

Lasso-based forecast combinations for forecasting realized variances

Wilms I, Rombouts J, Croux C.



Lasso-based forecast combinations for forecasting realized variances

Ines Wilms^{a*}, Jeroen Rombouts^b, Christophe Croux^a

^a *Faculty of Economics and Business, KU Leuven*

^b *ESSEC Business School*

Abstract. Volatility forecasts are key inputs in financial analysis. While lasso based forecasts have shown to perform well in many applications, their use to obtain volatility forecasts has not yet received much attention in the literature. Lasso estimators produce parsimonious forecast models. Our forecast combination approach hedges against the risk of selecting a wrong degree of model parsimony. Apart from the standard lasso, we consider several lasso extensions that account for the dynamic nature of the forecast model. We apply forecast combined lasso estimators in a comprehensive forecasting exercise using realized variance time series of ten major international stock market indices. We find the lasso extended “ordered lasso” to give the most accurate realized variance forecasts. Multivariate forecast models, accounting for volatility spillovers between different stock markets, outperform univariate forecast models for longer forecast horizons.

Keywords. Forecast combination, Hierarchical lasso, Lasso, Ordered Lasso, Realized variance, Volatility forecasting

*Corresponding author: ines.wilms@kuleuven.be

1 Introduction

Volatility forecasts of financial assets are of great importance in finance. For instance, risk management, asset allocation, option pricing and trading in financial markets (see e.g. Knight and Satchell, 2007) require reliable volatility forecasts. We use lasso-based forecast methods (Tibshirani, 1996) to obtain such volatility forecasts of various international stock market indices. The lasso has become a popular forecast method since it provides parsimonious forecast models (see e.g. Fan et al., 2011). Its usefulness for volatility forecasting, however, is underexplored. Our aim is to investigate (i) which lasso-based forecast method provides the most accurate volatility forecasts and, (ii) whether these volatility forecasts become more accurate if spillover effects between the different stock market indices are accounted for, and (iii) whether forecast combination in the lasso-based forecast methods improves further the forecast accuracy.

The inherent problem in obtaining volatility forecasts is that volatility is unobservable. Therefore, one traditionally considered latent variable models such as stochastic volatility models or GARCH models (Bollerslev, 1986). However, with the advent of high-frequency data, accurate estimates of volatility have become available (Andersen et al., 2001; Barndorff-Nielsen and Shephard, 2002). Realized variances are computed as the sum of squared intraday returns for a particular day, thereby making volatility observable. The advantage of working with realized variances is that traditional time series models like univariate AutoRegressive (AR) and multivariate Vector AutoRegressive (VAR) models can be used for forecasting.

This paper proposes the use of lasso-based estimators for AutoRegressive and Vector AutoRegressive models for log-realized variance of ten international stock market indices. The lasso (Tibshirani, 1996) provides sparse model estimates, meaning that many of the lagged autoregressive parameters are estimated as exactly zero. As such, a parsimonious model is obtained. Such parsimonious models reduce estimation error and therefore may yield more accurate forecasts. Recent papers by Audrino and Knaus (2016) and Audrino et al. (2015) discuss estimation of the AR model by the lasso to obtain realized variance forecasts. Our paper differs from theirs in that we not only consider lasso estimation of univariate (AR) models but also of multivariate (VAR) models, see e.g. Callot et al. (2016) for forecasting large realized covariance matrices. Especially when stock markets are going through a period of high uncertainty, like financial crises, their volatility interdependence tends to increase (e.g. Diebold and Yilmaz, 2015, Chapter 4). This justifies the use of VAR models over AR models to capture the spillover structure between the different stock market indices' volatilities (i.e. effects from one stock market index volatility on another). The parsimony advantage of the lasso becomes even more important for estimating such VAR models since they become intractable for large systems of time series, see e.g. Gelper et al. (2016). In fact, a parsimonious spillover structure is obtained since only the most important spillovers effects between different stock market indices are estimated non-zero.

While Audrino and Knaus (2016), Audrino et al. (2015) and Callot et al. (2016) only consider the use of the standard lasso, we also consider two extensions of the lasso that embed lag selection into the estimation procedure: the hierarchical lasso (e.g. Bien et al., 2013 for regression models, Nicholson et al., 2016a for time series models) and the ordered lasso (e.g. Suo and Tibshirani, 2015 for univariate models). In addition to the aim of attaining sparse estimates, the hierarchical and ordered lasso aim at attaining estimates with a low maximal lag order. Both estimators are very flexible since they allow each time series to have its own lag structure in each equation of the VAR. Most previous studies on multivariate models, by contrast, require the strong assumption of a single universal lag order.

An important choice that needs to be made for lasso-based estimators concerns the selection of the sparsity parameter that determines the degree of sparsity of the autoregressive estimates (or “model parsimony”), as measured by the number of non-zero coefficients. Previous studies either use information criteria or cross-validation to select the “optimal” sparsity parameter. Forecasts are then obtained from the lasso-based estimator with the selected sparsity parameter. We, in contrast, do not select an optimal sparsity parameter but rather combine forecasts obtained with different sparsity parameters in a forecast combination approach. Such forecast combinations are typically more robust to a wrongly chosen sparsity parameter.

We apply our forecast models to monthly log-realized variance of ten international stock market indices observed over January, 2000 to April, 2016. While most realized variance time series share similar patterns, the in-sample estimation results reveal that their respective optimal lag lengths selected by the different lasso methods can differ substantially. The ordered lasso has best in-sample fit over all markets. Studying the spillovers when estimating multivariate models allows to identify which markets drive others. By measuring out-of-sample forecast accuracy in terms of absolute forecast errors in an expanding window setup, it turns out that the ordered lasso also dominates. Forecast combination slightly improves the results, though applying the model confidence set approach shows that the difference with the ordered lasso method with a single sparsity parameter is not significant. Another interesting finding is that the multivariate forecasting models outperform univariate models for the longer forecast horizons we consider (i.e. 3 and 6 months) meaning that spillover effects have long term forecasting power.

The remainder of this paper is structured as follows. Section 2 discusses the univariate and multivariate forecast models, together with the three lasso-based estimators used to obtain log-realized variance forecasts. Section 3 describes the data and discusses the results on lag length estimation for the univariate forecast model and spillover structure for the multivariate forecast model. Section 4 details the out-of-sample forecast performance of the different forecast methods. Section 5 provides a range of robustness checks. Section 6 concludes.

2 Modeling and forecasting realized variance

First, we discuss the univariate and multivariate forecast models for the log-realized variances we study in this paper. Next, we introduce the three lasso-based estimators to obtain estimates and forecasts. Finally, we explain how we perform forecast combination.

2.1 Modeling realized variances

Univariate forecast model Denote the log-realized variance of stock market index i at time point t by $\log \text{RV}_t^{(i)}$. We consider an AutoRegressive (AR) Model of order p for the log-realized variance of stock market index i

$$\log \text{RV}_t^{(i)} = \beta_1 \log \text{RV}_{t-1}^{(i)} + \beta_2 \log \text{RV}_{t-2}^{(i)} + \dots + \beta_p \log \text{RV}_{t-p}^{(i)} + \epsilon_t^{(i)}, \quad (1)$$

for $t = p + 1, \dots, T$. The log-realized variance at time point t is modeled as a function of its own past values up to p periods ago. We assume, without loss of generality, that all time series are centered such that no intercept is included, and that the error terms ϵ_t are white noise with mean zero, variance σ^2 .

Multivariate forecast model As a multivariate extension, we consider the Vector AutoRegressive (VAR) Model of order p for the log-realized variances of the $q = 10$ stock market indices

$$\begin{cases} \log \text{RV}_t^{(1)} &= \sum_{j=1}^q \sum_{l=1}^p \beta_{1j,l} \log \text{RV}_{t-l}^{(j)} + \epsilon_t^{(1)} \\ \log \text{RV}_t^{(2)} &= \sum_{j=1}^q \sum_{l=1}^p \beta_{2j,l} \log \text{RV}_{t-l}^{(j)} + \epsilon_t^{(2)} \\ \vdots &= \vdots \\ \log \text{RV}_t^{(q)} &= \sum_{j=1}^q \sum_{l=1}^p \beta_{qj,l} \log \text{RV}_{t-l}^{(j)} + \epsilon_t^{(q)}, \end{cases} \quad (2)$$

where the log-realized variance of stock market index i at time point t is modeled as a function of its own past values and the past values of the other stock market indices up to p periods ago. We assume that the multivariate error time series with q components, $(\epsilon_t^{(1)}, \dots, \epsilon_t^{(q)})'$, is a multivariate white noise with mean $\mathbf{0}$, and covariance matrix Σ . The VAR model, in contrast to the AR model, captures spillovers between the different stock market indices. Indeed $\beta_{ij,l}$ measures the effect at lag l of stock market index j on stock market index i . We call this effect a ‘‘spillover’’ effect if stock market index $j \neq i$.

2.2 Lasso-based estimators and realized variance forecasts

We use three lasso-based estimators to obtain realized variance forecasts. They differ in terms of the sparsity structure they impose on the autoregressive parameters.

Lasso The lasso estimator with sparsity parameter $\lambda > 0$ is given by

$$\hat{\beta}_{\text{lasso},\lambda} = \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^q \left(\frac{1}{2} \sum_{t=p+1}^T \left(\log \text{RV}_t^{(i)} - \sum_{j=1}^q \sum_{l=1}^p \beta_{ij,l} \log \text{RV}_{t-l}^{(j)} \right)^2 \right) + \lambda \sum_{i=1}^q \sum_{j=1}^q \sum_{l=1}^p |\beta_{ij,l}|, \quad (3)$$

where the first term in equation (3) is a least squares criterion. For the univariate forecast model, $q = 1$ and the vector β collects all p estimated parameter values. For the multivariate model, β collects all $q \times q \times p$ estimated parameter values. The larger the sparsity parameter λ , the more elements of $\hat{\beta}$ will be put to zero, hence the more parsimonious the forecast model.

A drawback of the lasso is that it does not account for the dynamic nature of the AR/VAR model. Indeed, it allows higher order lags of a particular time series to be included in the model while lower order lags of the same time series are absent. We might want a higher order lag to be included in the model only if its lower order lags are already included. To accommodate this feature, we consider two extensions of the lasso.

Hierarchical Lasso The hierarchical lasso estimator (Bien et al., 2013; Nicholson et al., 2016a) with sparsity parameter $\lambda > 0$ is given by

$$\hat{\beta}_{\text{hierlasso},\lambda} = \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^q \left(\frac{1}{2} \sum_{t=p+1}^T \left(\log \text{RV}_t^{(i)} - \sum_{j=1}^q \sum_{l=1}^p \beta_{ij,l} \log \text{RV}_{t-l}^{(j)} \right)^2 \right) + \lambda \sum_{i=1}^q \sum_{j=1}^q \sum_{l=1}^p \|\beta_{ij,(l:p)}\|_2, \quad (4)$$

where $\|\cdot\|_2$ denotes the Euclidean norm of a vector and $\beta_{ij,(l:p)} = (\beta_{ij,l}, \dots, \beta_{ij,p})$ is an $p - l + 1$ vector collecting the furthest lags of time series j in equation i . The hierarchical lasso accounts for the dynamic nature of the AR/VAR by forcing lower order lagged coefficients of a time series in one of the equations to be selected before its higher order lagged coefficients: for all $l' > l$, $\hat{\beta}_{ij,l} = 0$ implies $\hat{\beta}_{ij,l'} = 0$ almost surely.

Ordered Lasso The ordered lasso estimator (Suo and Tibshirani, 2015) with sparsity parameter $\lambda > 0$ is given by

$$\hat{\beta}_{\text{orderedlasso},\lambda} = \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^q \left(\frac{1}{2} \sum_{t=p+1}^T \left(\log \text{RV}_t^{(i)} - \sum_{j=1}^q \sum_{l=1}^p \beta_{ij,l} \log \text{RV}_{t-l}^{(j)} \right)^2 \right) + \lambda \sum_{i=1}^q \sum_{j=1}^q \sum_{l=1}^p (\beta_{ij,l}^+ + \beta_{ij,l}^-), \quad (5)$$

subject to the constraints $\beta_{ij,1}^+ \geq \beta_{ij,2}^+ \geq \dots \geq \beta_{ij,p}^+ \geq 0$, and $\beta_{ij,1}^- \geq \beta_{ij,2}^- \geq \dots \geq \beta_{ij,p}^- \geq 0$, and each $\beta_{ij,l} = \beta_{ij,l}^+ - \beta_{ij,l}^-$, for all $1 \leq i, j \leq q$ and $1 \leq l \leq p$. Similar to the hierarchical lasso, the ordered lasso encourages lower order lagged coefficients to be selected before its higher order lagged coefficients. Additionally, the ordered lasso encourages the absolute values of the lagged effects of each time series j in each equation i to be monotone non-increasing as we move farther back in time.

Automatic lag selection An advantage of these lasso estimators is that they perform automatic lag selection due to the sparsity induced in $\hat{\beta}$. In the univariate forecast model, the lag length is estimated as

$$\hat{p} = \sup\{1 \leq l \leq p : \hat{\beta}_l \neq 0\},$$

the lag corresponding to the largest non-zero estimated parameter. Similarly in the multivariate forecast model, the lag length p_{ij} for each time series $1 \leq j \leq q$ in equation $1 \leq i \leq q$ of the VAR is estimated as

$$\hat{p}_{ij} = \sup\{1 \leq l \leq p : \hat{\beta}_{ij,l} \neq 0\}.$$

Each time series is allowed to have its own lag structure in each equation of the VAR.

Algorithmic details The lasso optimization problems (3), (4), and (5) can be solved efficiently using proximal gradient methods. Proximal gradient methods extend traditional gradients methods to non-smooth objective functions (see e.g. Bien and Witten, 2016). Starting from an initial value, the estimate of the parameter vector is updated until converge by evaluating the proximal operator at the gradient step we would take if we were minimizing the least squares criterion alone. The key step is thus to evaluate the proximal operator which involves solving a convex optimization problem. While the lasso and hierarchical lasso use the Fast Iterative Soft-Thresholding Algorithm to evaluate the proximal operator, the ordered lasso uses the Pool Adjacent Violators Algorithm. For more details, we refer to Suo and Tibshirani (2015) and Jenatton et al. (2011) and references therein. All computations are carried out in R version 3.2.1. The code of the algorithms relies on the R-packages `orderedLasso` (Friedman et al., 2015) and `BigVAR` (Nicholson et al., 2016b) and is available from the authors upon request.

2.3 Forecast combination

The selection of the sparsity parameter λ poses an important choice for lasso estimators. In contrast to previous studies on lasso-based forecasting, we do not select an “optimal” sparsity parameter. Instead, we combine forecasts obtained from lasso estimations with several sparsity parameters. We obtain several forecasts, one for each choice of the sparsity parameter λ , and weigh these individual forecasts to construct a final log-realized variance forecast. We expect this forecast combination approach to be more robust to a wrongly chosen sparsity parameter.

To construct our different forecasts, we consider a logarithmic spaced grid of sparsity parameters of length $L = 20$ in the interval $[\lambda_1, \lambda_L]$, with $\lambda_1 = 0$ and λ_L an estimate of the sparsity parameter that shrinks all parameters to zero, see Friedman et al. (2010) for the lasso. Similar expressions are obtained for the hierarchical and ordered lasso. If the number of parameters to be estimated is larger than the number of time points, we take $\lambda_1 = \lambda_L/10$. For each value λ_m , for $1 \leq m \leq L$, we estimate the (AR or VAR) model using either the lasso from equation (3), the hierarchical lasso from equation (4), or the ordered lasso from equation (5). This results in L different estimates $\widehat{\beta}_{\lambda_m}$ having their own degree of sparsity with corresponding h -step ahead direct forecasts of the log-realized variance for stock market index i

$$\widehat{\log RV}_{t+h, \lambda_m}^{(i)} = \sum_{j=1}^q \sum_{l=1}^p (\widehat{\beta}_{\lambda_m})_{ij, l} \log RV_{t+1-l}^{(i)}.$$

We use direct forecasts since these are more robust to model misspecification than iterated forecasts (Marcellino et al., 2006). The final forecast of the log-realized variance of stock market index i at time point $t+h$

is then given by

$$\widehat{\log \text{RV}}_{t+h}^{(i)} = \sum_{m=1}^L w_m \cdot \widehat{\log \text{RV}}_{t+h, \lambda_m}^{(i)}, \quad (6)$$

the weighted sum of the L individual log-realized variance forecasts obtained from the lasso, hierarchical or ordered lasso with sparsity parameter λ_m . As forecast combination weights w_m we take

$$w_m = \frac{\exp(-0.5\text{BIC}_{\lambda_m})}{\sum_{m=1}^L \exp(-0.5\text{BIC}_{\lambda_m})}. \quad (7)$$

Hence, the weights depend on the Bayesian Information Criterion BIC_{λ_m} of each forecast model $1 \leq m \leq L$ (e.g. Elliott and Timmermann, 2016, Chapter 14) as given by

$$\text{BIC}_{\lambda_m} = n \cdot \log(\text{Loss}_{\lambda_m}) + \text{df}_{\lambda_m} \cdot \log(n),$$

where $n = T - p$, Loss_{λ_m} is the first term in objective functions (3), (4), and (5) using sparsity parameter λ_m , and df_{λ_m} is the number of non-zero components of $\widehat{\beta}_{\lambda_m}$. When forecasting the log-realized variances, we take a maximum lag length of $p = 36$. The highest weight is given to the model having the smallest value of the BIC.

Our forecast combination approach is similar to Bayesian Model Averaging (BMA, see e.g. Koop, 2003, Chapter 11). BMA gives a final forecast by weighting different models according to their posterior model probability. Hence, BMA is mainly motivated as a way of dealing with model uncertainty. Note that the lasso can be seen as the mode of a Bayesian estimator with independent double-exponential prior (Tibshirani, 1996). The sparsity parameter λ enters as a hyperparameter in the prior distribution. Hence, we have different models, one for each choice of the hyperparameter λ entering the prior. As such, our forecast combination approach deals with uncertainty regarding the appropriate choice of the hyperparameter λ . Since we use BIC-based weights to construct our final forecast, we also give the highest weight to the model with the highest probability given the data.

3 Data and in-sample results

In this section, we discuss the data, the estimated lag lengths of the univariate models, and the spillover structure of the multivariate models.

3.1 Data

We use data on realized variances for ten different stock market indices. We consider a selection of American, Asian and European stock market indices: (1) the Amsterdam Exchange Index “AEX”, (2) the French stock index Cotation Assistée en Continu “CAC”, (3) the German stock index Deutscher Aktienindex “DAX”, (4) the American Dow Jones Industrial Average “DJIA”, (5) the stock index of the Eurozone “EUROSTOXX”,

Table 1: Descriptives: Mean, standard deviation, skewness, kurtosis and first order autocorrelation of each log-realized variance time series.

	Mean	St.Dev.	Skewness	Kurtosis	AutoCorr _{lag=1}
AEX	3.210	0.901	0.756	3.242	0.764
CAC	3.401	0.843	0.494	3.056	0.758
DAX	3.434	0.870	0.598	3.102	0.780
DJIA	2.760	0.871	0.878	3.760	0.745
EUROSTOXX	3.293	0.821	0.621	3.383	0.749
FTSE	2.864	0.865	0.681	3.425	0.770
NASDAQ	3.241	0.959	0.687	2.640	0.800
NIKKEI	3.361	0.608	0.357	3.945	0.659
SMI	2.882	0.864	0.928	3.660	0.727
SP500	2.776	0.896	0.773	3.462	0.767

(6) the United Kingdom Financial Times Stock Exchange Index 100 “FTSE”, (7) the American stock index “NASDAQ”, (8) the Japanese stock index “NIKKEI”, (9) the Swiss market index “SMI”, and (10) the American Standard & Poor’s 500 market index “SP500”.

Daily realized variance measures based on five minute returns are taken from Oxford-Man Institute of Quantitative Finance (publicly available on <http://realized.oxford-man.ox.ac.uk/data/download>). One of our objectives is the comparison of univariate and multivariate forecast models. We therefore aggregate the data to monthly frequency between January 2000 and April 2016, i.e. $T = 196$ observations. An advantage of working with monthly data instead of daily data is that we avoid the problem of non-overlapping trading days in the multivariate forecast model of the different international stock market indices. Note that in Section 6 below, we compare our monthly univariate forecasting results with the forecasts from models estimated on daily realized variances. Following standard practice, realized variance is log-transformed ensuring that their distribution more closely resembles the normal distribution and lives on the whole real line. Figure 1 plots the log-realized variances. The Augmented Dickey-Fuller test confirms that the time series are stationary. Descriptive statistics are given in Table 1. After log-transforming each time series, its skewness and kurtosis more resemble that of a normal distribution. The log-realized variance of NASDAQ is the most volatile and it has the highest first order autocorrelation. The log-realized variance of NIKKEI, on the contrary, is the least volatile and has the lowest first order autocorrelation.

3.2 Lag length estimation for univariate models

Table 2 reports the estimated lag lengths and the BIC values of the ten stock market indices for the forecast combined lasso, hierarchical lasso and ordered lasso applied to the full sample. The estimated lag length of

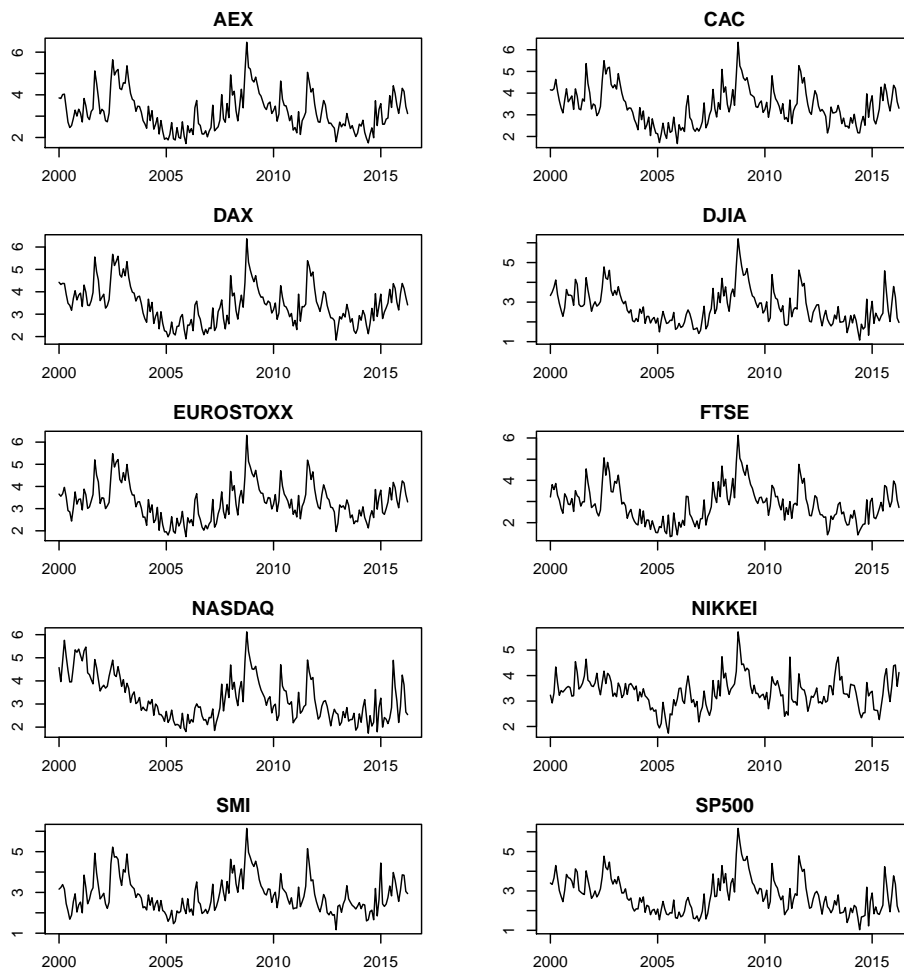


Figure 1: Log-realized variances for the ten stock market indices from January 2000 to April 2016.

Table 2: Estimated lag lengths \hat{p} and BIC (full sample).

Stock market index	Lasso		Hierarchical Lasso		Ordered Lasso	
	\hat{p}	BIC	\hat{p}	BIC	\hat{p}	BIC
AEX	4.63	173.44	5.00	-69.47	3.00	-160.07
CAC	31.00	-99.16	5.00	-84.11	2.99	-174.96
DAX	4.93	-140.66	4.00	-96.02	2.01	-181.90
DJIA	3.45	-42.32	6.02	-58.83	4.92	-149.00
EUROSTOXX	2.00	-112.83	6.00	-86.12	2.98	-178.05
FTSE	35.09	-77.24	6.01	-57.51	3.00	-182.95
NASDAQ	25.16	-71.52	7.00	-30.88	3.97	-154.92
NIKKEI	1.00	-148.62	3.85	-184.02	1.45	-232.36
SMI	1.00	147.33	4.00	-89.18	4.53	-157.50
SP500	3.34	142.56	5.00	-58.54	4.66	-147.03

each stock market index i is given by

$$\hat{p}_i = \sum_{m=1}^L w_m \hat{p}_{i,\lambda_m}.$$

In line with our expectations, the hierarchical and ordered lasso give much lower estimated lag lengths than the lasso for the majority of stock market indices (exceptions are EUROSTOXX, NIKKEI, SMI and SP500). For the hierarchical and ordered lasso, the lower order lags of the maximum non-zero lag are all non-zero, thereby encouraging a parsimonious, low maximal lag order model. For the lasso, in contrast, the estimated lag length is the highest order non-zero lagged coefficient; lower order lags can be both zero or non-zero. Among the three estimators, the ordered lasso consistently gives the best in-sample fit, i.e. lowest value of BIC for each stock market index.

The estimated lag lengths of the American stock market indices are typically the highest, followed by the European stock market indices. The American stock market indices have, overall, the highest volatility (see Table 1), and contain several volatility spikes. For more turbulent time series like these, higher order AR models typically have more explanatory power than lower order AR models. By contrast, the estimated lag length for the Japanese stock market index NIKKEI is the lowest. Since it also has the lowest stock market volatility, it is characterized by calmer market conditions. Simple, low lag order AR models then perform best.

3.3 Spillover structure in multivariate models

Figure 2 visualizes the lag matrix of estimated lag lengths for the forecast combined lasso, hierarchical lasso and ordered lasso applied to the full sample. For each cell ij , the estimated lag length is

$$\hat{p}_{ij} = \sum_{m=1}^L w_m \hat{p}_{ij,\lambda_m}.$$

If $\hat{p}_{ij} \neq 0$, there is a spillover effect from stock market index i to stock market index j , if $\hat{p}_{ij} = 0$ there is no spillover effect. For instance, the ordered lasso indicates a spillover effect from AEX to DAX since the estimated lag length is $\hat{p}_{13} = 1$. In contrast, there is no spillover effect from AEX to DJIA since $\hat{p}_{14} = 0$.

In line with the results from the univariate forecast models, (i) the lag matrix of the hierarchical and ordered lasso is much sparser with overall lower estimated lag lengths than the one of the lasso, and (ii) American stock market indices have, overall, high own (cfr. diagonal of lag matrix) lag lengths, NIKKEI has overall, a low own lag length. Furthermore, note that the diagonal of the lag matrix is non-zero (overall, and for all estimators), indicating that each stock market index log-realized variance is influenced by its own past values.

The lasso-based estimators allow each time series to have its own lag structure in each equation of the VAR, as can be seen from the lag matrices from Figure 2. It is important to note that the lag matrix is not a measure of the spillover strengths, but it still allows us to easily distinguish “net drivers” and “net receivers” of stock market index volatility. Net drivers influence more other stock market indices than that they are influenced by other indices. Net receivers are more influenced by other stock market indices than that they influence other indices. We focus on the results of the ordered lasso since it provides the best in-sample fit (i.e BIC= 628.57 versus 908.04 for the hierarchical lasso, and 1065.37 for the lasso). For some stock market indices, like FTSE, the own passed lags are the main driving factor (cfr. near zero FTSE column in lag matrix ordered lasso), others like AEX or EUROSTOXX are influenced by several other stock market indices. Similarly, some stock market indices, like NIKKEI, do not drive any of the other stock market indices’ volatility (cfr. zero NIKKEI row in lag matrix ordered lasso). Overall, we find that American stock market indices are net drivers, the European stock market indices are net receivers, and NIKKEI is a pure receiver of other stock market indices’ volatility. These findings are in line with previous research, see e.g. Diebold and Yilmaz (2015, Chapter 4) and references therein.

4 Forecast Accuracy

We use (i) four estimators: the lasso, hierarchical lasso, ordered lasso, and least squares; (ii) two forecast combination approaches: forecast combination versus no forecast combination; and (iii) two forecast models: univariate and multivariate. We compare the forecast performance along these three dimensions resulting in $16 = 4 \times 2 \times 2$ forecast methods.

Regarding the forecast combined least squares, (AR or VAR) models are estimated by least squares, each having lag order m , for $1 \leq m \leq p_{LS}$. The final forecast combines these forecasts by giving the highest weight to the model with the lowest BIC, see equation (6). An advantage of the lasso estimators compared to the least squares is that the former are computable if the number of parameters to be estimated exceeds

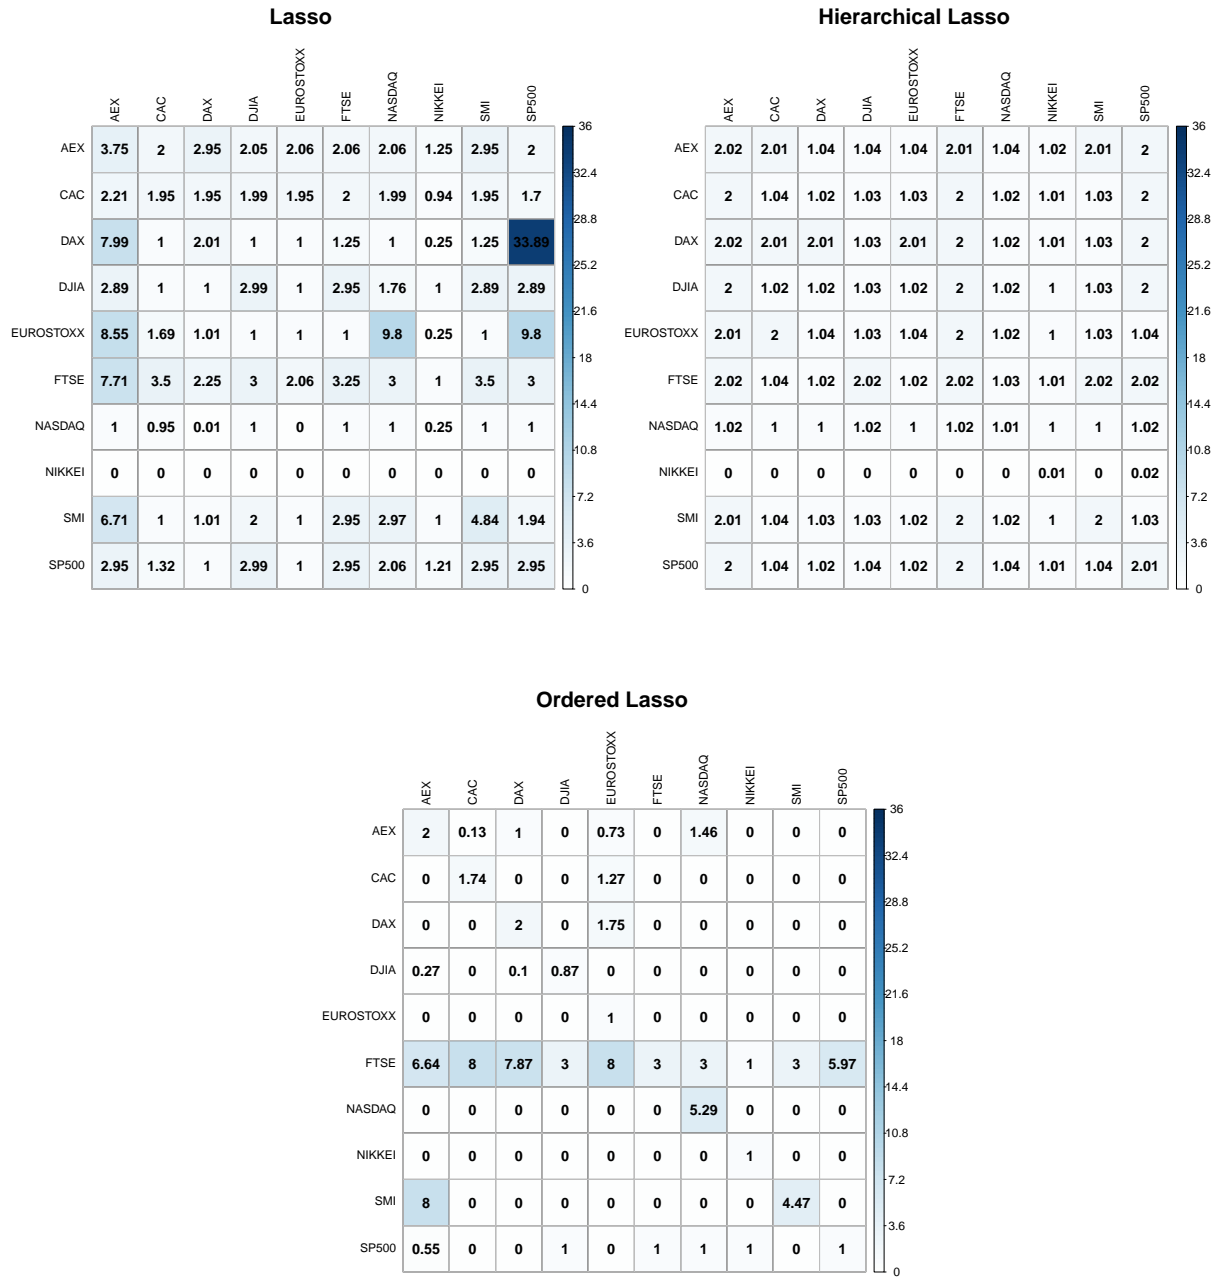


Figure 2: Lag matrix of estimated lag lengths. Each cell ij indicates the maximum number of non-zero lags of time series i (row) in equation j (column) of the VAR model estimated by the lasso (top left), hierarchical lasso (top right), or ordered lasso (bottom). The darker the cell, the higher the estimated lag length.

the number of time points, whereas the latter is not. For this reason, the maximum possible lag length p_{LS} estimable by least squares might be lower than p depending on the model (AR or VAR) and the number of time points under consideration. For instance, if the number of time points is $\lfloor T/2 \rfloor = 98$ (i.e. our smallest estimation window), the maximum possible lag length for the VAR estimable by least squares is $p = 9$ since then the number of parameters $p \times q = 9 \times 10 = 90$ per equation is still below the number of time points. Forecasts without combination are obtained by selecting the sparsity parameter selected based on the BIC for the lasso estimators, and with the lag length selected by the BIC for the least squares.

To evaluate the out-of-sample forecast accuracy of each method, we use an expanding window approach. We consider different values of the forecast horizon $h \in \{1, 2, 3, 6\}$. At each time point $t = \lfloor T/2 \rfloor, \dots, T-h$, we use the forecast methods to estimate the AR (in the univariate case) or the VAR model (in the multivariate case) with maximum lag length $p = 36$ months. As such, h -step ahead direct forecasts $\log \widehat{RV}_{t+h}^{(i)}$ and corresponding forecast errors $e_{t+h}^{(i)} = \log RV_{t+h}^{(i)} - \log \widehat{RV}_{t+h}^{(i)}$ are obtained for each stock market index $1 \leq i \leq 10$. The forecast performance of each stock market index i is then measured by its Absolute Forecast Error

$$AFE_h^{(i)} = \frac{1}{T-h-\lfloor T/2 \rfloor+1} \sum_{t=\lfloor T/2 \rfloor}^{T-h} |e_{t+h}^{(i)}|.$$

The overall forecast performance, averaged over the ten stock markets, is measured by the Mean Absolute Forecast Error

$$MAFE_h = \frac{1}{10} \sum_{i=1}^{10} \left(\frac{1}{T-h-\lfloor T/2 \rfloor+1} \sum_{t=\lfloor T/2 \rfloor}^{T-h} \frac{|e_{t+h}^{(i)}|}{\widehat{\sigma}_{1:t}^{(i)}} \right), \quad (8)$$

where we divide the forecast errors by $\widehat{\sigma}_{1:t}^{(i)}$, the standard deviation of the log-realized variance of stock market index i computed over the time points $1, \dots, t$, to adjust for the fact that the different series have different volatilities and predictabilities (see e.g. Carriero et al., 2011).

The AFEs for each stock market index are in Table 3 for the univariate forecast models, Table 4 for the multivariate forecast models. Table 5 summarizes the results by computing the average $MAFE_h$ as shown in equation (8) concentrating on the three dimensions Estimator, Forecast Combination, and Forecast Model.

Estimators Table 5 (top panel) shows that for all horizons, the ordered lasso attains the lowest, and thus best, MAFE. The least squares performs second best. Among the lasso-based estimators, only the ordered lasso performs better than the least squares. Hence, the largest improvements in forecast accuracy can be obtained by accounting for the dynamic nature of the AR model and imposing monotonicity constraints.

Forecast combination versus no forecast combination Table 5 (bottom left panel) shows that the forecast combined approaches produce slightly lower MAFEs than the non-forecast combined ones, and this for all horizons. Hence, the gains of protecting against a wrongly chosen sparsity parameter are, in the

Table 3: Univariate Forecast Models.

Stock market index	Horizon	Least Squares		Lasso		Hierarchical Lasso		Ordered Lasso	
		noFC	FC	noFC	FC	noFC	FC	noFC	FC
AEX	$h = 1$	0.460	0.458	0.550	0.546	0.581	0.581	0.457	0.458
	$h = 2$	0.567	0.553	0.599	0.599	0.626	0.626	0.554	0.553
	$h = 3$	0.599	0.598	0.631	0.628	0.675	0.675	0.619	0.620
	$h = 6$	0.670	0.671	0.742	0.743	0.717	0.717	0.690	0.692
CAC	$h = 1$	0.494	0.485	0.550	0.549	0.563	0.563	0.483	0.484
	$h = 2$	0.587	0.580	0.608	0.610	0.612	0.612	0.578	0.576
	$h = 3$	0.634	0.634	0.643	0.642	0.630	0.630	0.607	0.611
	$h = 6$	0.657	0.656	0.698	0.699	0.683	0.682	0.652	0.655
DAX	$h = 1$	0.477	0.472	0.522	0.523	0.549	0.549	0.468	0.469
	$h = 2$	0.531	0.531	0.577	0.577	0.585	0.585	0.538	0.538
	$h = 3$	0.599	0.598	0.649	0.648	0.621	0.621	0.590	0.596
	$h = 6$	0.650	0.650	0.678	0.844	0.681	0.681	0.665	0.662
DJIA	$h = 1$	0.567	0.567	0.591	0.593	0.637	0.637	0.548	0.548
	$h = 2$	0.650	0.646	0.670	0.669	0.706	0.706	0.638	0.640
	$h = 3$	0.667	0.665	0.688	0.688	0.725	0.725	0.675	0.669
	$h = 6$	0.733	0.733	0.785	0.786	0.799	0.799	0.753	0.755
EUROSTOXX	$h = 1$	0.477	0.467	0.514	0.517	0.527	0.527	0.469	0.468
	$h = 2$	0.552	0.556	0.576	0.573	0.563	0.563	0.544	0.540
	$h = 3$	0.608	0.608	0.593	0.588	0.595	0.595	0.588	0.588
	$h = 6$	0.631	0.629	0.658	0.951	0.658	0.658	0.632	0.630
FTSE	$h = 1$	0.463	0.458	0.526	0.529	0.550	0.550	0.462	0.462
	$h = 2$	0.565	0.559	0.584	0.587	0.595	0.595	0.551	0.550
	$h = 3$	0.615	0.613	0.628	0.619	0.627	0.627	0.606	0.606
	$h = 6$	0.677	0.677	0.699	0.702	0.690	0.690	0.667	0.669
NASDAQ	$h = 1$	0.521	0.505	0.659	0.658	0.710	0.710	0.498	0.499
	$h = 2$	0.611	0.608	0.751	0.736	0.756	0.756	0.601	0.604
	$h = 3$	0.637	0.634	0.703	0.705	0.776	0.776	0.640	0.644
	$h = 6$	0.671	0.672	0.841	0.843	0.832	0.832	0.718	0.723
NIKKEI	$h = 1$	0.625	0.632	0.732	0.729	0.659	0.659	0.638	0.637
	$h = 2$	0.719	0.721	0.713	0.705	0.680	0.680	0.684	0.692
	$h = 3$	0.775	0.778	0.767	0.770	0.727	0.727	0.774	0.766
	$h = 6$	0.762	0.762	0.722	0.723	0.738	0.738	0.741	0.736
SMI	$h = 1$	0.543	0.527	0.566	0.570	0.596	0.596	0.530	0.529
	$h = 2$	0.617	0.617	0.642	0.638	0.633	0.633	0.612	0.601
	$h = 3$	0.646	0.643	0.659	0.660	0.672	0.673	0.633	0.630
	$h = 6$	0.711	0.712	0.717	0.778	0.748	0.749	0.711	0.711
SP500	$h = 1$	0.544	0.535	0.602	0.602	0.654	0.654	0.530	0.530
	$h = 2$	0.661	0.652	0.699	0.699	0.723	0.723	0.639	0.643
	$h = 3$	0.686	0.685	0.718	0.721	0.734	0.734	0.682	0.681
	$h = 6$	0.741	0.741	0.850	0.845	0.803	0.803	0.763	0.765

Notes: Absolute Forecast Errors for the ten stock market indices, the four forecast horizons, and eight forecast methods. “no FC”: result from the optimal model selection by BIC. “FC”: result from Forecast Combination approach. The sample used is monthly data from January, 2000, to April, 2016, totaling 196 observations.

Table 4: Multivariate Forecast Models.

Stock market index	Horizon	Least Squares		Lasso		Hierarchical Lasso		Ordered Lasso	
		noFC	FC	noFC	FC	noFC	FC	noFC	FC
AEX	$h = 1$	0.472	0.472	0.624	0.622	0.653	0.651	0.468	0.466
	$h = 2$	0.591	0.591	0.650	0.649	0.674	0.672	0.563	0.562
	$h = 3$	0.632	0.632	0.637	0.637	0.678	0.677	0.599	0.598
	$h = 6$	0.700	0.700	0.690	0.690	0.709	0.710	0.678	0.678
CAC	$h = 1$	0.504	0.504	0.597	0.596	0.620	0.618	0.488	0.487
	$h = 2$	0.627	0.627	0.620	0.618	0.635	0.634	0.579	0.579
	$h = 3$	0.666	0.666	0.605	0.605	0.637	0.636	0.598	0.597
	$h = 6$	0.739	0.739	0.659	0.660	0.673	0.673	0.643	0.643
DAX	$h = 1$	0.488	0.488	0.606	0.604	0.625	0.623	0.478	0.477
	$h = 2$	0.599	0.599	0.621	0.620	0.634	0.633	0.533	0.533
	$h = 3$	0.628	0.628	0.615	0.615	0.638	0.637	0.586	0.585
	$h = 6$	0.677	0.677	0.670	0.670	0.673	0.674	0.652	0.652
DJIA	$h = 1$	0.573	0.573	0.688	0.685	0.715	0.713	0.534	0.533
	$h = 2$	0.665	0.665	0.706	0.705	0.732	0.730	0.624	0.623
	$h = 3$	0.645	0.645	0.687	0.687	0.735	0.734	0.666	0.665
	$h = 6$	0.724	0.724	0.758	0.758	0.787	0.787	0.763	0.763
EUROSTOXX	$h = 1$	0.482	0.482	0.569	0.567	0.590	0.589	0.475	0.475
	$h = 2$	0.607	0.607	0.588	0.588	0.600	0.599	0.550	0.551
	$h = 3$	0.668	0.668	0.575	0.575	0.603	0.602	0.575	0.575
	$h = 6$	0.695	0.695	0.621	0.622	0.633	0.633	0.602	0.602
FTSE	$h = 1$	0.463	0.463	0.561	0.559	0.597	0.594	0.467	0.465
	$h = 2$	0.628	0.628	0.591	0.590	0.614	0.612	0.546	0.546
	$h = 3$	0.648	0.648	0.612	0.612	0.629	0.628	0.598	0.598
	$h = 6$	0.728	0.728	0.658	0.658	0.679	0.680	0.666	0.666
NASDAQ	$h = 1$	0.526	0.526	0.731	0.727	0.740	0.739	0.534	0.530
	$h = 2$	0.628	0.628	0.726	0.724	0.754	0.753	0.645	0.643
	$h = 3$	0.613	0.613	0.719	0.719	0.755	0.755	0.673	0.673
	$h = 6$	0.656	0.656	0.768	0.769	0.783	0.783	0.748	0.748
NIKKEI	$h = 1$	0.673	0.673	0.744	0.744	0.739	0.738	0.699	0.693
	$h = 2$	0.751	0.751	0.722	0.722	0.717	0.716	0.731	0.732
	$h = 3$	0.783	0.783	0.717	0.717	0.725	0.724	0.739	0.739
	$h = 6$	0.887	0.887	0.726	0.726	0.728	0.728	0.723	0.723
SMI	$h = 1$	0.543	0.543	0.652	0.650	0.670	0.668	0.525	0.528
	$h = 2$	0.661	0.661	0.654	0.653	0.683	0.682	0.614	0.613
	$h = 3$	0.701	0.701	0.658	0.659	0.692	0.692	0.645	0.644
	$h = 6$	0.751	0.751	0.716	0.716	0.727	0.726	0.716	0.717
SP500	$h = 1$	0.547	0.547	0.692	0.691	0.720	0.718	0.537	0.536
	$h = 2$	0.675	0.675	0.725	0.723	0.742	0.741	0.650	0.650
	$h = 3$	0.678	0.678	0.703	0.704	0.744	0.743	0.684	0.684
	$h = 6$	0.739	0.739	0.757	0.757	0.790	0.790	0.769	0.769

Notes: Absolute Forecast Errors for the ten stock market indices, the four forecast horizons, and eight forecast methods. “no FC”: result from the optimal model selection by BIC. “FC”: result from Forecast Combination approach. The sample used is monthly data from January, 2000, to April, 2016, totaling 196 observations.

Table 5: Mean Absolute Forecast Errors (MAFE_h) for the different methods.

Horizon	Estimators			
	Least Squares	Lasso	Hierarchical Lasso	Ordered Lasso
$h = 1$	0.520	0.613	0.634	0.514
$h = 2$	0.624	0.650	0.663	0.599
$h = 3$	0.656	0.660	0.681	0.639
$h = 6$	0.710	0.734	0.727	0.698

Horizon	Forecast Combination		Forecast Model	
	No	Yes	Univariate	Multivariate
$h = 1$	0.571	0.570	0.552	0.590
$h = 2$	0.634	0.633	0.622	0.646
$h = 3$	0.659	0.659	0.658	0.660
$h = 6$	0.720	0.714	0.722	0.712

Notes: Regarding Estimators: we report the average over the univariate and multivariate forecast models, the forecast combined and non-forecast combined methods. Regarding Forecast Combination: we report the average over the four estimators, the univariate and multivariate forecast models. Regarding Forecast Model: we report the average over the four estimators, the forecast combined and non-forecast combined methods. The total sample is monthly data from January, 2000, to April, 2016, totaling 196 observations.

majority of considered cases, small.

Univariate versus multivariate forecast models Table 5 (bottom right panel) shows that for the short forecast horizons $h = 1$ and $h = 2$, the univariate forecast models outperform the multivariate forecast models. For $h = 3$, the MAFE is about equally good for both. For the longer forecast horizon $h = 6$, the multivariate forecast models outperform the univariate ones. In the long-run, the stock market indices move together, as can be seen from Figure 1. While it becomes more difficult to obtain accurate forecasts for longer horizons (i.e. MAFEs increase for longer forecast horizons), incorporating spillover effects between the different stock market indices helps to improve forecast accuracy for longer forecast horizons.

In sum, the MAFE measures indicate that (i) the ordered lasso is the best performing estimator; (ii) forecast combinations slightly improve the forecast performance of their non-forecast combined alternatives, and (iii) multivariate forecast models only outperform the univariate ones for longer forecast horizons. We further elaborate on these findings by inspecting the AFEs for each stock market index separately, see Table 3 and 4, and by computing the Model Confidence Set of Hansen et al. (2011), see Table 6.

AFE for each stock market index The ordered lasso also performs best when looking at the AFEs of each stock market index separately. For the univariate forecast models, it gives the best AFE in 24 out of 40 cases (60%), thereby clearly outperforming the least squares (best in 35% of the cases), the hierarchical lasso (best in 5%) and the lasso (best in 0%). For the multivariate forecast models, the ordered lasso gives

Table 6: Model Confidence set over the 16 forecast methods for the different forecast horizons.

Horizon	Univariate Forecast Models				Multivariate Forecast Models				Set size
	least squares	lasso	hierarchical lasso	ordered lasso	least squares	lasso	hierarchical lasso	ordered lasso	
$h = 1$	FC			no FC; FC					3
$h = 2$	FC			no FC; FC			no FC; FC		5
$h = 3$	no FC; FC			no FC; FC			no FC; FC		6
$h = 6$	no FC; FC			no FC; FC	no FC; FC		no FC; FC		8

Notes: The label “no FC” (“FC”) indicates that the non-Forecast Combined (Forecast Combined) approach of the corresponding estimator and forecast model is included in the MCS. The last column indicates the number of forecast methods in the MCS. The total sample is monthly data from January, 2000, to April, 2016, totaling 196 observations.

the best AFE in 25 out of 40 cases (63%), and performs more than twice as good as the second best which is the least squares (i.e. best in 25% of the cases).

In the majority of cases, the forecast combinations perform as good as their non-forecast combined alternatives. In some cases, their advantage is outspoken. In particular, the forecast combined least squares for the univariate forecast model gives the lowest AFE in 70% of the cases versus 30% for its non-forecast combined alternative. For the multivariate forecast model, the forecast combined ordered lasso attains the best AFE in 78% of the cases, the forecast combined hierarchical lasso in 85% of the cases.

Finally, comparing the best univariate to the best multivariate forecasts for each stock market index separately reveals that the univariate forecast models attain the best AFE for 8 out of 10 indices, both for horizon $h = 1$ and $h = 2$. In contrast, for $h = 3$ and $h = 6$, the multivariate forecast models attain the best AFE for respectively 8 out of 10 and 7 out of 10 cases. Hence, similar conclusions hold as for the MAFEs measures that are averaged over the ten stock market indices.

Model Confidence Set We use the Model Confidence Set (MCS) procedure to separate the best forecast methods with equal predictive ability from the others. Using the `MCSprocedure` function in R, we obtain a MCS that contains the best method with 75% confidence level. We use the TR test statistic computed from 5000 bootstrapped samples, see Catania and Bernardi (2014). The MCS for each forecast horizon is given in Table 6.

For $h = 1$, the MCS includes three (out of the 16) forecast methods: the forecast combined and non-forecast combined ordered lasso, and the forecast combined least squares, all for the univariate forecast model. Hence, univariate forecast models perform best for short horizons. For $h = 2$, the MCS includes five forecast methods. The forecast combined and non-forecast combined ordered lasso for the multivariate forecast model are added to the MCS of horizon $h = 1$. For $h = 3$, the MCS includes six forecast methods. The non-forecast combined least squares for the univariate forecast model is added to the MCS of horizon $h = 2$. Finally, for $h = 6$, the MCS includes eight forecast models. The forecast combined and non-forecast combined lasso for the multivariate forecast model are added to the MCS of horizon $h = 3$.

In sum, the MCS procedure confirms the excellent performance of the ordered lasso. As the forecast

horizon increases, it becomes more difficult to obtain more accurate forecasts, and more forecast methods are included in the MCS. The MCS contains more multivariate forecast models for longer horizons than for shorter horizons.

5 Robustness checks

We investigate the robustness of our findings in Section 4 to (i) the choice of forecast combination weights, (ii) the length of the sparsity grid L , (iii) the maximum lag length p of the AR/VAR models, and (iv) the use of rolling versus expanding window forecast approach.

For all robustness checks, similar overall findings to the ones from Section 4 are that (i) the ordered lasso is the best performing estimator, (ii) forecast combinations perform slightly better than their non-forecast combined alternatives, (iii) univariate forecast models perform best for short horizons; multivariate forecast models for longer horizons. Below we focus on the new findings of each robustness check. Detailed results are available from the authors upon request.

Forecast combination weights As alternative to the BIC-based forecast combination weights of equation (7), we take: (i) equal weights, and (ii) MSE-based weights, with weights proportional to the inverse of the individual (in-sample) Mean Squared Estimation Errors (MSE). The highest weight is given to the model with the smallest MSE.

The ordered lasso retains the best overall forecast performance, regardless of the choice of forecast combination weights. It performs about equally good for all forecast combination weights. The other estimators are, in contrast, more sensitive to an appropriate choice of forecast combination weights. BIC-based weights give the best forecast performance.

Length of sparsity grid In Section 4, we report the results for the length of the sparsity grid $L = 20$ since our numerical experiments show that this choice gives the lowest overall MAFEs. The main findings remain the same when varying L . The relative performance of the multivariate forecast models to the univariate ones is most influenced by the length of the sparsity grid. For a shorter sparsity grid $L = 10$, the multivariate forecast models perform much better than the univariate ones. For $h = 6$, for instance, their average MAFE is 0.75 compared to 0.92 for the univariate ones, a reduction of 22%. Since our sparsity grid is logarithmic spaced, a shorter sparsity grid means there are more parsimonious than non-parsimonious models to choose from. Especially the more parsimonious VAR models give more accurate log-realized variance forecasts. For a longer sparsity grid $L = p = 36$, multivariate forecast models perform relatively worse. Now there are more non-parsimonious than parsimonious models to choose from. The non-parsimonious VAR models give, on average, less accurate log-realized variance forecasts.

Maximum lag length We both shorten the maximum lag length from $p = 36$ to $p = 24$ and extend it to $p = 48$. The results of all forecast methods are very stable when varying p . Only the performance of the lasso deteriorates when using a larger maximum lag length. The lasso tends to pick up some higher order lags, the more so when p increases. This introduces more noise in the forecast performance and, hence, leads to lower overall AFEs. Hence, the lasso should only be used as a forecast method for models with lower maximum lag lengths. The hierarchical lasso and ordered lasso, in contrast, do not suffer from this as they encourage low maximal lag lengths in their estimation procedure.

Rolling versus expanding window. We replace the expanding window with a rolling window of size $S = \lfloor T/2 \rfloor = 98$. Results are very similar. Only when further decreasing the window size S , the performance of the least squares model starts to suffer. The least squares model becomes imprecise once a lot of parameters need to be estimated relative to the sample size, i.e. S . Lasso-based estimators, in contrast, perform regularization to deal with this overparametrization issue, and, as a consequence, retain their good performance. Hence, if few time points are available to make forecasts, the use of lasso-based estimators compared to the least squares becomes more interesting.

6 Discussion

We use lasso-based forecast combination methods to obtain log-realized variance forecasts for ten different stock market indices. Consistently across stock market indices, our results indicate that the ordered lasso is the best performing estimator. The forecast combined ordered lasso slightly improves the forecast accuracy of the non-forecast combined ordered lasso by hedging against the risk of selecting a wrong sparsity parameter. While the ordered lasso for the univariate forecast model performs best for short forecast horizons, the ordered lasso for the multivariate forecast model performs best for longer forecast horizons.

When evaluating the forecast performance of the different forecast methods in Section 4, we compare the predicted log-realized variances to the actual ones over the period February 2008-April 2016. The Global Financial Crisis, i.e. February 2008-February 2009, falls into this period. We also compute the MAFE for this subperiod only and investigate whether the relative forecast performance of our forecast methods is different during this subperiod. As expected, the MAFEs of all forecast methods are higher compared to the ones from Table 5 since it is more difficult to obtain reliable log-realized variance forecasts during this more turbulent period. Nevertheless, the forecast combined ordered lasso remains the best performing forecast method. Multivariate forecast models still have an advantage over the univariate ones for the longer forecast horizons.

One of the objectives of this paper is to compare multivariate with univariate forecast models. To this end, we opted to work with monthly data because of the non-overlapping trading days between the ten stock market indices. Since realized variance forecast are often made at the daily frequency as well, we also compared the forecast performance of the univariate forecast models to the Heterogeneous AutoRegressive (HAR) model (see e.g. Corsi, 2009) when using daily data. In line with the results of Section 4, we find the forecast combined ordered lasso to be the best performing forecast method. The forecast combined ordered lasso is competitive to the HAR model but it does not outperform it.

Acknowledgments. The authors gratefully acknowledge financial support from the FWO (Research Foundation-Flanders, contract number 12M8217N).

References

- Andersen, T. G.; Bollerslev, T.; Diebold, F. X. and Labys, P. (2001), “The distribution of realized exchange rate volatility,” *Journal of the American Statistical Association*, 96(453), 42–55.
- Audrino, F.; Camponovo, L. and Roth, C. (2015), “Testing the lag structure of assets’ realized volatility dynamics,” *Economics Working Paper Series*, 1501.
- Audrino, F. and Knaus, S. (2016), “Lassoing the HAR model: A model selection perspective on realized volatility dynamics,” *Econometric Reviews*, 35(8-10), 1485–1521.
- Barndorff-Nielsen, O. and Shephard, N. (2002), “Econometric analysis of realized volatility and its use in estimating stochastic volatility models,” *Journal of the Royal Statistical Society: Series B*, 64(2), 253–280.
- Bien, J.; Taylor, J. and Tibshirani, R. (2013), “A lasso for hierarchical interaction,” *The Annals of Statistics*, 41(3), 1111–1141.
- Bien, J. and Witten, D. (2016), *Penalized estimation in complex models.*, In: Bühlmann, P. and Drineas, P. and Kane, M. and van der Laan, M. (Eds.) *Handbook of Big Data*, Chapman & Hall/CRC.
- Bollerslev, T. (1986), “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 31(3), 307–327.
- Callot, L. A. F.; Kock, A. B. and Medeiros, M. C. (2016), “Modeling and forecasting large realized covariance matrices and portfolio choice,” *Journal of Applied Econometrics*, DOI: 10.1002/jae.2512.
- Carriero, A.; Kapetanios, G. and Marcellino, M. (2011), “Forecasting large datasets with Bayesian reduced rank multivariate models,” *Journal of Applied Econometrics*, 26(5), 735–761.
- Catania, L. and Bernardi, M. (2014), *MCS: Model Confidence Set Procedure*, R package version 0.1.1.
- Corsi, F. (2009), “A simple approximate long-memory model of realized volatility,” *Journal of Financial Econometrics*, 7(2), 174–196.
- Diebold, F. X. and Yilmaz, K. (2015), *Financial and macroeconomics connectedness: A network approach to measurement and monitoring*, Oxford University Press, New York, US.
- Elliott, G. and Timmermann, A. (2016), *Economic Forecasting*, Princeton University Press, Princeton, New Jersey.
- Fan, J.; Lv, J. and Qi, L. (2011), “Sparse high-dimensional models in economics,” *Annual Review of Economics*, 3, 291–317.

- Friedman, J.; Hastie, T. and Tibshirani, R. (2010), “Regularization paths for generalized linear models via coordinate descent,” *Journal of statistical software*, 33(1), 1–22.
- Friedman, J.; Suo, X. and Tibshirani, R. (2015), *Package: orderedLasso*, R package version 1.7.
- Gelper, S.; Wilms, I. and Croux, C. (2016), “Identifying demand effects in a large network of product categories,” *Journal of Retailing*, 92(1), 25–39.
- Hansen, P. R.; Lunde, A. and Nason, J. M. (2011), “The model confidence set,” *Econometrica*, 79(2), 453–497.
- Jenatton, R.; Audibert, J. and Bach, F. (2011), “Structured variable selection with sparsity-inducing norms,” *Journal of Machine Learning Research*, 12, 2777–2824.
- Knight, J. and Satchell, S. E. (2007), *Forecasting volatility in the financial markets.*, Quantitative Finance Series, Butterworth-Heinemann.
- Koop, G. (2003), *Bayesian Econometrics*, Wiley-Interscience.
- Marcellino, M.; H., S. J. and W., W. M. (2006), “A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series,” *Journal of Applied Econometrics*, 135, 499–526.
- Nicholson, W. B.; Bien, J. and Matteson, D. S. (2016a), “High dimensional forecasting via interpretable vector autoregression,” *arXiv*, 1412.5250v2.
- Nicholson, W. B.; Matteson, D. S. and Bien, J. (2016b), *Package: BigVAR*, R package version 1.0.1.
- Suo, X. and Tibshirani, R. (2015), “An ordered lasso and sparse time-lagged regression,” *Technometrics*, DOI: 10.1080/00401706.2015.1079245.
- Tibshirani, R. (1996), “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society Series B*, 58(1), 267–288.

FACULTY OF ECONOMICS AND BUSINESS
Naamsestraat 69 bus 3500
3000 LEUVEN, BELGIË
tel. + 32 16 32 66 12
fax + 32 16 32 67 91
info@econ.kuleuven.be
www.econ.kuleuven.be

