

A new combination approach to reducing forecast errors with an application to volatility forecasting

Till Weigt und Bernd Wilfling[†]

46/2016

[†] Department of Economics, University of Münster, Germany

A new combination approach to reducing forecast errors with an application to volatility forecasting

TILL WEIGT, BERND WILFLING

*Westfälische Wilhelms-Universität Münster, Department of Economics
(CeNoS and CQE), Am Stadtgraben 9, 48143 Münster, Germany**

(Date of this version: April 6, 2016)

Abstract

This paper formally establishes a new forecast combination approach, which is based on VAR modeling of the forecast errors resulting from alternative forecast models. We apply our approach to volatility forecasting by combining several structural time series models with implied volatility. Using a multi-currency data set, we conduct in-sample and out-of-sample forecasting analyses in order (a) to demonstrate the statistical significance of our approach, and (b) to assess its forecasting superiority over alternative forecasting models and combinations.

Keywords: Forecast combination, volatility forecasting, realized volatility, implied volatility, exchange rates

JEL classification: C53, G17

**E-mail addresses:* till.weigt@wiwi.uni-muenster.de (T. Weigt),
bernd.wilfling@wiwi.uni-muenster.de (B. Wilfling).

1 Introduction

Since the seminal paper by Bates and Granger (1969), the combination of alternative forecast models has become an important and active field of research, but with many issues remaining unresolved. In particular, problems arising from substantial estimation errors and the presence of structural breaks have encouraged many authors to formulate a plethora of alternative forecast combination schemes (see Timmermann, 2006, for an in-depth survey).

In this paper, we propose a new and intuitive forecast combination approach, which interrelates the forecast errors of alternative forecast models within a Vector Autoregressive (VAR) framework. In contrast to the existing forecast combination schemes, our approach does not merge alternative multiple forecasts into a single newly combined forecast. Instead, our methodology attempts to reduce the forecast errors of each single forecast model involved, by exploiting the interrelationships between the forecast errors of all forecast models under consideration.

In an empirical study, we apply our new approach to volatility forecasting, by explicitly interrelating the errors from realized and implied volatility forecasts, an important issue in the valuation of financial and derivative contracts (Andersen et al., 2003; Atak and Kapetanios, 2013). Using a multi-currency data set, we show that our forecast combination scheme performs convincingly with respect to in-sample, and, most notably, to out-of-sample forecast comparisons with alternative forecasting procedures.

Section 2 formally establishes our new forecasting procedure, while Section 3 shows how to apply it to forecasting realized volatility, using implied volatility measurements. Section 4 presents the empirical results and Section 5 concludes.

2 Vector Autoregressive Forecast Error Combination (VAFEC) approach

For $t = 0, \pm 1, \pm 2, \dots$ we consider the univariate time-series variable y_t and for $h > 0$, we denote a forecast of y_{t+h} , based on information available at date t , by $\hat{y}_{t+h|t}$. We assume given M alternative forecast models and, associated with each model, the corresponding forecasts $\hat{y}_{t+h|t,1}, \dots, \hat{y}_{t+h|t,M}$, which we collect in the $(M \times 1)$ vector

$\hat{\mathbf{y}}_{t+h|t} = (\hat{y}_{t+h|t,1}, \dots, \hat{y}_{t+h|t,M})'$. The information set at date t , $\mathcal{F}_t$, consists of (a) these M forecasts, (b) the entire history of these M forecasts, and (c) the entire history of our time-series variable, i.e. $\mathcal{F}_t = \{\dots, \hat{\mathbf{y}}_{t-1+h|t-1}, \hat{\mathbf{y}}_{t+h|t}, \dots, y_{t-1}, y_t\}$. Furthermore, we collect the forecast errors associated with our forecast models, $e_{t+h|t,i} = y_{t+h} - \hat{y}_{t+h|t,i}$ for $i = 1, \dots, M$, in the $(M \times 1)$ vector $\mathbf{e}_{t+h|t} = (e_{t+h|t,1}, \dots, e_{t+h|t,M})'$.

In order to contrast our new approach with the existing literature, we briefly review the general parametric forecast combination approach (see, *inter alia*, Timmermann, 2006), which essentially consists of mapping the M forecasts collected in $\hat{\mathbf{y}}_{t+h|t}$ to the single new forecast $\hat{y}_{t+h|t}^{\text{comb}}$. Denoting the $(M \times 1)$ vector of combination weights by $\boldsymbol{\omega}_{t+h|t} = (\omega_{t+h|t,1}, \dots, \omega_{t+h|t,M})'$, we may represent the forecast combination itself and the corresponding forecast error via the real-valued vector function $g : \mathbb{R}^M \rightarrow \mathbb{R}$ as

$$\hat{y}_{t+h|t}^{\text{comb}} = g(\hat{\mathbf{y}}_{t+h|t}; \boldsymbol{\omega}_{t+h|t}) \quad \text{and} \quad e_{t+h|t}^{\text{comb}} = y_{t+h} - \hat{y}_{t+h|t}^{\text{comb}}. \quad (1)$$

The optimal weights $\boldsymbol{\omega}_{t+h|t}^* = (\omega_{t+h|t,1}^*, \dots, \omega_{t+h|t,M}^*)'$ of the forecast combination, subject to the mean-squared-error (MSE) loss function, are given by the solution to the minimization problem

$$\boldsymbol{\omega}_{t+h|t}^* = \arg \min_{\boldsymbol{\omega}_{t+h|t}} \mathbb{E} \left\{ [e_{t+h|t}^{\text{comb}}(\boldsymbol{\omega}_{t+h|t})]^2 | \mathcal{F}_t \right\}$$

with $\mathbb{E}(\cdot|\cdot)$ denoting the conditional expectation operator. Analytically closed-form formulae for the optimal weights $\boldsymbol{\omega}_{t+h|t}^*$, subject to the MSE loss function, are available in the literature and typically also include an additional constant weight $\omega_{t+h|t,0}$ in the above minimization problem, in order to account for potentially biased forecast errors (Timmermann, 2006).

Our new forecasting approach rests on the assumption that the forecast error vector $\mathbf{e}_{t+h|t}$ is governed by a covariance-stationary process. Following Lütkepohl (2005), we can thus model the dynamics of the forecast errors as the VAR(p) process

$$\mathbf{e}_{t+h|t} = \boldsymbol{\nu} + \mathbf{A}_1 \mathbf{e}_{t+h-1|t-1} + \dots + \mathbf{A}_p \mathbf{e}_{t+h-p|t-p} + \tilde{\mathbf{e}}_{t+h|t}, \quad (2)$$

where $\boldsymbol{\nu} = (\nu_1, \dots, \nu_M)'$ represents an $(M \times 1)$ vector of constants, $\mathbf{A}_1, \dots, \mathbf{A}_p$ denote $(M \times M)$ parameter matrices and $\tilde{\mathbf{e}}_{t+h|t}$ represents an $(M \times 1)$ independent white noise

process. Denoting the i th row vector of the matrix $\mathbf{A}_k$ (for $k = 1, \dots, p$) by $\mathbf{A}_{k,i}$, we obtain the optimal weights (parameter estimates) from Eq. (2), under the MSE loss function, as the solution to the M simultaneous minimization problems

$$(\boldsymbol{\nu}^*, \mathbf{A}_1^*, \dots, \mathbf{A}_p^*) = \arg \min_{\boldsymbol{\nu}, \mathbf{A}_1, \dots, \mathbf{A}_p} \mathbb{E} \left\{ [\tilde{e}_{t+h|t,i}(\nu_i, \mathbf{A}_{1,i}, \dots, \mathbf{A}_{p,i})]^2 | \mathcal{F}_t \right\} \quad \text{for } i = 1, \dots, M. \quad (3)$$

Since Eq. (2) represents a standard VAR(p) model, the optimal weights computed from the minimization problems in Eq. (3) coincide with the multivariate least squares estimator, closed-form expressions of which are given, *inter alia*, in Lütkepohl (2006, pp. 69-72). Owing to our VAR(p) modeling of the forecast errors, we refer to our forecasting approach as the Vector Autoregressive Forecast Error Combination of order p [in symbols: VAFEC(p)]. In view of Eq. (3), it is clear that the VAFEC approach minimizes the forecast error variances of all M forecast models under consideration:

$$\mathbb{E} \left\{ [\tilde{e}_{t+h|t,i}(\nu_i^*, \mathbf{A}_{1,i}^*, \dots, \mathbf{A}_{p,i}^*)]^2 | \mathcal{F}_t \right\} \leq \mathbb{E} \left\{ e_{t+h|t,i}^2 | \mathcal{F}_t \right\} \quad \text{for } i = 1, \dots, M. \quad (4)$$

Finally, we establish an explicit forecasting formula within our VAFEC framework. Our forecasting formula arises from adding the lagged VAFEC errors obtained from the minimization problem (3) to the initial M forecasts collected in the vector $\hat{\mathbf{y}}_{t+h|t}$. To present a closed-form expression, we denote the $(M \times M)$ identity matrix by $\mathbf{I}_M$ and the $(M \times M)$ matrix consisting of zeros by $\mathbf{0}_M$. Furthermore, we define the $(Mp \times 1)$ vectors

$$\mathbf{E}_{t+h|t} \equiv \begin{bmatrix} \mathbf{e}_{t+h|t} \\ \vdots \\ \mathbf{e}_{t+h-p+1|t-p+1} \end{bmatrix}, \quad \tilde{\mathbf{E}}_{t+h|t} \equiv \begin{bmatrix} \tilde{\mathbf{e}}_{t+h|t} \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \mathbf{N} \equiv \begin{bmatrix} \boldsymbol{\nu} \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

the $(Mp \times Mp)$ matrix

$$\mathbf{A} \equiv \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \dots & \mathbf{A}_{p-1} & \mathbf{A}_p \\ \mathbf{I}_M & \mathbf{0}_M & \dots & \mathbf{0}_M & \mathbf{0}_M \\ \mathbf{0}_M & \mathbf{I}_M & \dots & \mathbf{0}_M & \mathbf{0}_M \\ \vdots & & \ddots & \vdots & \vdots \\ \mathbf{0}_M & \mathbf{0}_M & \dots & \mathbf{I}_M & \mathbf{0}_M \end{bmatrix},$$

and rewrite the VAR(p) process from Eq. (2) in VAR(1) form:

$$\mathbf{E}_{t+h|t} = \mathbf{N} + \mathbf{A}\mathbf{E}_{t+h-1|t-1} + \tilde{\mathbf{E}}_{t+h|t}. \quad (5)$$

By recursive substitution, we obtain the conditional expectation

$$\mathbb{E} \{ \mathbf{E}_{t+h|t} | \mathcal{F}_t \} = [\mathbf{I}_{Mp} + \mathbf{A} + \mathbf{A}^2 + \dots + \mathbf{A}^{h-1}] \mathbf{N} + \mathbf{A}^h \mathbf{E}_{t|t-h}, \quad (6)$$

which is the minimum MSE predictor of the forecast-error process as specified in Eqs. (2) and (5). As in the vector $\mathbf{E}_{t+h|t}$, we collect the p initial ($M \times 1$) forecast vectors $\hat{\mathbf{y}}_{t+h|t}, \dots, \hat{\mathbf{y}}_{t+h-p+1|t-p+1}$ in the $(Mp \times 1)$ vector $\hat{\mathbf{Y}}_{t+h|t}$, and, analogously, denote the corresponding $(Mp \times 1)$ vector of our new forecasts derived from the VAFEC approach by $\tilde{\mathbf{Y}}_{t+h|t}$, that is

$$\hat{\mathbf{Y}}_{t+h|t} \equiv \begin{bmatrix} \hat{\mathbf{y}}_{t+h|t} \\ \vdots \\ \hat{\mathbf{y}}_{t+h-p+1|t-p+1} \end{bmatrix} \quad \text{and} \quad \tilde{\mathbf{Y}}_{t+h|t} \equiv \begin{bmatrix} \tilde{\mathbf{y}}_{t+h|t} \\ \vdots \\ \tilde{\mathbf{y}}_{t+h-p+1|t-p+1} \end{bmatrix}.$$

Then, using the minimum MSE predictor of the forecast-error process from Eq. (6), we express our new forecasts as

$$\begin{aligned} \tilde{\mathbf{Y}}_{t+h|t} &= \hat{\mathbf{Y}}_{t+h|t} + \mathbb{E} \{ \mathbf{E}_{t+h|t} | \mathcal{F}_t \} \\ &= \hat{\mathbf{Y}}_{t+h|t} + [\mathbf{I}_{Mp} + \mathbf{A} + \mathbf{A}^2 + \dots + \mathbf{A}^{h-1}] \mathbf{N} + \mathbf{A}^h \mathbf{E}_{t|t-h}. \end{aligned} \quad (7)$$

In order to extract the relevant M new VAFEC forecasts $\tilde{\mathbf{y}}_{t+h|t} = (\tilde{y}_{t+h|t,1}, \dots, \tilde{y}_{t+h|t,M})'$ from $\tilde{\mathbf{Y}}_{t+h|t}$, we premultiply the $(M \times Mp)$ matrix $\mathbf{J} = [\mathbf{I}_M \quad \mathbf{0}_M \quad \dots \quad \mathbf{0}_M]$ to both sides of Eq. (7):

$$\begin{aligned} \tilde{\mathbf{y}}_{t+h|t} &= \mathbf{J} \left(\hat{\mathbf{Y}}_{t+h|t} + [\mathbf{I}_{Mp} + \mathbf{A} + \mathbf{A}^2 + \dots + \mathbf{A}^{h-1}] \mathbf{N} + \mathbf{A}^h \mathbf{E}_{t|t-h} \right) \\ &= \hat{\mathbf{y}}_{t+h|t} + \mathbf{J} [\mathbf{I}_{Mp} + \mathbf{A} + \mathbf{A}^2 + \dots + \mathbf{A}^{h-1}] \mathbf{N} + \mathbf{J} \mathbf{A}^h \mathbf{E}_{t|t-h}. \end{aligned} \quad (8)$$

Replacing the theoretical parameter vector $\boldsymbol{\nu}$ and the matrix $\mathbf{A}$ in Eq. (8) with their multivariate least squares estimates $\boldsymbol{\nu}^*$ and $\mathbf{A}^*$, we obtain our explicit VAFEC forecasts.

3 VAFEC volatility forecasting using implied volatility

In order to demonstrate how to apply our VAFEC approach formally to volatility forecasting using implied volatility, we consider the following purely illustrative example. Let y_t constitute the realized volatility of a financial return variable x_t at date t , where realized volatility is defined as the square-root of the sum of n (equidistantly observed) squared intraday returns $x_{t:i}$ ($i = 1, \dots, n$):

$$y_t = \sqrt{\sum_{i=1}^n x_{t:i}^2}.$$

In what follows, we use realized volatility as a proxy for the true latent volatility σ_t , and, from now on, frequently use the phrases 'realized volatility' and 'true volatility' interchangeably.

Let $\hat{y}_{t+1|t} \equiv \hat{y}_{t+1|t,1}$ denote the 1-step realized-volatility forecast from any arbitrary structural time-series model (e.g. from a GARCH model) for date $t + 1$, using only information up to date t . Next, consider the implied volatility of the financial return defined as that value of return volatility, for which the theoretical return-option price equals the observed market price of the option. Let IV_t denote the implied volatility measured at date t and regard implied volatility as our second realized-volatility forecast model, that is $IV_t \equiv \hat{y}_{t+1|t,2}$. In line with Eq. (2), the VAFEC(1) specification under this setting is formally given by

$$\begin{pmatrix} y_{t+1} - \hat{y}_{t+1|t} \\ y_{t+1} - IV_t \end{pmatrix} = \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} + \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} y_t - \hat{y}_{t|t-1} \\ y_t - IV_{t-1} \end{pmatrix} + \begin{pmatrix} \tilde{e}_{t+1|t,1} \\ \tilde{e}_{t+1|t,2} \end{pmatrix},$$

where a_{ij} ($i, j = 1, 2$) denote the parameters of the VAR(1) coefficient matrix $\mathbf{A} \equiv \mathbf{A}_1$.

However, since we may want to exploit all available information at date $t + 1$, it appears straightforward to shift the values of implied volatility one period ahead, so as to yield the slightly more informative specification

$$\begin{pmatrix} y_{t+1} - \hat{y}_{t+1|t} \\ y_{t+1} - IV_{t+1} \end{pmatrix} = \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} + \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} y_t - \hat{y}_{t|t-1} \\ y_t - IV_t \end{pmatrix} + \begin{pmatrix} \tilde{e}_{t+1|t,1} \\ \tilde{e}_{t+1|t,2} \end{pmatrix}. \quad (9)$$

Our higher-order VAFEC specifications in Section 4 make systematic use of this time shift in the implied volatility measurements.

In our subsequent empirical application, we separately combine three distinct structural time-series models with implied volatility via our VAFEC approach, in order to forecast realized volatility. The first model is the standard GARCH(1,1) process (Bollerslev, 1986) that specifies the financial return x_t as

$$x_t = \sigma_t \cdot \epsilon_t, \quad (10)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \cdot x_{t-1}^2 + \beta_1 \cdot \sigma_{t-1}^2, \quad (11)$$

where ϵ_t constitutes an i.i.d. process of standard normal variables [$\epsilon_t \sim \text{NID}(0, 1)$] and with parameters $\alpha_0 > 0, \alpha_1, \beta_1 \geq 0$. Assuming $\alpha_1 + \beta_1 < 1$ (wide-sense stationarity), we use the square-root of the 1-step-ahead conditional-variance forecasts $\hat{\sigma}_{t+1|t}^2$ from the estimated GARCH(1,1) models as the 1-step-ahead forecasts of realized volatility (that is we set $\hat{y}_{t+1|t, \text{GARCH}} \equiv \hat{\sigma}_{t+1|t}$).

In contrast to the GARCH(1,1) process, the other two time-series models used in Section 4 directly specify the dynamics of the realized volatility process y_t . The Heterogeneous-Auto-Regressive (HAR) model (Corsi, 2009) specifies realized volatility as

$$y_t = \alpha_0 + \alpha_1 \cdot y_{t-1} + \alpha_2 \cdot y_{t-1}^w + \alpha_3 \cdot y_{t-1}^m + \epsilon_t, \quad (12)$$

where

$$y_{t-1}^w \equiv \frac{1}{5} \sum_{i=1}^5 y_{t-i}, \quad y_{t-1}^m \equiv \frac{1}{22} \sum_{i=1}^{22} y_{t-i},$$

$\epsilon_t \sim \text{NID}(0, \sigma^2)$ and with parameters $\alpha_0, \dots, \alpha_3$. The (AutoRegressive, Fractionally Integrated, MovingAverage) ARFIMA(1, d , 1) model (see Granger and Joyeux, 1980; Hosking, 1981) uses the lag operator L and the unconditional expectation $\mathbb{E}(\cdot)$ to write the realized volatility process y_t as

$$(1 - \phi L)(1 - L)^d [y_t - \mathbb{E}(y_t)] = (1 + \theta L)\epsilon_t, \quad (13)$$

where $\epsilon_t \sim \text{NID}(0, \sigma^2)$ and with parameters d, ϕ, θ restricted by $0 < d < 0.5$ and $|\phi|, |\theta| < 1$ (stationarity and invertibility conditions). Via the conditional expectations

of the right-hand sides of the Eqs. (12) and (13), it is straightforward to establish the optimal 1-step ahead forecasts $\hat{y}_{t+1|t,\text{HAR}}$ and $\hat{y}_{t+1|t,\text{ARFIMA}}$ subject to the MSE loss function (see, *inter alia*, Corsi 2009; Doornik and Ooms, 2004).

4 Empirical application

4.1 Data

We analyze the following 8 currencies vis-à-vis the US dollar between January 1, 2004 and December 31, 2011: the euro (EUR), the British pound (GBP), the Japanese yen (JPY), the Canadian (CAD) and the New Zealand dollar (NZD), the Swiss franc (CHF), and the Norwegian (NOK) and the Swedish kroner (SEK). All exchange rates, provided by the *Dukascopy Bank SA*, were recorded at a 5-minute frequency (288 observations per 24-hours trading day). For each exchange rate, we use daily at-the-money implied volatilities with one week to maturity, as provided by *Thompson Reuters Datastream*. All econometric procedures were implemented within the software package R, using the sub-packages 'rugarch' and 'MTS'.

4.2 In-sample analysis

We demonstrate that the forecast errors from the models embedded in our VAFEC combination approach interrelate significantly with each other. We focus on the US-dollar/euro exchange rate.¹

Our first learning sample covers the period between January 1, 2004 and December 31, 2007, from which we initially estimated the parameters for the GARCH, HAR and ARFIMA models. We implemented a rolling time window between January 1, 2008 and December 31, 2011, from which we (a) sequentially updated the parameter estimates, and (b) computed the 1-step-ahead forecasts $\hat{y}_{t+1|t,\text{GARCH}}$, $\hat{y}_{t+1|t,\text{HAR}}$, $\hat{y}_{t+1|t,\text{ARFIMA}}$ and the forecast errors $e_{t+1|t,\text{GARCH}}$, $e_{t+1|t,\text{HAR}}$, $e_{t+1|t,\text{ARFIMA}}$. We then combined each separate forecast-error series with the (time-shifted) implied volatility (IV) measurement errors and, in line with Eq. (2), fitted VAR(3) specifications to each of these three

¹The in-sample results for the other 7 currencies are qualitatively similar and available upon request.

combinations, using the data between January 1, 2008 and December 31, 2011.²

Table 1 about here

Table 1 displays the VAR(3) estimation results for the three respective combinations VAFEC(3)-GARCH-IV, VAFEC(3)-HAR-IV, VAFEC(3)-ARFIMA-IV. The relevant coefficients governing the interrelationships between the forecast errors and the IV measurement errors are $\mathbf{A}_{1,12}$, $\mathbf{A}_{1,21}$, $\mathbf{A}_{2,12}$, $\mathbf{A}_{2,21}$, $\mathbf{A}_{3,12}$, $\mathbf{A}_{3,21}$. Except for the three $\mathbf{A}_{2,12}$ coefficients, all interrelating parameters are significant at least at the 5% level across all combinations. This essentially supports our VAFEC approach.

4.3 Out-of-sample analysis

We computed our VAFEC(3) out-of-sample forecasts with a 2-step procedure. (1) We generated the 1-step-ahead forecasts and forecast errors $\hat{y}_{t+1|t,\text{GARCH}}$, $\hat{y}_{t+1|t,\text{HAR}}$, $\hat{y}_{t+1|t,\text{ARFIMA}}$, $e_{t+1|t,\text{GARCH}}$, $e_{t+1|t,\text{HAR}}$, $e_{t+1|t,\text{ARFIMA}}$, as in our in-sample analysis. (2) We generated the 1-step-ahead VAFEC(3)-GARCH-IV, VAFEC(3)-HAR-IV, VAFEC(3)-ARFIMA-IV forecasts by (a) estimating the VAR(3) coefficients over the period between January 1, 2008 and December 31, 2009, and (b) by sequentially updating the VAR(3) coefficients during the period between January 1, 2010 and December 31, 2011. Using these updated estimation results from our rolling window, we generated our ultimate out-of-sample VAFEC(3)-GARCH-IV, VAFEC(3)-HAR-IV, VAFEC(3)-ARFIMA-IV forecasts according to Eq. (7).

Table 2 about here

We assess the forecasting performance of our models and model combinations in terms of the (sampling) MSEs. Table 2 displays the MSEs of the out-of-sample forecast from the GARCH, HAR, ARFIMA models and the implied volatility measurements (IV). While the HAR forecasts perform best for 7 out of 8 currencies, the GARCH forecasts perform unambiguously worst for all currencies. The ARFIMA and the IV

²We used the standard criteria to select our lag length $p = 3$.

forecasts take intermediate ranks, with the ARFIMA forecasts slightly outperforming the IV measurements for 4 out of 7 currencies.

Table 3 about here

Table 3 displays the out-of-sample forecasting performance of the three combinations VAFEC(3)-GARCH-IV, VAFEC(3)-HAR-IV, VAFEC(3)-ARFIMA-IV. All MSEs in Table 3 are expressed relative to the MSEs of the forecasts from the corresponding (uncombined) time series models.³ The MSEs of the three VAFEC(3)-combination forecasts are all less than 1. Thus, for all currencies, the forecasts from each VAFEC(3) combination outperform the respective forecasts from the (uncombined) GARCH, HAR, ARFIMA models. The direct comparison of the MSEs from the VAFEC(3) combinations with the MSEs from column 'IV' in Table 3 also reveals the unambiguous superiority of the VAFEC(3)-combination forecasts over the implied volatility forecasts.

Finally, we compare the out-of-sample forecasting performance of our VAFEC(3) combinations with an alternative forecast combination method, namely with the (equally weighted) averages between the (uncombined) time-series and the implied volatility forecasts (denoted by $\frac{1}{2} \cdot (\hat{y}_{t+1|t, \text{GARCH}} + \text{IV}_t)$, and so forth in Table 3). Our VAFEC(3)-GARCH-IV and VAFEC(3)-ARFIMA-IV combination forecasts clearly outperform the corresponding forecast averages for all currencies. Except for the British pound (GBP) and the Norwegian kroner (NOK), our VAFEC(3)-HAR-IV combination forecasts also outperform the HAR-IV forecast averages. Overall, our VAFEC(3) combination forecasts outperform the forecast average combinations in 22 out of the 24 cases considered in this multi-currency analysis.

5 Concluding remarks

This paper formally establishes a new forecast combination approach that reduces the forecast errors of any single involved forecast model. This is accomplished by VAR

³More precisely, in the upper (middle/lower) block, all MSEs were divided by the currency-specific MSEs from the (uncombined) GARCH (HAR/ARFIMA) forecasts.

modeling of the dynamic interrelationships between all model-specific forecast errors. The theoretical properties of the VAR process ensure that the (theoretical) MSEs of our improved forecast errors cannot exceed the MSEs of the original forecast error processes [see Eq. (4)].

In a multi-currency analysis, we apply our new combination approach to forecasting realized volatility, using implied volatility measurements and three alternative time-series forecasting models. The in-sample and out-of-sample results of our new approach unambiguously reflect (a) the empirical relevance of our new approach, and (b) its forecasting superiority over the uncombined and an average-combination forecast.

In line with the results presented in Pong et al. (2004), our empirical study focuses on the analysis of (short-term) 1-step-ahead forecasting performance. A systematic and extensive investigation of multiple-step-ahead volatility forecasting performance should be considered in future research. Overall, we believe that our new forecast combination approach offers a promising complement to the existing forecasting literature.

References

- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* 71, 579-625.
- Atak, A., Kapetanios, G., 2013. A factor approach to realized volatility forecasting in the presence of finite jumps and cross-sectional correlation in pricing errors. *Economic Letters* 120, 24-28.
- Bates, J.M., Granger, C.W.J., 1969. The combination of forecasts. *Operational Research Quarterly* 20, 451-468.
- Bollerslev, T., 1896. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307-327.
- Corsi, F., 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7, 174-196.
- Doornik, J.A., Ooms, M., 2004. Inference and forecasting for ARFIMA models with an application to US and UK inflation. *Studies in Nonlinear Dynamics & Econometrics* 8(2), Article 14, 1-23. *Statistics* 13, 253-263.

- Granger, C.W.J., Joyeux, R., 1980. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 15-39.
- Hosking, J.R.M., 1981. Fractional differencing. *Biometrika* 68, 165-176.
- Lütkepohl, H., 2005. *N Introduction to Multiple Time Series Analysis*. Springer-Verlag Berlin Heidelberg, Germany.
- Pong, S., Shackleton, M.B., Taylor, S.J., Xu, X., 2004. Forecasting currency volatility: A comparison of implied volatilities and AR(FI)MA models. *Journal of Banking & Finance* 28, 2541-2563.
- Timmermann, A., 2006. Forecast combinations. *Handbook of Economic Forecasting* 1, 135-196.

Tables

Table 1
VAR(3) parameter estimates from the VAFEC(3) combinations
for the US-dollar/euro exchange rate

	ν_1	$\mathbf{A}_{1,11}$	$\mathbf{A}_{1,12}$	$\mathbf{A}_{2,11}$	$\mathbf{A}_{2,12}$	$\mathbf{A}_{3,11}$	$\mathbf{A}_{3,12}$
	ν_2	$\mathbf{A}_{1,21}$	$\mathbf{A}_{1,22}$	$\mathbf{A}_{2,21}$	$\mathbf{A}_{2,22}$	$\mathbf{A}_{3,21}$	$\mathbf{A}_{3,22}$
VAFEC(3)-GARCH-IV combination:							
GARCH	−0.022*** (0.006)	1.473*** (0.112)	0.613*** (0.105)	−0.236 (0.148)	−0.175 (0.139)	−0.194* (0.114)	−0.326*** (0.107)
IV	−0.026*** (0.006)	−1.124*** (0.119)	−0.257** (0.111)	0.362** (0.153)	0.313** (0.144)	0.248** (0.123)	0.379** (0.115)
VAFEC(3)-HAR-IV combination:							
HAR	−0.015*** (0.005)	1.087*** (0.105)	0.512*** (0.097)	0.023 (0.116)	0.034 (0.107)	−0.232*** (0.068)	−0.298*** (0.063)
IV	−0.021*** (0.005)	−1.142*** (0.113)	−0.157 (0.104)	0.478*** (0.150)	0.303** (0.138)	0.277*** (0.100)	0.365*** (0.092)
VAFEC(3)-ARFIMA-IV combination:							
ARFIMA	−0.013** (0.006)	1.440*** (0.116)	0.600*** (0.107)	−0.231 (0.151)	−0.073 (0.139)	−0.340*** (0.111)	−0.444*** (0.103)
IV	−0.018*** (0.005)	−1.227*** (0.124)	−0.255** (0.115)	0.472*** (0.167)	0.287* (0.155)	0.414*** (0.129)	0.530*** (0.120)

Note: Standard errors are in parentheses. $\mathbf{A}_{k,ij}$ (for $k = 1, \dots, 3$ and $i, j = 1, 2$) denotes the parameters of the VAR(3) coefficient matrix $\mathbf{A}_k$ from Eq. (2). *, ** and *** denote significance at 10, 5, and 1% levels, respectively.

Table 2

MSEs of the out-of-sample forecasts from the (uncombined) time series models

	GARCH	HAR	ARFIMA	IV
EUR	0.025	0.016	0.020	0.020
GBP	0.026	0.015	0.015	0.013
JPY	0.068	0.052	0.053	0.054
CAD	0.036	0.021	0.024	0.022
CHF	0.061	0.035	0.039	0.038
NOK	0.056	0.029	0.034	0.037
NZD	0.074	0.042	0.042	0.050
SEK	0.050	0.027	0.032	0.033

Table 3

MSE comparisons between out-of-sample forecast errors

MSE comparisons relative to the (uncombined) GARCH model:				
	GARCH	IV	VAFEC(3)-GARCH-IV	$\frac{1}{2} \cdot (\hat{y}_{t+1 t, \text{GARCH}} + \text{IV}_t)$
EUR	1.000	0.796	0.516	0.818
GBP	1.000	0.491	0.482	0.617
JPY	1.000	0.784	0.689	0.826
CAD	1.000	0.617	0.507	0.755
CHF	1.000	0.618	0.549	0.747
NOK	1.000	0.650	0.488	0.770
NZD	1.000	0.671	0.532	0.775
SEK	1.000	0.657	0.455	0.772
MSE comparisons relative to the (uncombined) HAR model:				
	HAR	IV	VAFEC(3)-HAR-IV	$\frac{1}{2} \cdot (\hat{y}_{t+1 t, \text{HAR}} + \text{IV}_t)$
EUR	1.000	1.252	0.853	1.006
GBP	1.000	0.884	0.857	0.857
JPY	1.000	1.027	0.923	0.947
CAD	1.000	1.038	0.857	0.916
CHF	1.000	1.068	0.942	0.963
NOK	1.000	1.259	0.973	0.971
NZD	1.000	1.183	0.949	0.967
SEK	1.000	1.212	0.879	0.969
MSE comparisons relative to the (uncombined) ARFIMA model:				
	ARFIMA	IV	VAFEC(3)-ARFIMA-IV	$\frac{1}{2} \cdot (\hat{y}_{t+1 t, \text{ARFIMA}} + \text{IV}_t)$
EUR	1.000	1.016	0.686	0.943
GBP	1.000	0.861	0.847	0.854
JPY	1.000	1.018	0.917	0.972
CAD	1.000	0.919	0.743	0.884
CHF	1.000	0.961	0.840	0.919
NOK	1.000	1.087	0.846	0.943
NZD	1.000	1.174	0.936	0.994
SEK	1.000	1.029	0.747	0.938

Note: All MSEs in the upper (middle/lower) block were divided by the currency-specific MSEs from the (uncombined) GARCH (HAR/ARFIMA) forecasts.