Erschienen in: Jahrbücher für Nationalökonomie und Statistik ; 231 (2011), 1. - S. 134-152

Forecasting Multivariate Volatility using the VARFIMA Model on Realized Covariance Cholesky Factors*

By Roxana Halbleib**, Bruxelles, and Valeri Voev***, Aarhus

JEL C32; C53; G11

Forecasting; fractional integration; stochastic dominance; portfolio optimization; realized covariance.

Summary

This paper analyzes the forecast accuracy of the multivariate realized volatility model intro duced by Chiriac and Voev (2010), subject to different degrees of model parametrization and economic evaluation criteria. By modelling the Cholesky factors of the covariance matrices, the model generates positive definite, but biased covariance forecasts. In this paper, we provide empirical evidence that parsimonious versions of the model generate the best covariance fore casts in the absence of bias correction. Moreover, we show by means of stochastic dominance tests that any risk averse investor, regardless of the type of utility function or return distribu tion, would be better off from using this model than from using some standard approaches.

1 Introduction

In recent years much of the financial econometrics literature has dealt with modelling and forecasting multivariate volatility processes, which play a central role in many financial applications, including risk management and portfolio management. Many studies focus on developing multivariate extensions of the univariate generalized auto-regressive conditional heteroscedasticity (GARCH) approach of Bollerslev (1986) in order to capture the joint dynamics of daily return correlations and volatilities. Although

* Previous versions of this paper circulated under the title „Dynamic Modelling and Forecasting Realized Covariance Matrices“ and „Long Memory Modelling of Realized Covariance Ma trices“. The paper was written while Roxana Halbleib was research fellow at Department of Economics and Center of Finance and Econometrics (CoFE) at the University of Konstanz, Ger many. Financial support from the German Science Foundation within the project „Price , Liquid ity and Credit Risks: Measurement and Allocation“ is gratefully acknowledged. We would like to thank Ingmar Nolte, Winfried Pohlmeier and Joachim Grammig for helpful comments. Any remaining errors and inaccuracies are ours.

** The author gratefully acknowledge financial support from the Belgian National Bank and the IAP P6/07 contract, from the IAP programme (Belgian Scientific Policy), 'Economic policy and finance in the global economy'. The author is a member of ECORE, the recently created association between CORE and ECARES. Previous versions of this paper circulated under the maiden name of the author, Roxana Chiriac.

*** Financial support by the Center for Research in Econometric Analysis of Time Series, CREATES, funded by the Danish National Research Foundation, is gratefully acknowledged.

very popular, these approaches, known in literature as multivariate GARCH's (MGARCH), suffer from curse of dimensionality and encounter difficulties in ensuring the positive definiteness of the covariance forecasts (e.g. Bauwens et al. 2006).
As alternative to the GARCH approaches, a large body of literature focuses on measuring daily volatilities on basis of high frequency data. These new approaches, known as realized volatility (RV) estimators, exploit the richness of intraday information in order to increase the accuracy of daily volatility measures. Moreover, subject to the new estimators, daily volatility becomes observable and exhibits highly persistent dynamics, which are usually captured by means of autoregressive fractional integrated moving average (ARFIMA) models (e.g. Andersen et al. 2003; Oomen 2001) or heterogenous autoregressive (HAR) models (Corsi 2009). These approaches are mainly applied to univariate RV series, while modelling the dynamics of realized covariance matrices is not widely explored.
One contribution in this direction constitutes the approach developed by Chiriac and Voev (2010), who propose to model the dynamics of the Cholesky factors of realized covariance matrices by means of vector ARFIMA approach. This methodology has the advantage of capturing the long memory property of volatility, while simultaneously guaranteeing the positive definiteness of the resulting forecasts, without imposing any parameter restrictions. However, the implementation of the model carried out by Chiriac and Voev (2010) raises some questions, which we try to answer in the present paper by means of empirical evidence.
For empirical purposes, Chiriac and Voev (2010) choose to implement a very tightly parameterized version of the model, which implies only three parameters to estimate, regardless of the dimension of the covariance matrix. The fact that their approach implies a non-linear transformation of the covariance matrices leads to a non-zero expectation of the forecasting error. To account for this bias, they derive theoretically the necessary correction to ensure the unbiasedness of the covariance forecasts. However, when implementing the model, they ignore the heteroscedasticity property of model error terms, which plays a crucial role in deriving precise bias corrections. Moreover, the optimal portfolio problem which Chiriac and Voev (2010) use to evaluate the forecasting performance of their model against popular multivariate volatility approaches relies on heavy distributional assumptions of the portfolio return or certain functional forms of the investor's utility.
The purpose of this paper is to asses the quality of the multivariate volatility approach proposed by Chiriac and Voev (2010) with respect to the choice of model parametrization, bias correction and more general evaluation criteria. We provide empirical evidence that restricted versions of the model provide the best covariance matrix forecasts, without imposing any bias correction which serves as a robustness check of the implementation of Chiriac and Voev (2010). Furthermore, we use second order stochastic dominance tests to show that the proposed model outperforms standard volatility approaches. The stochastic dominance approach is very powerful in the sense that it implies optimality for any risk averse investor regardless of the functional form of the utility function and the distribution of returns.

Following this introduction, Section 2 presents the model of Chiriac and Voev (2010) whose extensions are thoroughly analyzed within two empirical applications in Section 3. Section 4 concludes.

2 The model

Let Y_t of dimension $n \times n$ be the realized covariance matrix at time t, where n represents the number of assets considered. The Cholesky decomposition of the matrix Y_t is given by the upper triangular matrix P_t, for which $P_t' P_t = Y_t$. Let $X_t = vech(P_t)$ be the vector obtained by stacking the upper triangular components of the matrix P_t in a vector. Chiriac and Voev (2010) propose to model the dynamics of the vector X_t of dimension $m \times 1$, where $m = n(n+1)/2$, by using a Vector Autoregressive Fractionally Integrated Moving Average or VARFIMA(p, d, q) model, as follows:[1]

$$\Phi(L)D(L)[X_t - BZ_t] = \Theta(L)\varepsilon_t, \qquad \varepsilon_t \sim N(0, \Sigma_t) \tag{1}$$

where X_t is the vector of dimension $m \times 1$ formed from the elements of the Cholesky decomposition of the realized covariance matrix Y_t. Z_t is a vector of exogenous variables of dimension $k \times 1$, B is a matrix of coefficients of dimension $m \times k$, $\Phi(L) = I_m - \Phi_1 L - \Phi_2 L^2 - \ldots - \Phi_p L^p$, $\Theta(L) = I_m + \Theta_1 L + \Theta_2 L^2 + \ldots + \Theta_q L^q$ are matrix lag polynomials with $\Phi_i, i = 1, \ldots, p$ and $\Theta_j, j = 1, \ldots, q$ – the AR- and MA-coefficient matrices, and $D(L) = diag\{(1 - L)^{d_1}, \ldots, (1 - L)^{d_m}\}$, where $d_1, \ldots, d_m$ are the degrees of fractional integration of each of the m elements of the vector X_t. Σ_t is the covariance matrix of ε_t. The roots of $\Phi(L)$ and $\Theta(L)$ are assumed to lie outside the unit circle and X_t is stationary if $d_i < 0.5$, for all $i = 1, \ldots, m$.

In this paper we employ an extension of the VARFIMA-Cholesky model introduced by Chiriac and Voev (2010) by allowing for conditional heteroscedasticity of the error term, as there is some evidence that volatility of volatility is time-varying: e.g., Corsi et al. (2008) show that the residuals of ARFIMA models fitted on univariate time series of realized (co)variances exhibit non-Gaussianity and volatility clustering. Consequently, they extend the ARFIMA framework by including a GARCH component on the volatility of ARFIMA residuals, which substantially improves the goodness-of-fit. Furthermore, following the arguments of Chiriac and Voev (2010), we refrain from including exogenous variables in our study and apply the following model:

$$\Phi(L)D(L)[X_t - c] = \Theta(L)\varepsilon_t, \qquad \varepsilon_t \sim N(0, \Sigma_t) \tag{2}$$

where c is a vector of constants of dimension $m \times 1$. Section 3 contains a discussion on the model specification of Σ_t given in Equation (2).

Similar to Chiriac and Voev (2010), we estimate the model in the final equation form (Lütkepohl 2005), restricting the AR polynomial to be a scalar polynomial. We use quasi maximum likelihood estimation, which, under the normality assumption (see Equation (2)), certain regularity conditions and a well specified conditional mean function, generates consistent estimators of the model parameters (Gourieroux/Monfort 1995).

Contrary to Chiriac and Voev (2010), who choose to implement a very restricted version of the original model: a common long memory parameter d, scalar AR and scalar MA polynomials[2], in the current study we aim at identifying and quantifying the conse-

[1] In this paper, for ease of exposition, we will call the approach introduced by Chiriac and Voev (2010) the VARFIMA Cholesky model.

[2] Chiriac and Voev (2010) pre test only the equality among the degrees of fractional integration, with out further testing the constraints on the elements of Θ. Furthermore they apply the estimation method proposed by Beran (1995), which circumvents the estimation of Σ_t and pre estimate the mean vector c as the sample mean of the vector X_t.

quences of gradually imposing parameter restrictions on the model performance both in- and out-of-sample.

3 Empirical application

This section presents the data used in the empirical study, the estimation of the model and the results of the forecasting evaluation against two standard MGARCH approaches.

3.1 Data

We use tick-by-tick bid and ask quotes from the NYSE Trade and Quotations (TAQ) database sampled from 9:45 until 16:00 over the period January 1, 2001 to June 30, 2006 (1381 trading days). Although the NYSE market opens at 9:30, we filter out the quotes recorded in the first 15 minutes in order to eliminate the opening auction effect on the price process. For the current analysis, we select the following six stocks: American Express Inc. (AXP), Citigroup (C), Home Depot Inc. (HD), Hewlett-Packard (HWP), International Business Machines (IBM) and JPMorgan Chase & Co (JPM), which are highly liquid. We use the previous-tick interpolation method, described in Dacorogna et al. (2001) in order to obtain a regularly spaced sequence of midquotes, which are thus sampled at the 5-minute and daily frequency, from which 5-minute and daily log returns are computed. Thus we obtain for each day a total of 75 intraday observations which are used to compute the realized variance-covariance matrix of that day. Table A1 in Appendix A reports summary statistics of both 5-minute and daily returns. We observe typical stylized facts such as overkurtosis and tendency for negative skewness of intra-daily and daily returns (across all six stocks, the average kurtosis of 5-minute return series is about 269.2, while of daily returns is about 10.9). For estimation, we scale up the daily and intradaily returns by 100, i. e., we consider percentage returns.

For each $t = 1, \dots, 1381$, we construct series of daily realized covariance matrices, Y_t as follows:

$$Y_t = \sum_{j=1}^{M} r_{j,t} r'_{j,t} \tag{3}$$

where $M = 75$ and $r_{j,t}$ is the $n \times 1$ vector of 5-minute returns computed as

$$r_{j,t} = p_{j\Delta,t} - p_{(j-1)\Delta,t}, \quad j = 1, \dots, M$$

where $\Delta = 1/M$ and $p_{j\Delta,t}$ is the log midquote price at time $j\Delta$ in day t. By construction, the realized covariance matrices are symmetric and, for $n < M$, they are positive definite. Similar to Chiriac and Voev (2010), we refine the estimator Y_t by averaging over realized covariance estimators computed on 30 regularly 300 seconds-spaced subgrids starting at seconds $1, 11, 21, \dots, 291$. As we are interested in the covariance matrix of the whole day (close-to-close), and Y_t estimates only its open-to-close portion, we use the scaling method introduced by Hansen and Lunde (2005) to adapted to the multivariate case: we scale each (co)variance estimate corresponding to the trading period by an average scaling factor, which incorporates the overnight information over all series. This procedure preserves the positive-definiteness of the resulting covariance matrix. The descriptive statistics of the 21 realized variance and covariance series are presented in Table A2 in Appendix A.

After computing the series of realized covariance matrices, we construct the series of Cholesky factors, which inherit the long memory property of realized (co)variances documented by Andersen and Bollerslev (1997) and Andersen et al. (2001). Evidence of this fact is presented in Figure A1 in Appendix B: the sample autocorrelations of the Cholesky factors of the realized covariance matrix of the six stocks decay at a slow rate, similar to the autocorrelations of the realized (co)variance series.

3.2 Forecasting using the VARFIMA-Cholesky model with heteroscedastic error terms

In this subsection we assess the forecasting accuracy of the VARFIMA-Cholesky model under different degrees of model parametrization and with consideration of bias correction, aspects which are only touched upon in Chiriac and Voev (2010). In this sense, the following results can be thought of as a robustness check of their results with respect to different model specifications.

In order to analyze the impact of imposing parameter restrictions on the performance of the VARFIMA-Cholesky model, we employ first the model in Equation (2) with AR and MA polynomials of order one and a constant c of dimension $m \times 1$ (Model 1). This unrestricted version of the model implies a total number of $\frac{(n^2+n+2)^2}{4}$ parameters. Secondly, we proceed with imposing equal degrees of fractional integration for all $X_{i,t}$ series, $i = 1, \ldots, m$ (Model 2), which leads to a total number of $\frac{(n^2+n+1)^2+7}{4}$ parameters to estimate. Chiriac and Voev (2010) implement this approach as result of pre-testing, however in a more restrictive framework than the one considered here: they impose a scalar MA polynomial, which might affect the inference. Thirdly, besides equality among the degrees of fractional integration, we also impose zero off-diagonal restrictions on the parameter matrix Θ (Model 3), which implies $n(n+1)+2$ parameters to estimate. Finally, we set the MA polynomial to a scalar parameter while imposing only one degree of fractional integration for all Cholesky series (Model 4). This restricted version involves a total of $\frac{n(n+1)+6}{2}$ parameters to estimate. Gradually imposing parameter restrictions aims at identifying the model specification which generates the best multivariate volatility forecasts.

As already mentioned in Section 2, there is empirical evidence at the univariate level that the volatility of volatility is time varying. Besides properly describing the dynamic properties of realized covariance matrices, this result may play a considerably role in the decision of bias correcting the covariance forecasts based on the VARFIMA-Cholesky approach, discussed at length in Chiriac and Voev (2010). For this purpose, we choose to parameterize Σ_t in Equation (2) by means of the BEKK approach introduced by Engle and Kroner (1995). However, as a result of the large number of parameters involved by the estimation of the heteroscedastic model, we choose to apply the diagonal BEKK$(1,1,1)$ specification given by:

$$\Sigma_t = C'C + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + B'\Sigma_{t-1}B, \tag{4}$$

where C is a $m \times m$ upper triangular parameter matrix and A and B are $m \times m$ diagonal parameter matrices.

For forecasting purposes, we apply the forecasting procedure described by Chiriac and Voev (2010), whereby we account for the time-varying variance of the forecast errors. Since ε_t in Equation (2) are assumed to be normally distributed, the s-step ahead forecast errors of X_t are also normally distributed as:

$$u_{t,t+s} \equiv X_{t+s} - E_t[X_{t+s}] \sim N(0, \Sigma_{t,s}), \tag{5}$$

where $E_t[X_{t+s}] \equiv \hat{X}_{t+s}$ is the optimal s-step ahead predictor of X_t given the information set consisting of all relevant information up to and including t and

$$\Sigma_{t,s} = E[(X_{t+s} - E_t[X_{t+s}])(X_{t+s} - E_t[X_{t+s}])'] = E[u_{t,t+s}u'_{t,t+s}] = \sum_{i=0}^{s-1} \Psi_i \Sigma_{t+s-i} \Psi'_i \tag{6}$$

and

$$\Sigma_{t+s-i} = C'C + A'\varepsilon_{t+s-i-1}\varepsilon'_{t+s-i-1'}A + B'\Sigma_{t+s-i-1}B.$$

It follows that the forecast errors of the one-step ahead forecast, $u_{t,t+1}$, are normally distributed with zero mean and variance-covariance matrix $\Sigma_{t,1} = \Sigma_t$.

Chiriac and Voev (2010) show that although the forecast errors of X_t have zero mean (unbiased $\hat{X}_{t+s}$), the mean of the forecast errors of Y_t, which is a quadratic transformation of $\hat{X}_t$, is no longer zero (biased $\hat{Y}_{t+s}$) and depends on the forecast error variance, which in our context is time-varying. Thus, the bias correction for $E_t[Y_{ij,t+s}]$ is obtained from the elements of matrix $\Sigma_{t,s}$ using the following formula:

$$\sigma^*_{t,ij} = \sum_{l=1+\frac{i(i-1)}{2}}^{\frac{i(i+1)}{2}} \sigma_{t,s(l,l+\frac{j(j-1)}{2}-\frac{i(i-1)}{2})}, \tag{7}$$

where $j \geq i$, $i = 1, \ldots, n$ and $\sigma_{t,s(v,u)}$ is the v,u-th element of the matrix $\Sigma_{t,s}$. The high parametrization involved and the time-varying property of the covariance matrix motivate Chiriac and Voev (2010) to avoid estimating Σ_t. As an alternative, they suggest to apply a data-driven bias correction of $E_t[Y_{ij,t+s}]$, which turns out to be ineffective in their empirical example. However, this result cannot be generalized, without further assessment.

In the present empirical application we test whether the flexible specification of Σ_t leads to empirical gains via bias correcting the forecasts of daily multivariate volatility based on the VARFIMA-Cholesky model. As the correction is based on estimated parameters, it introduces some estimation noise to the forecasts. Since it is not clear whether this effect dominates the effect of the reduced bias, we consider forecasts with and without bias correction. It is important to note that while the bias correction may theoretically lead to non-positive definite forecasts, in the application we present below all forecasts remain positive definite.

In order to keep the estimation tractable, for this empirical application we choose arbitrarily two out of the six stocks considered: HWP and JPM[3] and compute the realized covariance matrices without the refinement described in section (3.1). Table A3 in Appendix A reports the results from estimating the four versions of the VARFIMA-Cholesky model described above. In Model 1, the long memory property is captured in the values of d-parameters around 0.37 for the diagonal elements of the Cholesky decomposition matrices and around 0.29 for the off-diagonal elements. However, when estimating only one parameter (Model 2, 3 and 4), the degrees of fractional integration

[3] Empirical results for other stock combinations can be obtained from the authors on request.

stabilizes around 0.36. Regardless of the model choice, the AR parameter is significantly positive, while the diagonal MA-parameters are significantly negative. Similar results are obtained by Oomen (2001) in the univariate modelling of log realized volatilities. The fact that the off-diagonal elements of the MA matrix are in general not significant at 5% level, validates the constraints imposed in Model 3, while the similarity[4] among the diagonal elements Θ in Model 3 endorse the parametric specification in Model 4.

The value of the multivariate Ljung-Box statistics at lag 30 (last row of Table A3) decreases as we relax parameter restrictions which is to be expected, as the in-sample fit improves by adding flexibility to the model. Though one could argue that the values of the statistic are high, the reduction is considerable when compared to its value of 9803.994 computed on the original series.

In terms of interdependencies among the variance and covariance series, it is important to note that the diagonal and scalar specifications do not exclude the possibility of spillover effects among the series. This is due to the non-linearity of the Cholesky transformation. Even if the X_t series are estimated independently of each other, the resulting Y_t components are functions of the X_t series, and therefore related to each other.

To check the robustness of the results with respect to the sampling frequency of the returns used to construct the realized covariance matrices, we carry out the above analysis with realized covariances based on 30-minute returns, which should be less susceptible to market microstructure noise. The results we obtained are qualitatively similar, so we refrain from reporting them here.[5]

For forecasting purposes, we divide the overall sample of 1381 daily observations in two subsets: an in-sample period on which we estimate the model, and an out-of-sample period which serves to evaluate the forecasting performance. The in-sample period contains initially the first 1181 observations. In each forecasting step, the in-sample period is increased by one observation, the models are re-estimated and a new one-step ahead forecast is made. This procedure is carried out 200 times, and as a result we obtain a total of 200 one-step ahead forecasts. In this empirical exercise we evaluate the performance of one-day ahead forecasts of the four model specifications before and after the bias correction. As a proxy for the unobserved volatility of day $t+1$ we consider the realized volatility Y_{t+1}, which for statistical evaluation purposes, is proven to be a better estimator of daily volatility than an estimator constructed from daily data (Patton/Sheppard 2009).

To evaluate the forecasting performance of the four model specifications we employ Mincer-Zarnovitz type regressions as in, e.g., Andersen and Bollerslev (1998) and Andersen et al. (2005). The regression is given by:

$$Y_{ij,t+1} = \alpha_{ij} + \beta_{ij} \hat{E}_t[Y_{ij,t+1}] + v_{ij,t+1}, \quad i,j = 1,\dots,n, \quad t = 1181,\dots,1380, \tag{8}$$

where $v_{ij,t+1}$ is the error term of the regression.

The corresponding values of the R^2 statistics for the four model forecasts are contained in Table 1 and the corresponding α and β coefficients are reported in Table A4 in Appendix A. From the entries of Table 1 it is evident that, in all cases, the R^2 measure based on forecasts from Model 4 is higher or equal to the R^2 values for the other model specifica-

4 The null hypothesis of equality among the elements of the diagonal Θ in Model 3 is not rejected at 5 % level.

5 The results can be obtained on request from the authors.

Table 1 Results of forecast evaluations. The table reports the R^2 of the Mincer-Zarnovitz regressions in Equation (8) and the RMSE of daily covariance matrix forecasts based on the VARFIMA approach applied on realized covariance matrices of dimension 2 × 2. The total number of forecasts is 200.

	Bias Correction				**No Bias Correction**			
Model	1	2	3	4	1	2	3	4
Y_{11}	0.180	0.183	0.196	0.199	0.180	0.182	0.196	0.198
Y_{12}	0.167	0.167	0.177	0.179	0.170	0.170	0.179	0.182
Y_{22}	0.255	0.256	0.258	0.259	0.258	0.260	0.260	0.260
RMSE	3.403	3.399	3.410	3.414	3.284	3.279	3.279	3.280

tions. This suggests that imposing restrictions of equal degrees of fractional integration for all series and scalar MA parameters improves the forecasting ability of the proposed model.

Regarding the magnitude of the R^2 statistic, similar results have been obtained by Andersen et al. (2003) for the out-of-sample one-step-ahead exchange rates univariate volatility forecast based on an ARFIMA$(1, 0.4, 0)$ specification. The R^2's they report range between 20 % and 25 %. It should be noted also that exchange rates have much lower volatility (and variation in volatility) compared to stocks, and hence we can a priori expect higher predictive power of the model for exchange rate volatility than for stock volatility. Since the realized covariance is subject to estimation noise (due to market microstructure noise, and, additionally for the covariance, due to non-synchronicity), the R^2's understate the true predictive power of the models. Based on the entries of Table A4, we cannot reject the hypotheses that $a_{ij} = 0$ and $\beta_{ij} = 1$ in all cases of Model 4 without bias correction, and in most of the cases for the other models.

The results on the root mean squared error (RMSE) of the out-of-sample forecasts based on the Frobenius norm (see the last row of Table 1) complete the outcomes from the Mincer Zarnovitz regression approach: all models without bias correction deliver the smallest forecasting errors. This suggests that unbiased predictors are not necessarily the optimal choice when forecasting the daily covariance; the forecasts with no bias correction are, on average, closer to the true realizations than the bias-corrected ones. Therefore, we can conclude that restricted model specifications such as the ones given by Model 4 with no bias correction generally provide the best predictions of positive definite and symmetric variance-covariance matrices. These results validate the assumptions and the choice of model specification implemented by Chiriac and Voev (2010).

3.3 Economic evaluation of VARFIMA-Cholesky forecasts: A comparison approach

Based on the results of the previous empirical application, in what follows we aim at assessing the forecasting performance of Model 4 with no bias correction against some standard models by means of economic criteria.

To assess the merits of the model, similar to Chiriac and Voev (2010) we consider a risk-averse investor who faces the problem of optimal portfolio selection among the six stocks considered, subject to different covariance forecasts. Alternatively to Chiriac and Voev (2010), who evaluate the forecast performance by means of standard criteria, such as mean-variance efficiency, which is adequate only if the investor has a quadratic utility function or the return distribution is fully described by its first two moments (e.g normal distribution), we employ here a much more powerful criteria, which is applicable for any concave utility function and any return distribution. We also evaluate the portfolio performance under short selling restrictions.

In order to keep the estimation tractable, we set the intercept vector c equal to the sample mean of X_t. As a result of pre-estimating c, the resulting „second-step“ QML standard errors of the estimated parameters are incorrect. Therefore, to assure a robust inference we derive the standard errors by employing the subsampling method developed by Politis and Romano (1994a) and Politis et al. (1999) for dependent and cross-correlated time series.

For our comparative study we consider two popular MGARCH approaches for the conditional covariance matrix: the DCC model (Engle 2002) and the diagonal BEKK model (Engle/Kroner 1995). We assume here that the conditional mean of daily returns, μ is constant and we estimate it along with the MGARCH parameters.

DCC-GARCH

The DCC-GARCH model, proposed by Engle (2002), is a multivariate GARCH model with univariate GARCH(1,1) conditional variances, $h_{ii,t}$, and time-varying conditional correlations:

$$H_t = D_t R_t D_t,$$

where $D_t = diag(h_{11,t}^{1/2} \dots h_{nn,t}^{1/2})$ and

$$h_{ii,t} = w_i^* + \alpha_i^* \varepsilon_{i,t-1}^2 + \beta_i^* h_{ii,t-1},$$

where $w_i^*, \alpha_i^*, \beta_i^* \geq 0$ and $\alpha_i^* + \beta_i^* < 1, \forall i = 1, \dots, n$. $\varepsilon_{i,t}$ is the residual from the mean specification $\varepsilon_{i,t} = r_{i,t} \quad \mu$.

$$R_t = (diag(Q_t))^{\ \frac{1}{2}} Q_t (diag(Q_t))^{\ \frac{1}{2}},$$

where Q_t is an $n \times n$ symmetric and positive definite matrix given by:

$$Q_t = (1 \quad \theta_1^* \quad \theta_2^*)Q + \theta_1^* u_{t-1} u_{t-1}' + \theta_2^* Q_{t-1},$$

where u_t is the vector of standardized residuals with elements

$$u_{i,t} = \frac{\varepsilon_{i,t}}{\sqrt{h_{ii,t}}}, \quad i = 1, \dots, n$$

and Q is the unconditional covariance of u_t. For $n = 6$ assets, the DCC model implies a total of 26 parameters, which are estimated by maximizing the normal pseudo-likelihood.

Diagonal BEKK

Engle and Kroner (1995) suggest a multivariate GARCH model, where the conditional return covariance matrix H_t is parameterized as a function of its lags and lagged squared innovations:

$$H_t = C^{*'} C^* + A^{*'} \varepsilon_{t-1} \varepsilon_{t-1}' A^* + B^{*'} H_{t-1} B^*,$$

where C^* is an upper $n \times n$ triangular matrix and A^* and B^* are $n \times n$ parameter matrices. Under certain restrictions, described in Engle and Kroner (1995), the resulting covariance matrices are assured to be positive definite and stationary. In the present paper, we estimate the diagonal specification of the model, where A^* and B^* are diagonal matrices. The model includes 39 parameters, which are estimated by maximum likelihood.

Before turning to the forecasting evaluation, we present here briefly the estimation results for the full sample of data. The results of the DCC and diagonal BEKK models are reported in Table A5. Table A6 reports the estimated values of the three parameters implied by the restricted version of VARFIMA-Cholesky approach given by Model 4 along with their bootstrapped standard errors. The results are in line with the ones described in the previous empirical exercise: all parameters are significant at 5% level and the estimated value of the degree of fractional integration d is smaller than 0.5, which indicates that the series are stationary; the autoregressive parameter is significantly positive, while the moving average parameter is significantly negative.

In order to assess the economic value of the three model forecasts, we construct portfolios which are supposed to maximize the utility of a risk-averse investor. If the utility function is second degree polynomial or logarithmic and/or the return distribution is completely characterized by its first two moments (e.g. normal distribution), the portfolio optimization is reduced to finding the asset weights which minimize the portfolio volatility while aiming for a target expected return or maximize the portfolio return while targeting a certain volatility (Markowitz 1952).

We assume that our investor minimizes his portfolio volatility subject to a fixed expected return (10% p.a.). He is allowed (Scenario 1) or prohibited (Scenario 2) to sell assets he does not own (short selling). In this context, the optimal portfolio is given by the solution to the following quadratic problem:

$$\min_{w_{t+1|t}} \quad w'_{t+1|t}\hat{H}_{t+1|t}w_{t+1|t}$$

subject to:

Scenario 1: $w'_{t+1|t}E_t[R_{t+1}] + (1 \quad w'_{t+1|t}\iota)R_f = R^*$

Scenario 2: $w'_{t+1|t}E_t[R_{t+1}] + (1 \quad w'_{t+1|t}\iota)R_f = R^*$, $w_{t+1|t} \geq 0$,

where $\hat{H}_{t+1|t}$ is the covariance forecast at day t for day $t+1$, $w_{t+1|t}$ is the $n \times 1$ vector of portfolio weights chosen at day t for day $t+1$, ι is an $n \times 1$ vector of ones, R_f is the risk free rate (3% p.a.) and R^* is the target expected return (10% p.a.).

Given that there is hardly any predictable return variation at the daily level, we assume that the expected returns are constant as in Fleming et al. (2001) and Fleming et al. (2003). Having solved for the optimal weights based on the three different conditional covariance forecasts, we compute the ex-post daily portfolio returns and the corresponding Sharpe ratios, given by:

$$SR = \frac{R_p \quad R_f}{\hat{\sigma}_{R_p}},$$

where R_p is the sample mean and $\hat{\sigma}_{R_p}$ is the sample standard deviation of the ex-post realized portfolio return series.

Table 2 reports the annualized realized Sharpe ratios and standard deviations of the three sets of minimum-covariance portfolios. The numbers in this table should be interpreted

Table 2 Annualized Sharpe ratios and standard deviations of out-of-sample realized portfolio returns

Portfolio	VARFIMA	DCC	BEKK
		Sharpe Ratios	
Scenario 1	0.849	0.618	0.493
Scenario 2	0.447	0.327	0.363
		Standard Deviations	
Scenario 1	12.30	12.95	13.15
Scenario 2	15.10	16.86	16.83

simply as indicative that for the considered sample the VARFIMA-Cholesky-based portfolio delivers a smaller standard deviation and a higher Sharpe ratio than the GARCH-based ones. We relegate the formal comparison of these results by means of significance tests to the following discussion on stochastic dominance which is a much more general way of assessing whether a given return distribution is „better“ than another one.

The assumption of a „mean-variance“ investor is rather restrictive from an economic point of view. A more meaningful evaluation of the optimality of the portfolios can be achieved by comparing the whole distribution of the portfolio returns as opposed to just the first two moments. For example, the skewness and the shape of the tails of the return distribution are relevant in the investment decision process. Therefore, in what follows, we compare the VARFIMA-Cholesky-, DCC- and BEKK-based portfolio return distributions by means of stochastic dominance tests. To this end we need an additional definition.

Defination 1: Let X_1 and X_2 be two real random variables. It is said that X_1 s-th order stochastically dominates X_2 ($X_1 \succeq_s X_2, s > 0$) if and only if $F^s_{X_1}(x) \leq F^s_{X_2}(x)$ for all x with strict inequality for some x, where $F^s_{X_i}(x) = \int_{-\infty}^{x} F^{s-1}_{X_i}(t)dt$ for $s \geq 2$, $F^1_{X_i}(x) = F_{X_i}(x)$ and $F_{X_i}(x)$ is the cumulative distribution function (CDF) of X_i, $i = 1, 2$.

Fishburn (1980) and Bawa (1975), among others, show that X_1 s-th order stochastically dominates X_2 if and only if $E[u(X_1)] \geq E[u(X_2)]$ (with strict inequality for some x from the common support of X_1 and X_2) for every function u with $(\; 1)^{j+1} u^{(j)}(x) \geq 0$ for all $j \in 1, \ldots, s$ where $u^{(j)}(x)$ stands for the j-th derivative of $u(x)$. The implications of this for our analysis are as follows: Let us have two optimal portfolio strategies (forecasting models), A and B and $R_{p,A}$ and $R_{p,B}$ be the realized returns of the two minimum-variance portfolios with CDF's $F_A(x)$ and $F_B(x)$. A risk-averse investor with an increasing utility function $u(x)$, translating into $u^{(1)}(x) \geq 0$ and $u^{(2)}(x) \leq 0$, chooses portfolio A over portfolio B if and only if portfolio A second order stochastically dominates portfolio B, i. e., $\int_{-\infty}^{r} F_A(x)dx \leq \int_{-\infty}^{r} F_B(x)dx$ for $r \in \Pi$, where Π is the common support of $R_{p,A}$ and $R_{p,B}$, with strict inequality for at least one $r \in \Pi$. In this case the investor has a larger expected utility from portfolio A than from portfolio B, $E[u(R_{p,A})] \geq E[u(R_{p,B})]$.

Comparing the integrated cumulative distributions (i. e., $F^2(\cdot)$) of the VARFIMA-Cholesky-based portfolio pairwise against the DCC- and BEKK-based ones, we find that the former is strictly smaller for each value of the common return support, which is a first indication that the VARFIMA-Cholesky-based portfolio second order stochastically dominates the other two portfolios. To check the robustness of these results, we apply a number of stochastic dominance tests on the estimated distributions.

Table 3 P-values of the LMW and KRS tests for 2nd order stochastic dominance. Portfolio *A* denotes the minimum covariance portfolio based on the VARFIMA-Cholesky forecasts. The critical values of the tests are derived from bootstrap procedures which account for serial and cross dependence of the observations: subsampling bootstrap („Sub“) and stationary bootstrap („SB“). The subsampling size is b 30 observations. The „block“ length of the stationary bootstrap is driven by the average value of the first order serial correlation of the series.

		Scenario 1		***Scenario 2***	
Test/Portfolio B		DCC	BEKK	DCC	BEKK
		LMW Test			
Sub	$H_0 : A \succeq_2 B$	0.436	0.383	0.389	0.447
	$H_0 : B \succeq_2 A$	0.441	0.319	0.005	0.000
		KRS Test			
SB	$H_0 : A \not\succeq_2 B$	0.235	0.137	0.019	0.078
	$H_0 : B \not\succeq_2 A$	0.990	0.990	0.990	0.990

The literature on stochastic dominance tests is separated into two groups: one group (McFadden 1989; Klekan et al. 1991; Barett and Donald 2003; Linton et al. 2005) tests the null hypothesis of dominance ($H_0 : A \succeq_2 B$) against the alternative of non-dominance ($H_1 : A \not\succeq_2 B$), while the other group (Kaur et al. 1994; Davidson/Duclos 2000) tests the null hypothesis of non-dominance, against the alternative hypothesis of dominance. Most of these tests are developed on the assumptions of i. i. d. and cross-independent observations. Due to the fact that we deal with serially (due to GARCH effects) and cross-dependent portfolio returns, we apply here two tests which account for these features: the Linton et al. (2005) (LMW) test and Kaur et al. (1994) (KRS) test. We apply the subsampling procedure (Sub) of Politis and Romano (1994a) and Politis et al. (1999) and the stationary bootstrap (SB) procedure of Politis and Romano (1994b) to obtain consistent critical values for the test.

Table 3 reports the *p*-values of the LMW and KRS tests for various null hypotheses described in the first column. Regardless of the investment strategy, all tests with the null hypothesis of stochastic dominance of the VARFIMA-Cholesky portfolio against the other two portfolios have a *p*-value well in excess of 35 % indicating a strong support for the null hypothesis. Changing the testing direction, we reject the null hypothesis of dominance of MGARCH portfolios against the VARFIMA-Cholesky for Scenario 2 at 10 % significance level.

Similar results are obtained from the KRS test with null hypotheses of non-dominance. Generally, for Scenario 2 we find ample evidence for the dominance of the VARFIMA-Cholesky-based portfolio, while for Scenario 1 the data is inconclusive, but still delivers support for the VARFIMA-Cholesky approach. Referring again to Table 2, it is evident that for Scenario 2, the differences in the variance of the portfolio distributions are substantial, which is the reason for the much more clear-cut test results compared to Scenario 1. The relevance of the constrained portfolio optimization problem in Scenario 2 is supported by the fact that many institutional investors are forbidden by law from short selling. Furthermore, a recent study of Boehmer et al. (2008) reveals that on the NYSE only up to 2% of short sales are undertaken by individual traders.

Thus, we conclude that the VARFIMA-Cholesky approach is a worthwhile strategy to pursue, as it has the potential of providing added economic value, regardless of the investor's utility function form or return distribution assumption.

4 Conclusion

In this paper we bring empirical evidence on the quality of multivariate volatility forecasts based on the VARFIMA-Cholesky model proposed by Chiriac and Voev (2010), subject to different investment conditions and evaluation criteria. Introduced to capture the dynamics of realized covariance matrices by modelling their Cholesky factors with the vector fractional integrated ARMA approach, the model explicitly accounts for the long memory of financial volatility and guarantees the positive definiteness of forecasts without imposing parameter restrictions.

One potential shortcoming of this approach concerns the bias of the covariance matrix forecasts, originating in the nonlinear transformation of Cholesky factor forecasts. In the present study, we aim at assessing the robustness of the forecasting performance of VARFIMA-Cholesky approach subject to different model specifications and bias correction with time varying volatility of volatility. Thus, through gradually imposing parameter restrictions and accounting for the model heteroscedasticity by means of diagonal BEKK approach, we show that restricted versions of the model provide generally the best daily covariance matrix forecasts without imposing any bias correction.

Similar to Chiriac and Voev (2010), we assess the forecasting performance of the model, by applying it to an optimal portfolio selection problem. Hoverer, contrary to Chiriac and Voev (2010), who apply standard evaluation criteria implying restrictive assumptions, we show by means of stochastic dominance tests, that *any* risk averse investor would achieve the highest (among the models considered) expected utility by using the VARFIMA-Cholesky forecasts to optimize his portfolio, regardless of the investment constraints.

Appendix

Table A1 Summary statistics of 5-minute and daily stock returns from 1st January 2001 to 30th June 2006. The means are scaled by 10^4.

Stock	Mean	Max	Min	Std. dev	Skew	Kurt
			5-minute returns			
AXP	0.0113	0.0703	-0.1843	0.0022	-4.7063	485.0690
HWP	-0.0016	0.1112	-0.1597	0.0031	-1.0157	256.3915
JPM	-0.0037	0.0774	-0.1186	0.0025	-0.9637	137.7105
HD	-0.0241	0.1082	-0.1271	0.0024	-2.2291	270.6422
C	0.0006	0.0845	-0.1035	0.0022	-0.4016	157.5951
IBM	-0.0119	0.1086	-0.1071	0.0019	1.5253	307.8203
			Daily returns			
AXP	1.1391	0.1034	-0.1464	0.0193	-0.2277	8.5927
HWP	0.3494	0.1567	-0.2066	0.0267	-0.0234	10.7708
JPM	-0.2844	0.1578	-0.2019	0.0218	0.0683	13.7154
HD	-1.7161	0.1228	-0.1509	0.0210	-0.2066	9.2915
C	0.1761	0.1178	-0.1726	0.0184	-0.4100	13.2778
IBM	-0.7115	0.1173	-0.1106	0.0177	0.4465	10.2498

Table A2 Summary statistics of realized variances and realized covariances of the stocks AXP, C, HWP, JPM HD and IBM. The realized variances and covariances are calculated from 5-minute intraday returns, as described in the main text. The realized variances and covariances are scaled by 10^2

Stock	Mean	Max	Min	Std. dev	Skew	Kurt
			Realized Variance			
AXP	0.0390	0.8339	0.0011	0.0635	5.4105	46.2969
HWP	0.0656	1.4397	0.0028	0.0961	6.6996	75.0095
JPM	0.0490	2.8130	0.0017	0.1083	14.9024	334.1691
HD	0.0413	0.7317	0.0012	0.0533	4.9629	41.7344
C	0.0386	1.4113	0.0013	0.0738	10.0771	151.6528
IBM	0.0267	0.8111	0.0013	0.0387	8.1390	131.8510
			Realized Covariance			
AXP-HWP	0.0154	0.4085	-0.0145	0.0290	6.3709	61.8298
AXP-JPM	0.0169	0.6035	-0.0791	0.0325	8.3144	117.3552
AXP-HD	0.0143	0.3223	-0.0060	0.0256	5.5456	48.0592
AXP-C	0.0171	0.4900	-0.0130	0.0312	5.8290	50.4219
AXP-IBM	0.0128	0.3288	-0.0185	0.0226	5.0769	43.5852
HWP-JPM	0.0170	0.4047	-0.0054	0.0294	6.2420	59.3477
HWP-HD	0.0150	0.3183	-0.1175	0.0249	7.0555	82.6114
HWP-C	0.0171	0.2913	-0.0473	0.0270	4.7059	34.1161
HWP-IBM	0.0150	0.3334	-0.0026	0.0233	14.3477	317.0420
JPM-HD	0.0152	0.3637	-0.0345	0.0268	5.8616	56.1417
JPM-C	0.0221	1.2769	-0.0552	0.0498	6.2820	61.3870
JPM-IBM	0.0141	0.4329	-0.0098	0.0253	5.3664	48.0463
HD-C	0.0156	0.4063	-0.0051	0.0269	7.4878	92.5154
HD-IBM	0.0127	0.2234	-0.0037	0.0195	4.8518	35.7024
C-IBM	0.0142	0.4839	-0.0151	0.0252	7.8865	110.4219

Table A3 Estimation results of the VARFIMA(1,*d*,1)-diagonal BEKK(1,1,1) model. P-values based on QML standard errors are reported in parenthesis.

Parameter	**Model 1**	**Model 2**	**Model 3**	**Model 4**
d_1	0.374 (0.000)			
d_2	0.287 (0.000)	0.365 (0.000)	0.390 (0.000)	0.364 (0.012)
d_3	0.389 (0.000)			
ϕ	0.499 (0.001)	0.452 (0.000)	0.668 (0.017)	0.212 (0.178)
θ_{11}	-0.592 (0.000)	-0.538 (0.000)	-0.693 (0.001)	
θ_{12}	0.029 (0.596)	0.007 (0.876)		
θ_{13}	0.093 (0.017)	0.105 (0.009)		
θ_{21}	0.015 (0.352)	0.011 (0.470)		
θ_{22}	-0.634 (0.000)	-0.663 (0.000)	-0.772 (0.000)	-0.301 (0.170)
θ_{23}	0.123 (0.000)	0.113 (0.000)		
θ_{31}	0.022 (0.256)	0.024 (0.234)		
θ_{32}	0.020 (0.639)	0.034 (0.391)		
θ_{33}	-0.500 (0.000)	-0.429 (0.000)	-0.672 (0.007)	
c_1	1.895 (0.001)	0.562 (0.000)	0.750 (0.000)	0.659 (0.186)
c_2	0.351 (0.033)	0.003 (0.671)	0.034 (0.510)	0.004 (0.098)
c_3	0.377 (0.447)	0.168 (0.105)	0.151 (0.058)	0.207 (0.119)
C_{11}	0.062 (0.000)	0.018 (0.000)	0.015 (0.000)	0.019 (0.005)
C_{12}	0.043 (0.052)	0.015 (0.086)	0.016 (0.086)	0.018 (0.014)
C_{13}	0.058 (0.009)	-0.017 (0.013)	0.008 (0.013)	0.018 (0.012)
C_{22}	0.169 (0.000)	0.052 (0.000)	0.049 (0.000)	0.053 (0.012)
C_{23}	0.024 (0.725)	0.001 (0.860)	-0.004 (0.860)	0.004 (0.031)
C_{33}	0.000 (0.992)	0.000 (0.995)	0.000 (0.998)	0.000 (0.000)
A_{11}	0.075 (0.006)	0.076 (0.005)	-0.056 (0.005)	0.076 (0.027)
A_{22}	0.225 (0.000)	0.232 (0.000)	0.201 (0.000)	0.232 (0.092)
A_{33}	0.366 (0.000)	0.362 (0.000)	0.370 (0.000)	0.362 (0.057)
B_{11}	0.995 (0.000)	0.995 (0.000)	0.996 (0.000)	0.995 (0.001)
B_{22}	0.972 (0.000)	0.971 (0.000)	-0.978 (0.000)	0.971 (0.025)
B_{33}	0.921 (0.000)	0.923 (0.000)	-0.924 (0.000)	0.923 (0.019)
LB(30) of $\hat{\varepsilon}_t$	335.116	342.465	435.266	432.548

Table A4 Estimated parameters of the Mincer-Zarnovitz regression. Round parentheses report standard errors. $\hat{\alpha}_{ij}$ and their standard errors are scaled by 10^5.

Model	1	2	3	4	1	2	3	4
	Bias Correction				No Bias Correction			
	$\hat{\alpha}_{ij}$							
Y_{11}	2.531	2.396	2.001	1.911	1.787	1.644	1.217	1.123
	(1.976)	(1.981)	(1.957)	(1.951)	(2.088)	(2.093)	(2.068)	(2.062)
Y_{12}	0.481	0.535	0.320	0.551	0.369	0.432	0.219	0.453
	(0.519)	(0.510)	(0.526)	(0.489)	(0.531)	(0.521)	(0.536)	(0.499)
Y_{22}	1.250	1.254	0.325	-0.058	0.925	0.931	0.178	-0.102
	(0.962)	(0.958)	(1.060)	(1.100)	(0.990)	(0.986)	(1.071)	(1.104)
	$\hat{\beta}_{ij}$							
Y_{11}	0.998	1.008	1.050	1.058	0.989	0.999	1.040	1.048
	(0.152)	(0.152)	(0.152)	(0.151)	(0.151)	(0.151)	(0.150)	(0.150)
Y_{12}	1.026	1.012	1.120	1.029	1.013	0.997	1.099	1.011
	(0.163)	(0.161)	(0.172)	(0.157)	(0.160)	(0.157)	(0.167)	(0.153)
Y_{22}	1.158	1.159	1.335	1.415	1.062	1.061	1.186	1.239
	(0.141)	(0.141)	(0.161)	(0.170)	(0.128)	(0.128)	(0.143)	(0.149)

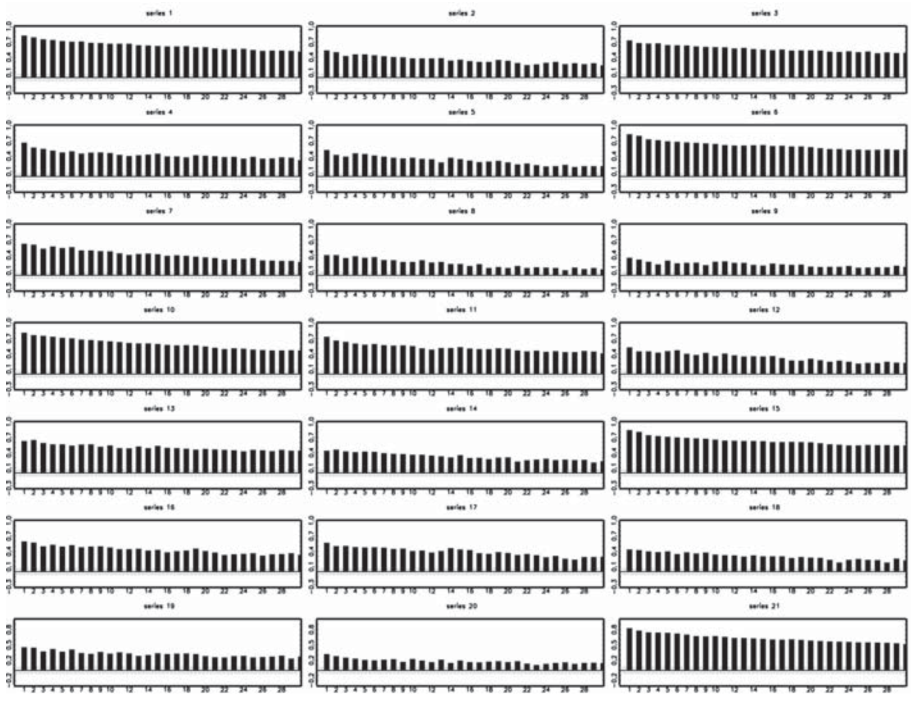

Figure A1 Autocorrelograms of the Cholesky factor series X_t.

Table A5 Estimation results of the diagonal BEKK(1,1,1) and DCC model. QML standard errors are reported in parenthesis.

Parameter/Stock	AXP	HWP	JPM	HD	C	IBM
			diagonal BEKK(1,1,1)			
μ_i	0.1039	0.0578	-0.1042	0.0556	-0.0349	-0.0886
	(0.0689)	(0.0523)	(0.0542)	(0.0417)	(0.0425)	(0.0290)
C^*	0.0420	-0.0221	-0.0301	-0.0892	0.0485	-0.0139
	(0.0219)	(0.0386)	(0.0246)	(0.0441)	(0.0223)	(0.0278)
		-0.0538	-0.0060	0.0706	0.0394	-0.0393
		(0.0235)	(0.0128)	(0.0173)	(0.0253)	(0.0456)
			-0.0042	-0.0171	0.0032	-0.1054
			(0.0288)	(0.0207)	(0.0136)	(0.0389)
				0.1607	0.0850	0.1946
				(0.0851)	(0.0117)	(0.0485)
					0.1499	0.1777
					(0.0304)	(0.0292)
						0.1412
						(0.0340)
diag(A^*)	0.9845	0.9947	0.9788	0.9870	0.9814	0.9857
	(0.0153)	(0.0014)	(0.0097)	(0.0056)	(0.0054)	(0.0068)
diag(B^*)	0.0617	0.0620	0.0353	0.0336	0.0203	0.0180
	(0.0410)	(0.0611)	(0.0358)	(0.0440)	(0.0309)	(0.0350)
			DCC of Engle (2002)			
μ_i	0.0717	0.0490	0.0313	0.0182	0.0264	0.0262
	(0.0354)	(0.0589)	(0.0340)	(0.0418)	(0.0320)	(0.0534)
w_i^*	0.0236	0.0144	0.0117	0.0155	0.0167	0.0273
	(0.0157)	(0.0163)	(0.0079)	(0.0139)	(0.0137)	(0.0550)
α_i^*	0.0867	0.0097	0.0658	0.0403	0.0670	0.0714
	(0.0341)	(0.0046)	(0.0270)	(0.0138)	(0.0359)	(0.1191)
β_i^*	0.9087	0.9871	0.9315	0.9549	0.9266	0.9194
	(0.0314)	(0.0065)	(0.0252)	(0.0165)	(0.0373)	(0.1259)
	θ_1	0.0067		θ_2	0.9776	
		(0.0031)			(0.0139)	

Table A6 Estimation results of VARFIMA-Cholesky model. Bootstrapped standard errors are reported in parenthesis.

AR	0.4664 (0.0052)
MA	-0.3190 (0.0065)
fractional integration	0.4664 (0.0063)

References

Andersen, T. G., T. Bollerslev (1997), Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long Run in High Frequency Returns. Journal of Finance 52: 975 1005.

Andersen, T. G., T. Bollerslev (1998), Answering the Skeptics: Yes, Standard.Volatility Models Do Provide Accurate Forecasts, International Economic Review 39: 885 905.

Andersen, T. G., T. Bollerslev, F. X. Diebold, H. Ebens (2001), The Distribution of Stock Return Volatility. Journal of Financial Economics 61: 43 76.

Andersen, T. G., T. Bollerslev, F. X. Diebold, P. Labys (2003), Modeling and Forecasting Realized Volatility. Econometrica 71(2): 579 625.

Andersen, T. G., T. Bollerslev, N. Meddahi (2005), Correcting the Errors: Volatility Forecast Evaluation Using High Frequency Data and Realized Volatilities. Econometrica 73(1): 279 296.

Barett, G., S. Donald (2003), Consistent Tests for Stochastic Dominance. Econometrica 71: 71 104.

Bauwens, L., S. Laurent, J. Rombouts (2006), Multivariate GARCH Models: a Survey, Journal of Applied Econometrics 21: 79 109.

Bawa, V. S. (1975), Optimal Rules for Ordering Uncertain Prospects. Journal of Financial Eco nomics 2: 95 121.

Beran, J. (1995), Maximum Likelihood Estimation of the Differencing Parameter for Invertible Short and Long Memory Autoregressive Integrated Moving Average Models. Journal of the Royal Statistical Society 57: 659 672.

Boehmer, E., C. Jones, X. Zhang (2008), Which Shorts are Informed? Journal of Finance 63: 491 527.

Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity. Journal of Econometrics 31: 307 327.

Chiriac, R., V. Voev (2010), Modelling and Forecasting Multivariate Realized Volatility. Journal of Applied Econometrics, forthcoming.

Corsi, F. (2009), A Simple Approximate Long Memory Model of Realized Volatility. Journal of Financial Econometrics 7: 174 196.

Corsi, F., U. Kretschmer, S. Mittnik, C. Pigorsch (2008), Volatility of Realized Volatility. Econo metric Reviews 27: 46 78.

Dacorogna, M. M., R. Gençay, U. A. Müller, R. B. Olsen, O. V. Pictet (2001), An Introduction to High Frequency Finance. San Diego Academic Press.

Davidson, R., J. Y. Duclos (2000), Statistical Inference for Stochastic Dominance and for the Measurement of Poverty and Inequality, Econometrica, 68: 1435 1464.

Engle, R. (2002), Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroscedasticity Models. Journal of Business and Economic Statistics 20: 339 350.

Engle, R., K. Kroner (1995), Multivariate Simultaneous GARCH. Econometric Theory 11: 122 150.

Fishburn, P. C. (1980), Continua of Stochastic Dominance Relations for Unbounded Probability Distributions. Journal of Mathematical Economics 7(3): 271 285.

Fleming, J., C. Kirby, B. Ostdiek (2001), The Economic Value of Volatility Timing. Journal of Finance 56: 329 352.

Fleming, J., C. Kirby, B. Ostdiek (2003), The Economic Value of Volatility Timing using Realized Volatility. Journal of Financial Economics 67: 473 509.

Gourieroux, C., A. Monfort (1995), Statistics and Econometric Models. Vol.1. Cambridge Uni versity Press, Cambridge.

Hansen, P. R., A. Lunde (2005), A Realized Variance for the Whole Day Based on Intermittent High Frequency Data. Journal of Financial Econometrics 3: 525 554.

Kaur, A., B. L. S. P. Rao, H. Singh (1994), Testing for Second Order Stochastic Dominance of Two Distributions. Econometric Theory 10(5): 849 866.

Klekan, L., R. McFadden, D. McFadden (1991), A Robust Test for Stochastic Dominance. Working Paper, Department of Economics, MIT.

Linton, O., E. Maasoumi, Y. J. Whang (2005), Consistent Testing for Stochastic Dominance under General Sampling. Review of Economic Studies 72: 735 765.

Lütkepohl, H. (2005), New Introduction to Multiple Time Series Analysis. Springer, Berlin.

Markowitz, H. (1952), Portfolio Selection. Journal of Finance 7: 77 91.

McFadden, D. (1989), Testing for Stochastic Dominance. Pp. 113 134 in: T. Fomby, T. Seo (eds.), Studies in the economics of Uncertainty. Springer Verlag.

Oomen, R. (2001), Using High Frequency Stock Market Index Data to Calculate, Model and Forecast Realized Return Variance. European University, Economics Discussion Paper No. 2001/6.

Patton, A., K. Sheppard (2009), Evaluating Volatility and Correlation Forecasts. In: T. Andersen, R. Davis, J. P. Kreiss, T. Mikosch (eds.), Handbook of Financial Time Series. Springer.

Politis, D., J. Romano (1994a), Large Sample Confidence Regions based on Subsamples under Minimal Assumptions. Annals of Statistics 22: 2031 2050.

Politis, D., J. Romano (1994b), The Stationary Bootstrap, Journal of the American Statistical Association 89: 1303 1313.

Politis, D., J. Romano, M. Wolf (1999), Subsampling. Springer Verlag, New York.

Roxana Halbleib, European Center for Advanced Research in Economics and Statistics (EC ARES), Université libre de Bruxelles, Solvay Brussels School of Economics and Management, Avenue F. Roosevelt, 50, CP114/04, 1050 Brussels, Belgium, and CoFE.
Roxana.Halbleib@ulb.ac.be

Valeri Voev, School of Economics and Management, Aarhus University, 8000 Aarhus C, Den mark, and CREATES.
vvoev@creates.au.dk