then, a weight is obtained for each model such that the convex combination of predictions (or distributions) minimizes the criterion. Instead of using LOO-CV, in this work, we use leave-future-out (LFO) predictive performance as the criterion for stacking. For example, if the predictive criterion is LFO-MSE, then we can define the loss function as the mean squared error of the future value from the expected (predicted) future value. That is, the LFO predictor $\hat{f}_k^{t+1}(y_{y_{1:t}}) = E_{\theta_k}(y_{t+1}|y_{1:t})$ is obtained for each future data point $t + 1$ and every model k in the list of models, and a set of stacking weights (w_s) are obtained that minimize the LFO-MSE

$$\hat{w}_s = \arg \min_{w_s} \sum_{t=1}^{N-1} \left(y_{t+1} - \sum_{k=1}^K w_{sk} \hat{f}_k^{t+1}(y_{1:t}) \right)^2. \quad (17)$$

The above corresponds to a stacking of means (Yao et al., 2018). Note that although this is a valid performance criterion, it fails to optimize the marginal predictive *distribution*, instead optimizing for expected values only. For example, the MGARCH directly models entire predictive distributions, comprising of expectations, variances, and correlations — ideally, a scoring rule would take into account these predictive distributions when assessing forecasting performance. If the goal is to stack a set of K predictive distributions built from the M_k models, then the criterion to optimize becomes a (proper) scoring function S , such as the log scoring rule. In this case, the goal is to find a set of weights for each K predictive distribution that maximizes the scoring rule. If a set is defined as $\mathcal{S}_1^K = \{w_s \in [0, 1]^K : \sum_{k=1}^K w_{sk} = 1\}$, then the weights we seek are

$$\hat{w}_s = \arg \max_{w_s \in \mathcal{S}_1^K} S \left(\sum_{k=1}^K w_{sk} p(\cdot|y, M_k), p_t(\cdot|y) \right). \quad (18)$$

In particular, when the logarithmic scoring rule is used, it is called *stacking of predictive*

distributions. For the ELPD_{LFO} the weights are given by

$$\hat{w}_s = \arg \max_{w_s \in \mathcal{S}_1^K} \frac{1}{N-1} \sum_{t=L}^{N-1} \log \sum_{k=1}^K w_{sk} p(y_{t+1} | y_{1:t}, M_k) \quad (19)$$

Model Averaged Forecasting

While it has been known for a long time that combining forecasts over candidate models increases forecasting performance with respect to the location of a future data point (Bates & Granger, 1969), recent research suggests that this is also the case for forecasted variances and covariances. Specifically, univariate (Li, Huang, & Zhang, 2013) and multivariate GARCH models (Vrontos, 2012; Vrontos et al., 2000; Vrontos, Dellaportas, & Politis, 2003) have been shown to be either at par or outperform the forecasts from the single best model. So far, however, averaged MGARCH forecasts are typically based on marginal likelihood BMA techniques to obtain model weights. To the best of our knowledge, we are the first to use the Bayesian stacking approach to average forecasts.

With the model weights at hand, we weigh the forecasted distributions accordingly. Given that the weights form a standard simplex (i.e., they are non-negative and sum to one), the marginal predictive distribution from the ensemble of models is equivalent to the weighted sum of each candidate model's predictive distribution:

$$\begin{aligned} p(y_{t+1} | y_{1:t}) &= \sum_{k=1}^K w_{sk} p_k(y_{t+1} | y_{1:t}) \\ &= \sum_k^K w_{sk} \int p_k(y_{t+1} | \theta_k, y_{1:t}) p_k(\theta_k | y_{1:t}) d\theta. \end{aligned}$$

Again, the integration required to attain the posterior predictive distributions analytically is untenable. Instead, samples can be drawn from the marginal predictive distribution of interest by weighting the samples of the candidate models' predictive distributions.

Assuming that the number of MCMC samples, S , is equivalent across all models:

$$\begin{aligned}\theta_k^{(s)} &\sim p_k(\theta_k|y_{1:t}) \\ \tilde{y}_{k,t+1}^{(s)} &\sim p_k(y_{t+1}|\theta_k^{(s)}) \\ \tilde{y}_{t+1}^{(s)}|y_{1:t} &= \sum_k^K w_{sk} \tilde{y}_{k,t+1}^{(s)}.\end{aligned}$$

Here, the resulting MCMC draws are samples from the ensemble posterior predictive distribution. From these draws, forecasted expected values and credible intervals can be computed for point estimates and uncertainty expression, respectively.

Forecasting Evaluation

Evaluating the forecasting performance for the future m realizations $\tilde{y}_{N+m}$ is straight forward once we obtain the data. A common approach is to compute the root mean squared forecasting error (RMSE) across all samples S from the whole predictive distribution

$$RMSE_m = \sqrt{\frac{\sum_{s=1}^S (\tilde{y}_{N+m}^s - y_{N+m})^2}{S}}. \quad (20)$$

The other components of interest, the conditional variance and covariance, however, are not quantities that we can directly observe. In order to evaluate the forecasting performance [Andersen, Bollerslev, Diebold, and Labys \(2003\)](#) proposed the use of realized variance (RV) estimates. They show that “the quadratic covariation constitutes, in theory, a unique and invariant ex-post realized volatility measure that is essentially model free” and that quadratic variation can be approximated by “cumulating cross-products of high-frequency returns.” (p. 585).

There is a limitation to this approach: In order for the RV to be computed, we need data on a more granular scale than the one we used to fit our model. For example, in our illustration, we average walking speed and computer use across two weeks. While we fit our MGARCH models to the averaged data where one unit change represents two weeks, we can compute the RV on the basis of 14 days worth of data. That is, each two week period

is divided into $\Delta = 1/14$ intervals, where Δ coincides to one day. The daily observations are defined as $y_{t,j} = y_{t+(j-1)\Delta}$, $j = 1, \dots, 14$, $y_t = E_j(y_{t,j}) = \mu_t$, and $\mathbf{y}_t$ is the $J = 1/\Delta$ by D matrix of daily observations (rows) in the two week window for time period t for the D series (columns). RV can then be defined as (Andersen et al., 2003; Liu & Maheu, 2009):

$$RV_t = \frac{1}{J^2}(\mathbf{y}_t - \mu_t)'(\mathbf{y}_t - \mu_t). \quad (21)$$

Simply put, RV_t is estimated as the covariance of the observations in time window t , divided by J . This yields the expected covariance of the aggregate value estimated by MGARCH. For one-ahead forecasting, we will refer to RV_t as RV , such that RV is the realized variability of the $N + 1$ period.

Having an RV estimate now allows us to compare the realized and the forecasted variability $\hat{\mathbf{H}}_m$. Given that we are comparing two $d \times d$ dimensional covariance matrices, we can estimate Kullback-Leibler divergence (D_{KL}) for each replicate $s \in \{1 \dots S\}$ and each forecasted period m as

$$D_{KL,m,s} = \frac{1}{2} \left[\text{tr}(\hat{\mathbf{H}}_m^{-1} RV) + \log \frac{\det \hat{\mathbf{H}}_m}{\det RV} - d + (\hat{\boldsymbol{\mu}} - \boldsymbol{\mu})' \hat{\mathbf{H}}_m^{-1} (\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) \right]. \quad (22)$$

This provides a measure of information loss between RV and the forecasted covariance matrix $\hat{\mathbf{H}}_m$. $D_{KL,m,s}$ yields a distribution across all replicated S samples and we will report the means as well as the lower 2.5 and upper 97.5 quantiles.

Illustrative Example

In the illustrative example we estimate four MGARCH models (BEKK, pdBEKK, DCC, and CCC). The models reported here also estimate an ARMA(1,1) for the expected means μ alongside the MGARCH for the conditional variances so that $\mathbf{y}_t \sim MST(\boldsymbol{\mu}_t = \boldsymbol{\phi}_0 + \boldsymbol{\phi} \mathbf{y}_{t-1} + \boldsymbol{\theta}(\mathbf{y}_{t-1} - \boldsymbol{\mu}_{t-1}), \mathbf{H}_t, \nu)$ with $\mathbf{H}_t$ being one of the four parameterizations. Given that we assume a multivariate Student- t (MST) distribution, we restrict the degrees of freedom ν in the MST to be larger than 2, to ensure existence of the variance matrix (Bauwens & Laurent, 2005; Rast et al., 2020). Moreover, to ensure

unbiased estimates, the ARMA and MGARCH parameters are estimated simultaneously (Carnero & Eratalay, 2014). Once we estimate all models, we compute predictive accuracy according to the PSIS-LFO-CV method and obtain stacking weights. These weights are used to create the model averaged composite forecasts. The weighted forecasts for the expected means and conditional covariance are compared with forecasts from the single best model. Here, we will forecast five periods ahead at once – alternatively one could use a rolling window and forecast only one period ahead and repeat that five times.

The data for this illustration come from the longitudinal Intelligent Systems for Assessing Aging Change (ISAAC) study (Dodge et al., 2015, weekly averaged data is openly available following the links in that publication). The study comprised up to 265 participants from the Portland, Oregon, metropolitan area and covered up to 4 years of daily measurements. All participants were 70 years of age or older at time of recruitment in 2007. They were all living individually (without the need of a caregiver), were in average health and not demented. All participants agreed to be in-home monitored by a wide array of unobtrusive sensors. Here we focus on two variables, *walking speed* and *computer use*, that have been shown to be sensitive markers of cognitive change (Dodge et al., 2015).

Walking speed was calculated using a line of four motion sensors positioned in a series on the ceiling. Speed was measured as m/s and, for the purpose of the present illustration we averaged walking speed across 14 days. The other variable, minutes of computer use, was averaged as well over 14 days. Both variables were z -standardized within individuals.

We selected two individuals who had complete data on their time series with a length of at least 105 consecutive two-week periods. All models were fit to the first 100 periods, covering 3.8 years. The estimation of predictive accuracy was based on that training sample. The remaining 5 periods, covering approximately 2.5 months, served as test sample to compare the forecasting performance via RMSE and D_{KL} . Due to the need to also estimate realized variability (RV) for the test sample (to compute D_{KL}), the participants in question needed to have full *daily* information during that period of time.

That is, RV is computed for every two weeks on the basis of the underlying daily high-frequency data. We identified two individuals who fulfilled all these desiderata. These individuals will be referred to as Individual 1 and Individual 2. Figure 1 shows the time series of both individuals across 105 time points. The blue lines represent walking speed and the red lines computer use on a z metric. The vertical line represents the point where the time-series are split into training and test sample to be forecasted.

All models were estimated using the **bmgarch**¹ package in R. All prior specifications are described in detail elsewhere (Rast et al., 2020).

Results. The first step was to fit the sequence of four models, BEKK, pdBEKK, DCC, and CCC separately to both individuals and compute the model weights from the $\text{ELPD}_{\text{LFO-CV}}$ method. The quality of all models was assessed by assuring that all chains mixed well and converged on the typical set. This was done visually for the a subset of parameters in each model and by keeping all potential scale reduction factors $\hat{R} \leq 1.1$ (Gelman & Rubin, 1992). Overall, for every fitted model, we used four chains with 2000 iterations, where the first 1000 were discarded as warm-up. Given that the focus here is on the forecasted portion of the models, we will not report the specific estimated parameters.

After fitting all four models, we computed the posterior likelihood and the model weights with the stacking method. Table 1 shows the stacking weights (w_s) for both individuals along with LFO-CV estimates and the differences in the $\text{ELPD}_{\text{LFO-CV}}$'s with respect to the best model.

It is noteworthy that the pattern of weights was distinct for both individuals even though CCC happened to come out as the model with the highest weight and the lowest LFO-CV. Anecdotally, when running these models on other participants (who did not have full data to compute RV), the weights came out differently for each participant, with different models receiving the most weight. A model selection approach based on LFO-CV would select the model with the smallest deviance. Note, however, that the standard errors

¹ Available at <https://github.com/ph-rast/bmgarch>

of the model differences were relatively large in se metric which means that the predictive accuracy among these models is not clearly distinguishable. This is especially pronounced for Individual 2, where se_{Δ} is very large with respect to the estimated difference ($ELPD_{\Delta}$).

With the stacking weights at hand, we computed the composite forecasts. That is, the weights were applied to the full distribution of forecasted draws for each model and summed up to yield the final weighted forecast samples. We forecasted the expected realizations $y_{N+1:5}$ as well as the conditional variances, and covariances in $\hat{\mathbf{H}}_{N+1:5}$ for five periods ahead, covering a total of 10 weeks.

Forecasting Means. Tables 2 and 3 contain the forecasts, $\hat{y}_{N+1:5}$, for both Individuals 1 and 2. Alongside the actually observed values, we report the averaged forecasts in the top section and the forecasts from the single best model (CCC for both individuals) in the lower section. Given that we obtain a whole posterior predictive distribution of forecasts, we report the mean, standard deviation, median and lower 2.5 and upper 97.5 quantiles.

Generally, means and medians were almost identical which suggests that the distributions were mostly symmetrical. The means across the averaged forecast and the forecast from the CCC were very close. However the SD 's and the 95% predictive intervals (PrI)'s were always substantially larger for the single CCC model, across both variables. In terms of RMSE, the lower precision in the CCC estimates resulted in larger values. Notably, the RMSE's across all time points for the composite forecasts were always smaller, indicating higher accuracy of these weighted estimates.

Figure 2 shows the averaged forecasts with 95% PrI bands (turquoise) together with the previous 50 observations (red) for both individuals and both variables. The black dots show the actual, observed data values.

Forecasting the Conditional Variances and Covariances. Next, we report the forecasted variances in Tables 4 and 5, alongside with the RV estimate. RV is the estimate of the realized variance from the 14 data points that make up each of five the 2-week

periods. These estimates will be used to compute the distance in terms of D_{KL} of the forecast to the actual target, the RV.

Again, the estimates from both, the model averaged and the single best model, were very similar with respect to their means and the medians were mostly at the same location as the means. Yet, the SD 's and the PrI's from the model averaged forecasts were smaller than those from the single best model. This was the case for both individuals and both variables and again, all SD 's from the model averaged forecasts were smaller than those from the CCC model – the same holds for the 95% interval width of the PrI's.

Figure 3 shows the resulting forecasts from the averaged model. The turquoise band represents the 95% PrI of the forecasts. The black dots show the RV values. Across both variables, the forecasted variability seemed to be slightly higher than the observed one. Especially walking speed tended to produce larger forecasted variability compared to the RV estimate.

We now turn to the forecasted conditional covariances among walking speed and PC usage. All estimates are reported in Tables 6 and 7. A similar pattern emerged as the one reported before, with respect to the means and medians suggesting that the distributions were symmetrical. Individual 2 showed a large discrepancy among SD 's and PrI's across the two modes of forecasting where ensemble forecasts were much more precise compared to the forecasts from the CCC model. This discrepancy in the spread of the forecasting distribution was not evident in Individual 1. Figure 4 shows the observed and predicted forecasts for the conditional covariance together with the 95% PrI's of the expected averaged forecasts, marked by the turquoise band. Notably, the forecasts for Individual 1 were expected to be mostly positive, when in fact, most of the observed covariances were negative or practically zero (see panel A). The forecasts for Individual 2 seemed to be largely in line with the observations.

In order to compare the forecasts of the conditional covariance matrix $\hat{\mathbf{H}}_m$ from both, the model averaged approach and the CCC model, we computed D_{KL} divergences for each

forecasted time point. D_{KL} computes a distance between the realized covariance matrix (RCV_m) for a given period and the forecasted covariance matrix $\hat{\mathbf{H}}_m$. We use RCV , a $d \times d$ matrix that corresponds to the dimension of $\mathbf{H}$, to underscore that we compute the distance among two covariance matrices. Similarly to the RMSE, the model averaged forecasts outperformed the forecasted covariance matrix for all but one condition. That is, at best, the CCC model was at par with the averaged forecasts (slightly better in fact), as seen in Table 8, for period 101 in Individual 2, where the D_{KL} for the CCC model (denoted as D_{CCC}) outperformed the one of from the averaged models (KL_w). Overall, the differences in the average D_{KL} estimates were sizable, especially for Individual 1. Moreover, for that same individual, there was a large difference in the size of the credible intervals (CrI) for D_{KL} , as illustrated in Figure 5. The red 95% credible band from the CCC model almost completely overlapped with the turquoise D_{KL} estimate for the model averaged solution and was up to three times larger. The D_{KL} 's for Individual 2 were largely indistinguishable.

Discussion

The recent availability of high frequency data from intensive longitudinal designs or from personal devices and wearable sensors in the behavioral sciences, has provided new opportunities for not only explaining behavior *post-hoc* but generating expectations about future behavior. As such, forecasting, in addition to explaining behavior opens up the door not only to validating a model but to also testing its ability to predict unseen behavior. We would concur with Dawid (1984) that “One of the major purposes of statistical analysis is to make *forecasts for the future*.” (p. 278). Especially, forecasting individual behavior might turn out to be fruitful in applied settings. For example, one might generate expectations about future behavior on cognitive or physical fitness for a single individual. As long as forecasted behavior and future observations coincide we may conclude that this individual is stable and behaving in a consistent and predictable manner. In turn, if the forecast and observed behavior are incongruent, either in the actual realization or the forecasted

(co-)variation of variables, we may conclude that this persons behavior has changed in an unexpected manner. Behavioral inconsistency may be linked to more dramatic changes down the road ([MacDonald, Hultsch, & Dixon, 2008](#); [Rast, MacDonald, & Hofer, 2012](#)).

Along these lines, of forecasting behavioral (in)consistency, the focus here was mainly on models that are able to generate conditional variances, or volatility, forecasts from multiple time-series for individuals. The other great promise intensive data holds, is the ability to tailor models to individuals and to being able to making statements about a specific individual. Population based models, even if they include highly granular random effects do not allow one to make statements about single individuals as the necessary ergodicity assumption does not typically hold ([Hamaker, 2012](#); [Molenaar, 2004](#)).

As outlined in the introduction, models that forecast volatility are fundamental in econometrics, and they can be reparameterized in a way that one can interpret their parameters in a behavioral context ([Rast et al., 2020](#)). First and foremost, however, MGARCH models are being developed with mainly one goal in mind: to forecast the volatility and variance spill overs in multivariate time-series ([Bauwens et al., 2006](#)). As such, these models seem ideal candidates for forecasting variability in individual time-series.

Here, we focused not only on the forecasting aspect of these models but also introduced a Bayesian model averaging approach that turned out to deliver more precise estimates compared to a single model. To this end, we presented a novel aspect in computing the stacking weights for time-series, by maximizing the scoring rule based on the $ELPD_{LFO-CV}$ instead of $ELPD_{LOO-CV}$. The key to being able to apply these kinds of algorithms to real data is to reduce computation time. The PSIS-LFO-CV method ([Bürkner et al., 2019](#)) makes this approach feasible. For example, running all four model parameterizations and computing the stacking weights via the PSIS-LFO-CV algorithm, which refits approximately 5-7 models per parameterization, takes approximately 2 hours with four chains of each 2000 iterations and a total of 100 time-points on a standard desktop PC.

Model averaging has the known advantage that composite forecasts tend to be more accurate but it also underscores the differences in individuals. That is, model weights can allocate higher probability to different models that captures the idiosyncrasies of one given individual. In fact, the model weights that we obtained here indicated that both individuals behavior was predicted by a unique combination of model weights. Had we chosen a model selection approach, we would have settled on the CCC giving it 100% of the weight, when for example, Individual 2's model weight was at less than $p \approx .4$ for that parameterization. As noted earlier, every individual that we modeled (not reported) yielded a different configuration of model weights that was distinct from another. This underscores the necessity for tailored approach to modeling individuals instead of a one-size fits all strategy.

In terms of quality, the model averaged forecast always yielded higher predictive accuracy for the actual realizations and the conditional variances and covariances, compared to the single best model. This feature of averaged forecasts has been reported before, with other methods and other variables, and the advantage of combining forecasts to obtain better predictive performance has long been known ([Bates & Granger, 1969](#)). Generally, the forecasts seemed to be in line with the observations but it seemed to be a more difficult task to predict the RV instead of the actual datapoint. The forecasts for the variances and covariances seemed to be generally a bit too large, especially so for Individual 1.

While we have noted that these models and this approach might be fruitfully exploited for behavioral forecasting, it needs to be pointed out that the data demand can be quite drastic. Especially if one seeks to compare the forecasts to RV estimates, as we did here, the data demand quickly rises beyond the size and granularity of most behavioral data sets. However, studies that operate with sensor data or multiple daily measurements over an extended period of time should be able to use this approach. That is, while a time-series length of 100 for the training set is probably at the lower end of the data needs, the generated forecasts were still surprisingly close to the observed values.

References

- Akaike, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. In B. N. Petrov & F. Csaki (Eds.), *Second international symposium on information theory* (pp. 267–281).
- Almeida, D. M., Piazza, J. R., & Stawski, R. S. (2009). Interindividual differences and intraindividual variability in the cortisol awakening response: An examination of age and gender. *Psychology and Aging*, 24(4), 819–827. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/a0017910> doi: 10.1037/a0017910
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, 71(2), 579–625. Retrieved from <http://doi.wiley.com/10.1111/1468-0262.00418> doi: 10.1111/1468-0262.00418
- Bates, J. M., & Granger, C. W. J. (1969, dec). The Combination of Forecasts. *Journal of the Operational Research Society*, 20(4), 451–468. Retrieved from <https://www.tandfonline.com/doi/full/10.1057/jors.1969.103> doi: 10.1057/jors.1969.103
- Bauwens, L., & Laurent, S. (2005). A New Class of Multivariate Skew Densities, With Application to Generalized Autoregressive Conditional Heteroscedasticity Models. *Journal of Business & Economic Statistics*, 23(3), 346–354. Retrieved from <http://www.tandfonline.com/doi/abs/10.1198/073500104000000523> doi: 10.1198/073500104000000523
- Bauwens, L., Laurent, S., & Rombouts, J. V. K. (2006). Multivariate GARCH models: a survey. *Journal of Applied Econometrics*, 21(1), 79–109. Retrieved from <http://doi.wiley.com/10.1002/jae.842> doi: 10.1002/jae.842
- Bernardo, J. M., & Smith, A. F. M. (2000). *Bayesian Theory*. Chichester, UK: Wiley.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/0304407686900631> doi:

10.1016/0304-4076(86)90063-1

Bollerslev, T., Engle, R. F., & Wooldridge, J. M. (1988). A Capital Asset Pricing Model with Time-Varying Covariances. *Journal of Political Economy*, 96(1), 116–131.

Retrieved from <http://doi.wiley.com/10.3982/ECTA9240> doi: 10.1086/261527

Bürkner, P.-C., Gabry, J., & Vehtari, A. (2019). Approximate leave-future-out cross-validation for Bayesian time series models. Retrieved from

<http://arxiv.org/abs/1902.06281>

Carnero, M. A., & Eratalay, M. H. (2014). Estimating VAR-MGARCH models in multiple steps. *Studies in Nonlinear Dynamics & Econometrics*, 18(3), 339–365. Retrieved

from [https://www.degruyter.com/view/j/snde.2014.18.issue-3/](https://www.degruyter.com/view/j/snde.2014.18.issue-3/snde-2012-0065/snde-2012-0065.xml)

[snde-2012-0065/snde-2012-0065.xml](https://www.degruyter.com/view/j/snde.2014.18.issue-3/snde-2012-0065/snde-2012-0065.xml) doi: 10.1515/snde-2012-0065

Castro-Schilo, L., & Ferrer, E. (2013). Comparison of Nomothetic Versus

Idiographic-Oriented Methods for Making Predictions About Distal Outcomes From Time Series Data. *Multivariate Behavioral Research*, 48(2), 175–207. Retrieved from

<http://www.tandfonline.com/doi/abs/10.1080/00273171.2012.736042> doi:

10.1080/00273171.2012.736042

Chen, W. Y., & Gerlach, R. H. (2019, oct). Semiparametric GARCH via Bayesian Model Averaging. *Journal of Business & Economic Statistics*, 1–16. Retrieved from

<https://doi.org/10.1080/07350015.2019.1668796>[https://](https://www.tandfonline.com/doi/full/10.1080/07350015.2019.1668796)

[www.tandfonline.com/doi/full/10.1080/07350015.2019.1668796](https://doi.org/10.1080/07350015.2019.1668796) doi:

10.1080/07350015.2019.1668796

Cleveland, W. S., Denby, L., & Liu, C. (2002). *Random scale effects* (Tech. Rep.). Murray Hill, NJ: Statistics Research Dept. Bell Labs. Retrieved from

<https://citeseerx.ist.psu.edu/viewdoc/>

[download?doi=10.1.1.141.2403&rep=rep1&type=pdf](https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.141.2403&rep=rep1&type=pdf)

Dawid, A. P. (1984). The Prequential Approach. *Journal of the Royal Statistical Society. Series A (General)*, 147(2), 278–292. Retrieved from

<https://www.jstor.org/stable/2981683>

Dawid, A. P. (1992). Prequential Data Analysis. *Lecture Notes-Monograph Series*, 17, 113–126. Retrieved from www.jstor.org/stable/4355629

de Almeida, D., Hotta, L. K., & Ruiz, E. (2018). MGARCH models: Trade-off between feasibility and flexibility. *International Journal of Forecasting*, 34(1), 45–63.

Retrieved from

<https://www.sciencedirect.com/science/article/pii/S0169207017300894>

doi: 10.1016/J.IJFORECAST.2017.08.003

de Finetti, B. (1975). *Theory of Probability*. New York: Wiley.

Dodge, H. H., Zhu, J., Mattek, N. C., Austin, D., Kornfeld, J., & Kaye, J. A. (2015, sep).

Use of High-Frequency In-Home Monitoring Data May Reduce Sample Sizes Needed in Clinical Trials. *PLOS ONE*, 10(9), e0138095. Retrieved from

<https://dx.plos.org/10.1371/journal.pone.0138095> doi:

10.1371/journal.pone.0138095

Draper, D. (1995, jan). Assessment and Propagation of Model Uncertainty. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 45–70. Retrieved from

<http://doi.wiley.com/10.1111/j.2517-6161.1995.tb02015.x> doi:

10.1111/j.2517-6161.1995.tb02015.x

Eklund, J., & Karlsson, S. (2007, apr). Forecast Combination and Model Averaging Using Predictive Measures. *Econometric Reviews*, 26(2-4), 329–363. Retrieved from

<http://www.tandfonline.com/doi/abs/10.1080/07474930701220550> doi:

10.1080/07474930701220550

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4), 987. Retrieved from

<https://www.jstor.org/stable/1912773?origin=crossref> doi: 10.2307/1912773

Engle, R. F. (2002). Dynamic Conditional Correlation. *Journal of Business & Economic*

- Statistics*, 20(3), 339–350. Retrieved from
<http://www.tandfonline.com/doi/abs/10.1198/073500102288618487> doi:
10.1198/073500102288618487
- Engle, R. F., & Kroner, K. F. (1995). Multivariate simultaneous generalized arch.
Econometric Theory, 11(1), 122–150. Retrieved from
http://www.journals.cambridge.org/abstract{_}S0266466600009063 doi:
10.1017/S0266466600009063
- Engle, R. F., & Sheppard, K. (2001). *Theoretical and Empirical properties of Dynamic Conditional Correlation Multivariate GARCH* (Tech. Rep.). Cambridge, MA: National Bureau of Economic Research. Retrieved from
<http://www.nber.org/papers/w8554.pdf> doi: 10.3386/w8554
- Feldkircher, M. (2012, jul). Forecast Combination and Bayesian Model Averaging: A Prior Sensitivity Analysis. *Journal of Forecasting*, 31(4), 361–376. Retrieved from
<http://doi.wiley.com/10.1002/for.1228> doi: 10.1002/for.1228
- Ferrer, E., Gonzales, J. E., & Steele, J. (2013). Intra- and Interindividual Variability of Daily Affect in Adult Couples. *GeroPsych: The Journal of Gerontopsychology and Geriatric Psychiatry*, 26(3), 163–172. doi: 10.1024/1662-9647/a000095
- Fiske, D. W., & Rice, L. (1955). Intra-individual response variability. *Psychological Bulletin*, 52(3), 217–250. Retrieved from
<http://doi.apa.org/getdoi.cfm?doi=10.1037/h0045276> doi: 10.1037/h0045276
- Gelman, A., & Carlin, J. (2014). Beyond Power Calculations. *Perspectives on Psychological Science*, 9(6), 641–651. Retrieved from
<http://journals.sagepub.com/doi/10.1177/1745691614551642> doi:
10.1177/1745691614551642
- Gelman, A., & Rubin, D. B. (1992, nov). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4), 457–472. Retrieved from
<http://projecteuclid.org/euclid.ss/1177011136> doi: 10.1214/ss/1177011136

- Gneiting, T., & Raftery, A. E. (2007, mar). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477), 359–378.
Retrieved from
<http://www.tandfonline.com/doi/abs/10.1198/016214506000001437> doi:
10.1198/016214506000001437
- Hamaker, E. L. (2012). Why researchers should think "within-person": A paradigmatic rationale. In *Handbook of research methods for studying daily life* (pp. 43–61).
- Hamaker, E. L., Asparouhov, T., Brose, A., Schmiedek, F., & Muthén, B. (2018). At the Frontiers of Modeling Intensive Longitudinal Data: Dynamic Structural Equation Models for the Affective Measurements from the COGITO Study. *Multivariate Behavioral Research*, 53(6), 820–841. Retrieved from
<https://www.tandfonline.com/doi/full/10.1080/00273171.2018.1446819> doi:
10.1080/00273171.2018.1446819
- Hedeker, D., & Mermelstein, R. J. (2007). Mixed effects regression models with heterogeneous variance: Analyzing ecological momentary assessment (EMA) data of smoking. In T. D. Little, J. A. Bovaird, & N. A. Card (Eds.), (pp. 183–206). Mahwah, NJ.: Lawrence Erlbaum.
- Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2008). An Application of a Mixed-Effects Location Scale Model for Analysis of Ecological Momentary Assessment (EMA) Data. *Biometrics*, 64(2), 627–634. Retrieved from
<http://doi.wiley.com/10.1111/j.1541-0420.2007.00924.x> doi:
10.1111/j.1541-0420.2007.00924.x
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian Model Averaging: A Tutorial. *Statistical Science*, 14(4), 382–417. Retrieved from
https://projecteuclid.org/download/pdf_{_}1/euclid.ss/1009212519
- Kapetanios, G., Labhard, V., & Price, S. (2008). Forecasting Using Bayesian and Information-Theoretic Model Averaging. *Journal of Business & Economic Statistics*,

- 26(1), 33–41. Retrieved from
<http://www.tandfonline.com/doi/abs/10.1198/073500107000000232> doi:
10.1198/073500107000000232
- Le, T., & Clarke, B. (2017). A Bayes Interpretation of Stacking for M-Complete and M-Open Settings. *Bayesian Analysis*, 12(3), 807–829. Retrieved from
<https://projecteuclid.org/euclid.ba/1473276261> doi: 10.1214/16-BA1023
- Leckie, G., French, R., Charlton, C., & Browne, W. (2014). Modeling Heterogeneous Variance-Covariance Components in Two-Level Models. *Journal of Educational and Behavioral Statistics*, 39(5), 307–332. Retrieved from
<http://journals.sagepub.com/doi/10.3102/1076998614546494> doi:
10.3102/1076998614546494
- Li, Y., Huang, W.-P., & Zhang, J. (2013, sep). Forecasting volatility in the Chinese stock market under model uncertainty. *Economic Modelling*, 35, 231–234. Retrieved from
<http://dx.doi.org/10.1016/j.econmod.2013.07.006> doi:
10.1016/j.econmod.2013.07.006
- Liu, C., & Maheu, J. M. (2009). Forecasting realized volatility: a Bayesian model-averaging approach. *Journal of Applied Econometrics*, 24(5), 709–733. Retrieved from <http://doi.wiley.com/10.1002/jae.1070> doi: 10.1002/jae.1070
- MacDonald, S. W. S., Hultsch, D. F., & Dixon, R. A. (2008). Predicting impending death: Inconsistency in speed is a selective and early marker. *Psychology and Aging*, 23, 595–607. Retrieved from
<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2562863/> doi:
10.1037/0882-7974.23.3.595
- Molenaar, P. C. M. (2004). A Manifesto on Psychology as Idiographic Science: Bringing the Person Back Into Scientific Psychology, This Time Forever. *Measurement: Interdisciplinary Research & Perspective*, 2(4), 201–218. Retrieved from
http://www.tandfonline.com/doi/abs/10.1207/s15366359mea0204_{_}1 doi:

10.1207/s15366359mea0204_1

- Molenaar, P. C. M., Huizenga, H. M., & Nesselroade, J. R. (2003). The relationship between the structure of interindividual and intraindividual variability: A theoretical and empirical vindication of developmental systems theory. In U. M. Staudinger & U. Lindenberger (Eds.), (pp. 339–360). Dordrecht, Netherlands: Kluwer Academic Publishers. Retrieved from http://link.springer.com/chapter/10.1007/978-1-4615-0357-6_{ }15
- Molenaar, P. C. M., Sinclair, K. O., Rovine, M. J., Ram, N., & Corneal, S. E. (2009). Analyzing developmental processes on an individual level using nonstationary time series modeling. *Developmental psychology*, 45(1), 260–71. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19210007> doi: 10.1037/a0014170
- Natesan, P., & Hedges, L. V. (2017). Bayesian unknown change-point models to investigate immediacy in single case designs. *Psychological Methods*, 22(4), 743–759. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/met0000134> doi: 10.1037/met0000134
- Nesselroade, J. R., & Molenaar, P. C. M. (2010). Emphasizing Intraindividual Variability in the Study of Development Over the Life Span Concepts and Issues. In R. M. Lerner, M. E. Lamb, & A. M. Freund (Eds.), *The handbook of life-span development* (pp. 30–54). Hoboken, NJ, USA: John Wiley & Sons, Inc. doi: 10.1002/9780470880166
- Raftery, A. E., Madigan, D., & Volinsky, C. T. (1996). Accounting for Model Uncertainty in Survival Analysis Improves Predictive Performance. *Bayesian Statistics*.
- Ram, N., & Gerstorf, D. (2009). Time-structured and net intraindividual variability: Tools for examining the development of dynamic characteristics and processes. *Psychology and Aging*, 24(4), 778–791. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/a0017915> doi: 10.1037/a0017915
- Rast, P., & Ferrer, E. (2018). A Mixed-Effects Location Scale Model for Dyadic

- Interactions. *Multivariate Behavioral Research*, 53(5), 756–775. Retrieved from <https://www.tandfonline.com/doi/full/10.1080/00273171.2018.1477577> doi: 10.1080/00273171.2018.1477577
- Rast, P., Hofer, S. S., & Sparks, C. (2012). Modeling Individual Differences in Within-Person Variation of Negative and Positive Affect in a Mixed Effects Location Scale Model Using BUGS/JAGS. *Multivariate Behavioral Research*, 47(2), 177–200. doi: 10.1080/00273171.2012.658328
- Rast, P., MacDonald, S. W. S., & Hofer, S. M. (2012). Intensive Measurement Designs for Research on Aging. *GeroPsych*, 25(2), 45–55. Retrieved from <http://econtent.hogrefe.com/doi/abs/10.1024/1662-9647/a000054> doi: 10.1024/1662-9647/a000054
- Rast, P., Martin, S. R., Liu, S., & Williams, D. R. (2020). A New Frontier for Studying Within-Person Variability: Bayesian Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *PsyArXiv*. Retrieved from <https://psyarxiv.com/j57pk/>
- Rindskopf, D. (2014). Nonlinear Bayesian analysis for single case designs. *Journal of School Psychology*, 52(2), 179–189. Retrieved from <https://linkinghub.elsevier.com/retrieve/pii/S0022440513001143> doi: 10.1016/j.jsp.2013.12.003
- Rush, J., Rast, P., Almeida, D. M., & Hofer, S. M. (2019). Modeling long-term changes in daily within-person associations: An application of multilevel SEM. *Psychology and Aging*, 34(2), 163–176. Retrieved from <http://doi.apa.org/getdoi.cfm?doi=10.1037/pag0000331> doi: 10.1037/pag0000331
- Shadish, W. R. (2014). Statistical Analyses of Single-Case Designs: The Shape of Things to Come. *Current Directions in Psychological Science*, 23(2), 139–146. Retrieved from <http://journals.sagepub.com/doi/10.1177/0963721414524773> doi:

10.1177/0963721414524773

- Sullivan, K. J., Shadish, W. R., & Steiner, P. M. (2015). An introduction to modeling longitudinal data with generalized additive models: Applications to single-case designs. *Psychological Methods*. doi: 10.1037/met0000020
- Tsay, R. S. (2013). *An Introduction to Analysis of Financial Data with R*.
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27, 1413–1432. doi: 10.1007/s11222-016-9696-4
- Vehtari, A., & Ojanen, J. (2012). A survey of Bayesian predictive methods for model assessment, selection and comparison. *Statistics Surveys*, 6, 142–228. Retrieved from <http://projecteuclid.org/euclid.ssu/1356628931> doi: 10.1214/12-SS102
- Vehtari, A., Simpson, D., Gelman, A., Yao, Y., & Gabry, J. (2015). Pareto Smoothed Importance Sampling. Retrieved from <http://arxiv.org/abs/1507.02646>
- Vrontos, I. D. (2012). Evidence for hedge fund predictability from a multivariate Student's t full-factor GARCH model. *Journal of Applied Statistics*, 39(6), 1295–1321. doi: 10.1080/02664763.2011.644771
- Vrontos, I. D., Dellaportas, P., & Politis, D. N. (2000). Full bayesian inference for GARCH and EGARCH models. *Journal of Business and Economic Statistics*, 18(2), 187–198. doi: 10.1080/07350015.2000.10524861
- Vrontos, I. D., Dellaportas, P., & Politis, D. N. (2003, dec). A full-factor multivariate GARCH model. *Econometrics Journal*, 6(2), 312–334. Retrieved from <https://academic.oup.com/ectj/article/6/2/312-334/5073961> doi: 10.1111/1368-423X.t01-1-00111
- Williams, D. R., Mulder, J., Rouder, J., & Rast, P. (2020). Beneath the Surface: Unearthing Within-Person Variability and Mean Relations with Bayesian Mixed Models. *Psychological Methods*, 1–24. Retrieved from <https://psyarxiv.com/gwatq/> doi: 10.31234/OSF.IO/GWATQ

Woodrow, H. (1932). Quotidian variability. *Psychological Review*, *39*(3), 245–256. doi: 10.1037/h0073076

Yao, Y., Vehtari, A., Simpson, D., & Gelman, A. (2018, sep). Using Stacking to Average Bayesian Predictive Distributions. *Bayesian Analysis*, *13*(3), 917–1007. Retrieved from <https://projecteuclid.org/euclid.ba/1516093227> doi: 10.1214/17-BA1091

Table 1

Stacking Weights Across Four MGARCH Model Parameterizations

Model	Individual 1				Individual 2			
	w_s	LFO-CV	$ELPD_{\Delta}$	se_{Δ}	w_s	LFO-CV	$ELPD_{\Delta}$	se_{Δ}
BEKK	.155	104.55	-20.67	20.44	.000	198.51	-1.61	14.68
pdBEKK	.122	117.73	-27.26	27.05	.210	208.00	-6.35	14.75
DCC	.132	69.93	-3.36	4.40	.395	205.70	-5.20	13.84
CCC	.591	63.21	[†] 0.00	—	.395	195.30	[†] 0.00	—

Note. w_s are the stacking weights. $ELPD_{\Delta}$ is the difference between the model with the smallest LFO-CV and se_{Δ} is the standard error of the difference. [†]Denotes the best model in terms of LFO-CV.

Table 2

Forecasted and Observed Realizations for Individual 1

	PC Usage					Walking Speed				
Period	101	102	103	104	105	101	102	103	104	105
Observed	1.28	0.86	0.48	0.28	0.2	-0.51	-0.7	-0.82	-0.72	-0.56
Averaged Forecasts										
M	-0.06	0.03	0.06	0.07	0.09	-0.5	-0.45	-0.42	-0.42	-0.42
SD	0.36	0.32	0.29	0.29	0.31	0.13	0.16	0.18	0.19	0.21
Med	-0.05	0.02	0.05	0.07	0.09	-0.5	-0.45	-0.42	-0.42	-0.43
2.5%	-0.75	-0.61	-0.5	-0.49	-0.52	-0.75	-0.76	-0.78	-0.78	-0.82
97.5%	0.68	0.67	0.65	0.63	0.67	-0.24	-0.14	-0.08	-0.03	-0.01
RMSE	1.39	0.9	0.52	0.36	0.32	0.13	0.3	0.44	0.36	0.25
Forecasts from Single Best Model: CCC										
M	-0.08	0.02	0.05	0.06	0.09	-0.51	-0.46	-0.44	-0.44	-0.44
SD	0.54	0.49	0.45	0.44	0.46	0.18	0.21	0.23	0.25	0.27
Med	-0.07	0.01	0.04	0.05	0.09	-0.52	-0.46	-0.44	-0.44	-0.44
2.5%	-1.14	-0.94	-0.83	-0.78	-0.77	-0.87	-0.87	-0.89	-0.93	-0.95
97.5%	0.99	0.95	0.95	0.9	0.99	-0.16	-0.04	0.01	0.07	0.12
RMSE	1.46	0.97	0.62	0.49	0.47	0.18	0.33	0.45	0.38	0.3

Note. Forecasts for Individual 1. The 2.5% and 97.5% represent 95% predictive intervals PrI. Averaged forecasts across all four parameterizations, BEKK, pdBEKK, DCC, CCC as well as forecasts from the single best model, CCC, here.

Table 3

Forecasted and Observed Realizations for Individual 2

	PC Usage					Walking Speed				
Period	101	102	103	104	105	101	102	103	104	105
Observed	1.08	0.68	0.2	0.47	0.85	-0.12	0.35	-0.07	0.12	0.05
Averaged Forecasts										
<i>M</i>	0.57	0.54	0.53	0.53	0.52	-0.06	0	0.03	0.04	0.06
<i>SD</i>	0.2	0.22	0.24	0.26	0.27	0.22	0.24	0.26	0.25	0.26
<i>Med</i>	0.57	0.55	0.52	0.53	0.52	-0.07	0.01	0.03	0.04	0.06
2.5%	0.17	0.1	0.04	0.03	0	-0.49	-0.47	-0.46	-0.45	-0.45
97.5%	0.97	0.96	1.02	1.03	1.06	0.38	0.46	0.55	0.54	0.57
RMSE	0.55	0.26	0.41	0.27	0.43	0.23	0.42	0.28	0.27	0.26
Forecasts from Single Best Model: CCC										
<i>M</i>	0.56	0.54	0.51	0.51	0.5	0	0.06	0.1	0.13	0.12
<i>SD</i>	0.35	0.34	0.37	0.4	0.42	0.37	0.42	0.47	0.43	0.44
<i>Med</i>	0.56	0.53	0.52	0.51	0.51	-0.01	0.06	0.09	0.13	0.13
2.5%	-0.12	-0.13	-0.2	-0.27	-0.34	-0.72	-0.72	-0.75	-0.7	-0.75
97.5%	1.24	1.21	1.21	1.33	1.29	0.72	0.86	1.02	1	0.97
RMSE	0.62	0.37	0.48	0.4	0.55	0.39	0.5	0.5	0.43	0.44

Note. Forecasts for Individual 2. The 2.5% and 97.5% represent 95% predictive intervals PrI. Averaged forecasts across all four parameterizations, BEKK, pdBEKK, DCC, CCC as well as forecasts from the single best model, CCC, here.

Table 4

Forecasted and Realized Variability (RV) for Individual 1

	PC Usage					Walking Speed				
Period	101	102	103	104	105	101	102	103	104	105
RV	0.33	0.1	0.05	0.11	0.09	0.01	0.01	0.01	0.01	0.01
Averaged Forecasts										
<i>M</i>	0.29	0.19	0.16	0.15	0.18	0.04	0.05	0.05	0.06	0.06
<i>SD</i>	0.05	0.07	0.07	0.07	0.07	0.01	0.02	0.03	0.04	0.05
<i>Med</i>	0.28	0.17	0.14	0.14	0.16	0.04	0.04	0.05	0.05	0.05
2.5%	0.22	0.12	0.1	0.1	0.12	0.03	0.03	0.02	0.02	0.02
97.5%	0.4	0.37	0.31	0.26	0.3	0.07	0.09	0.12	0.15	0.17
Forecasts from Single Best Model: CCC										
<i>M</i>	0.28	0.2	0.16	0.15	0.17	0.03	0.03	0.03	0.03	0.03
<i>SD</i>	0.07	0.11	0.1	0.11	0.1	0.01	0.02	0.03	0.04	0.06
<i>Med</i>	0.27	0.17	0.14	0.14	0.16	0.02	0.02	0.02	0.02	0.02
2.5%	0.18	0.1	0.08	0.08	0.09	0.01	0.01	0.01	0.01	0.01
97.5%	0.46	0.5	0.41	0.32	0.36	0.06	0.08	0.1	0.11	0.13

Note. Forecasted variability for Individual 1. RV is the observed realized variability for each forecasted time period and serves as reference for the forecasts. The 2.5% and 97.5% represent 95% predictive intervals PrI.

Table 5

Forecasted and Realized Variability (RV) for Individual 2

	PC Usage					Walking Speed				
Period	101	102	103	104	105	101	102	103	104	105
RV	0.16	0.05	0.02	0.05	0.1	0.05	0.04	0.05	0.05	0.02
Averaged Forecasts										
<i>M</i>	0.08	0.08	0.09	0.09	0.09	0.09	0.09	0.11	0.1	0.1
<i>SD</i>	0.01	0.03	0.04	0.1	0.1	0.02	0.02	0.03	0.03	0.04
<i>Med</i>	0.07	0.07	0.08	0.08	0.08	0.09	0.09	0.1	0.1	0.1
2.5%	0.06	0.05	0.05	0.05	0.05	0.07	0.07	0.07	0.07	0.07
97.5%	0.11	0.14	0.18	0.17	0.19	0.13	0.14	0.17	0.16	0.16
Forecasts from Single Best Model: CCC										
<i>M</i>	0.08	0.07	0.08	0.08	0.08	0.09	0.1	0.12	0.11	0.1
<i>SD</i>	0.03	0.05	0.11	0.16	0.07	0.03	0.09	0.08	0.09	0.09
<i>Med</i>	0.08	0.06	0.06	0.07	0.07	0.09	0.09	0.1	0.1	0.09
2.5%	0.04	0.03	0.03	0.03	0.03	0.05	0.05	0.05	0.05	0.05
97.5%	0.15	0.15	0.21	0.21	0.21	0.15	0.27	0.26	0.24	0.22

Note. Forecasted variability for Individual 2. RV is the observed realized variability for each forecasted time period and serves as reference for the forecasts. The 2.5% and 97.5% represent 95% predictive intervals PrI.

Table 6

Forecasted and Realized Covariance (RCV) among Walking Speed and PC Usage for Individual 1

Realized Covariance					
Period	101	102	103	104	105
RCV	-0.032	0.004	-0.006	-0.011	0.004
Averaged Forecasts					
<i>M</i>	0.031	0.021	0.018	0.018	0.021
<i>SD</i>	0.009	0.009	0.009	0.010	0.014
<i>Med</i>	0.031	0.020	0.017	0.017	0.020
2.5%	0.014	0.006	0.005	0.004	0.006
97.5%	0.052	0.040	0.037	0.039	0.044
Forecasts from Single Best Model: CCC					
<i>M</i>	0.016	0.014	0.013	0.013	0.014
<i>SD</i>	0.010	0.010	0.010	0.010	0.011
<i>Med</i>	0.015	0.012	0.011	0.011	0.012
2.5%	-0.001	-0.001	-0.001	-0.001	-0.001
97.5%	0.038	0.037	0.036	0.036	0.037

Note. Forecasts of the covariance among PC Usage and Walking Speed. RCV is the realized covariance for the given time-period.

Table 7

Forecasted and Realized Covariance (RCV) among Walking Speed and PC Usage for Individual 2

Realized Covariance					
Period	101	102	103	104	105
RCV	0.036	0.013	-0.030	-0.017	0.002
Averaged Forecasts					
<i>M</i>	-0.004	-0.003	-0.001	-0.000	-0.001
<i>SD</i>	0.008	0.011	0.015	0.015	0.021
<i>Med</i>	-0.003	-0.002	-0.001	0.000	-0.000
2.5%	-0.021	-0.028	-0.030	-0.031	-0.030
97.5%	0.010	0.017	0.024	0.026	0.025
Forecasts from Single Best Model: CCC					
<i>M</i>	-0.004	-0.000	-0.001	-0.001	-0.000
<i>SD</i>	0.018	0.045	0.055	0.041	0.037
<i>Med</i>	-0.002	-0.000	0.001	0.000	0.001
2.5%	-0.045	-0.045	-0.048	-0.051	-0.053
97.5%	0.028	0.045	0.051	0.046	0.041

Note. Forecasts of the covariance among PC Usage and Walking Speed. RCV is the realized covariance for the given time-period.

Table 8

Kullback-Leibler Divergence

Period	101	102	103	104	105
Kullback Leibler Divergence and Ratio					
Individual 1:					
KL_w	27.602	97.399	107.349	29.619	40.600
KL_{CCC}	50.005	192.998	226.218	67.456	98.130
Ratio $^{KL_{CCC}/KL_w}$	1.812	1.982	2.107	2.277	2.417
Individual 2:					
KL_w	8.929	11.515	15.112	10.518	10.960
KL_{CCC}	8.774	14.154	19.023	12.054	13.388
Ratio $^{KL_{CCC}/KL_w}$	0.983	1.229	1.259	1.146	1.221

Note. KL_w is the Kullback-Leibler divergence for the model averaged conditional covariance matrix $\hat{\mathbf{H}}_t$, and KL_{CCC} is divergence for the single best model, the CCC parameterization. The Ratio $^{KL_{CCC}/KL_w}$ is an indication on how much larger the KL_{CCC} is with respect to KL_w .

Figure 1. Time series of walking speed (blue) and computer use (red), averaged over two weeks for both individuals. Both variables are standardized within individual. The time unit is two weeks. The first 100 observations are used as training set, the remaining 5 observations are used to evaluate the forecasts. The vertical line separates these two data sets.

Figure 2. Estimated and forecasted z -scores for both individuals and both variables, Walking Speed and Computer Use. The red lines and bands represent mean estimates with the according PrI's for the training data. The turquoise line and band represents the averaged forecasts with their respective 95% PrI's. Black dots are the actual observed values.

Figure 3. Model averaged forecasted variability for five time periods and both variables in Individual 1 and 2. The red lines and bands represent mean estimates with the according PrI's for the training data. The turquoise line and band represents the averaged forecasts with their respective 95% PrI's. Black dots represent the RV values.

Figure 4. Model averaged forecasts of the conditional covariance among Walking Speed and Computer Usage. The red lines and bands represent mean estimates with the according PrI's for the training data. The turquoise line and band represents the averaged forecasts with their respective 95% PrI's. Black dots represent the realized covariance values.

Figure 5. Kullback-Leibler divergence for each forecasted time-point. The solid lines represent the means while the bands capture 95% of the credible region.