DIGITALES ARCHIV

ZBW – Leibniz-Informationszentrum Wirtschaft
ZBW – Leibniz Information Centre for Economics

Gupta, Rangan; Pierdzioch, Christian

Book

Multi-task forecasting of the realized volatilities of agricultural commodity prices

Provided in Cooperation with:
University of Pretoria

Reference: Gupta, Rangan/Pierdzioch, Christian (2024). Multi-task forecasting of the realized volatilities of agricultural commodity prices. Pretoria, South Africa : Department of Economics, University of Pretoria.
https://www.up.ac.za/media/shared/61/WP/wp_2024_23.zp251849.pdf.

This Version is available at:
http://hdl.handle.net/11159/654453

Kontakt/Contact
ZBW – Leibniz-Informationszentrum Wirtschaft/Leibniz Information Centre for Economics
Düsternbrooker Weg 120
24105 Kiel (Germany)
E-Mail: *rights[at]zbw.eu*
https://www.zbw.eu/

Standard-Nutzungsbedingungen:
Dieses Dokument darf zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden. Sie dürfen dieses Dokument nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen. Sofern für das Dokument eine Open-Content-Lizenz verwendet wurde, so gelten abweichend von diesen Nutzungsbedingungen die in der Lizenz gewährten Nutzungsrechte. Alle auf diesem Vorblatt angegebenen Informationen einschließlich der Rechteinformationen (z.B. Nennung einer Creative Commons Lizenz) wurden automatisch generiert und müssen durch Nutzer:innen vor einer Nachnutzung sorgfältig überprüft werden. Die Lizenzangaben stammen aus Publikationsmetadaten und können Fehler oder Ungenauigkeiten enthalten.

Terms of use:
This document may be saved and copied for your personal and scholarly purposes. You are not to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public. If the document is made available under a Creative Commons Licence you may exercise further usage rights as specified in the licence. All information provided on this publication cover sheet, including copyright details (e.g. indication of a Creative Commons license), was automatically generated and must be carefully reviewed by users prior to reuse. The license information is derived from publication metadata and may contain errors or inaccuracies.

https://savearchive.zbw.eu/termsofuse

ZBW Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Mitglied der Leibniz-Gemeinschaft

University of Pretoria
Department of Economics Working Paper Series

Multi-Task Forecasting of the Realized Volatilities of Agricultural Commodity Prices

Rangan Gupta
University of Pretoria
Christian Pierdzioch
Helmut Schmidt University
Working Paper: 2024-23
June 2024

Department of Economics
University of Pretoria
0002, Pretoria
South Africa
Tel: +27 12 420 2413

Multi-Task Forecasting of the Realized Volatilities of Agricultural Commodity Prices

Submission: May 2024

Rangan Gupta[*], Christian Pierdzioch[†]

Abstract

Motivated by the comovement of realized volatilities (RVs) of agricultural commodity prices, we study whether multi-task forecasting algorithms improve the accuracy of out-of-sample forecasts of 15 agricultural commodities during the sample period from July 2015 to April 2023. We consider alternative multi-task stacking algorithms and variants of the multivariate Lasso estimator. We find evidence of in-sample predictability, but hardly evidence that multi-task forecasting improves out-of-sample forecasts relative to a classic univariate heterogeneous autoregressive (HAR) RV model. We also study an extended model that features the RVs of energy commodities and precious metals.

JEL Classifications: C22; C32; C53; Q11

Keywords: Agricultural commodities; Realized volatility; Multi-task forecasting

Conflicts of interest: The authors declare no conflict of interest.

Funding statement: The authors declare that they did not receive any funding for this research.

[*]Department of Economics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa. Email address: rangan.gupta@up.ac.za.

[†]Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany; Email address: macroeconomics@hsu-hh.de.

1 Introduction

Quite a number of empirical studies have been undertaken to shed light on the connectedness of volatility across agricultural commodities (see, for example, Trujillo-Barrera et al. (2012), Beckmann and Czudaj (2014), Gardebroek et al. (2016), Bonato (2019), Luo and Chen (2019), Cagli et al. (2023), Luo et al. (2023)), with Marfatia et al. (2022) highlighting that accounting for co-volatility of Chinese futures of five (corn, cotton, palm, wheat, and soybeans) agricultural commodities improves the accuracy of volatility forecasts in particular for corn, cotton, and wheat individually. We contribute to this area of research by exploring whether stacking algorithms that have been developed in the recent bioinformatics literature help to improve the accuracy of out-of-sample forecasts of the intraday data-based realized volatility (RV) of 15 important agricultural commodities during the daily sample period of July, 2015 to April, 2023.

An important advantage of using RV for our empirical analyses derives from the rich information contained in intraday data, besides being a consistent and asymptotically unbiased estimator of volatility (Andersen and Bollerslev, 1998; McAleer and Medeiros, 2008; Chatziantoniou et al., 2021). In addition, RV is an observable and unconditional metric of "volatility". This, in turn, is unlike the latent processes underlying the class of Generalized Autoregressive Conditional Heteroskedasticity (GARCH) and Stochastic Volatility (SV) models that have been widely used for predicting agricultural commodity price volatility (see, Degiannakis et al. (2022) and Luo et al. (2022) for reviews of this extensive literature). Moreover, the dynamics of RV can be easily modeled by means of the heterogeneous autoregressive (HAR-) RV model (Corsi, 2009). The HAR-RV model has been extensively studied in research on realized volatility, including that of agricultural commodities (as reviewed in Bonato et al. (2022, 2024, forthcoming)), because it is able to capture long-memory and multi-scaling properties of realized volatility, as has been reported by Gil-Alana et al. (2012), and Živkov et al. (2019, 2022). Because the HAR-RV model employs RVs at different time resolutions to model and predict RV, it can be interpreted as a simple empirical representation of the heterogeneous market hypothesis (HMH; Müller et al., 1997), which stipulates that asset markets (in our case, markets for agricultural commodities) are populated by various types of market participants such as, investors, speculators, and traders, who, in turn, in turn, vary in their sensitivity to information flows at high and low frequencies.

Another advantage of the HAR-RV model is that it can easily be adapted to a multi-task forecasting setting, i.e., a setting where a forecaster seeks to forecast not only the RV of a single agricultural commodity, but the RVs of several agricultural commodities simultaneously. One possibility to address such a multi-task forecasting problem is to consider as a modeling framework one of the multivariate HAR-RV models with heteroskedastic error structures, as has been studied by, for instance, Bubák et al. (2011), de Nicola et al. (2016), Luo and Chen (2019, 2020a, b), Marfatia et al. (2022), Luo et al. (forthcoming). The focus of many studies in this area, however, have been on modeling and forecasting covolatilities (see, for example, Asai and McAleer (2017), Čech and Baruník (2017), Asai et al. (2019, 2020), Luo and Chen (2019, 2020)). Moreover, applications of the HAR-RV-cum-heteroskedastic-errors models are often restricted to settings where the number of RVs to be analyzed is relatively small, as was the case in

Luo and Chen (2019) and Marfatia et al. (2022) involving, respectively, seven and five agricultural commodities. This is due to the fact that, in a multivariate setting, the number of parameters to be estimated rapidly increases in the dimension of the model, unless a researcher is willing to impose restrictions on parameters and/or functional forms so as to obtain a parsimonious representation of the heteroskedastic error structures.

In our case, the dataset comprises RVs of 15 agricultural commodities (and, in an extended model, in addition the RVs of three important energy commodities and the RVs of five precious metals), and so we use various computationally efficient multi-task stacking algorithms that have been proposed in the recent bioinformatics literature (along with a multivariate shrinkage estimator) to re-examine the out-of-sample predictability of the RVs of the commodities in our sample (for a recent application of stacking in a univariate forecasting exercise of stock returns, see Zhao and Cheng (2022)). Also, we focus on direct spillovers among the RVs as captured by a multi-task HAR-RV model, and do not consider the issue of forecasting co-volatility, which requires imposing further structure on the residuals. The multi-task stacking algorithms are easy to implement even when the dimension of the model is large. Moreover, they make it possible to employ and combine alternative popular machine-learning algorithms that make it possible to estimate a multi-task HAR-RV model in a data-driven way, that is, without imposing a priori any specific structure that restricts the spillover dynamics across the RVs. Finally, the multi-task stacking algorithms can be set up in a way such that the resulting statistical model captures potential nonlinear structures in the data, an issue that certainly deserves special attention in the wake of the type of sudden outbreaks and clustering of volatility typical of financial markets, and of markets for agricultural commodities as well. In the process, our paper adds to the growing literature on modeling and predicting RV of agricultural commodities by investigating the role of volatility spillovers, where researchers in this literature have thus far otherwise relied on realized moments (such as realized kurtosis and realized jumps) and various other predictors that relate, for example, to the state of financial and other (non-agricultural) commodity markets, investor sentiment, climate-change-related risks, and infectious disease-related uncertainty (see, for example, Tian et al. (2017a, b), Yang et al. (2017), Degiannakis et al. (2022), Luo et al., (2022), Marfatia et al. (2022), Shiba et al. (2022), Bonato et al. (2022, 2024, forthcoming)), Gupta and Pierdzioch (2023), Luo and Zhang (2024)).

Agricultural commodities have become increasingly financialized (Bahloul et al., 2018; Aït-Youcef, 2019; Ji et al., 2020). The process has caused institutional investors to increase their holdings in agricultural commodities relative to traditional assets. Naturally, besides the academic value of our work, accurate forecasts of the volatility of agricultural commodity prices are of key importance for investors, because volatility is a core input in investment and portfolio allocation decisions, risk management, derivatives pricing, and assessments of hedging performance (Poon and Granger, 2003; Rapach et al., 2008). In addition, agricultural commodities comprise a large proportion of household consumption spending, implying that price volatility in agricultural commodities markets is likely to have substantial consequences for food security, especially as far as the economically vulnerable groups of the population are concerned (FAO, 2010, 2011; Ordu, et al., 2018). Hence, from a policy perspective, it is important to produce accurate high-frequency forecasts of agricultural commodity prices volatility so that policies can be discussed and implemented in a timely manner to protect those vulnerable groups in particular from large and adverse food price fluctuations (Greb and Prakash, 2015; 2017).

In order to get to our empirical findings, we organize the rest of the paper as follows. In Section 2, we provide a description of the data we use in our study, while we outline in Section 3 our methods. In Section 4, we present our empirical results. In Section 5, we conclude.

2 Data

We use in our empirical analysis data on RV of 15 agricultural commodities. The data is available publicly for download from the internet page of Professor Dacheng Xiu.[1] The data are based on (Globex) futures data for the following 15 agricultural commodities: soybean oil futures (BO), cocoa futures (CC), corn futures (C), cotton no. 2 futures (CT), feeder cattle futures (FC), coffee C futures (KC), lumber futures (LB), live cattle futures (LC), lean hogs futures (LH), orange juice futures (OJ), oats futures (O), sugar #11 futures (SB), soybean meal futures (SM) , soybean futures (S), and wheat futures CBOT (W). After matching the data by date, the matched dataset starts on 27/07/2015 and ende on 28/04/2023. We plot the RV data of the agricultural commodities in Figure 1. The RVs display a discernible heterogeneity across the agricultural commodities, and they also

[1] Internet address: https://dachxiu.chicagobooth.edu/#risklab. Data downloaded in May 2024.

exhibit the type of clusters and sudden outbursts characteristic of many financial market volatilities.

– Figure 1 about here. –

In order to obtain a first glimpse of the comovement of the RVs, we plot their full-sample contemporaneous correlation matrix in Figure 2. The contemporaneous correlations vary from weakly negative to strongly positive, where the positive correlations are mainly collected in the lower part of the matrix. For example, we observe strong positive contemporaneous correlations between C and S, S and SM, and LC and FC, among others. While the full-sample contemporaneous correlations shed light on an important feature of the data, one should bear in mind that the correlations do not inform about the question whether the comovement of the RVs can be exploited in a multi-task out-of-sample forecasting exercise to improve predictive accuracy at various forecast horizons.

– Figure 2 about here. –

3 Methods

3.1 Forecasting Models

We frame our empirical analysis in terms of the popular HAR-RV model developed by Corsi (2009). This model can be specified by the following equation:

$$RV_{t+h} = \beta_0 + \beta_1 RV_t + \beta_2 RV_{IF,t} + \beta_3 RV_{LF,t} + u_{t+h}, \tag{1}$$

which we estimate by the ordinary-least-squares (OLS) technique, $\beta_j, j = 0,..,3$ are the coefficients to be estimated, u_{t+h} denotes a disturbance term, and RV_{t+h} denotes the average realized volatility over the forecast horizon, h. We study in our empirical research a short, two intermediate, and a long forecast horizon. To this end, we specify $h = 1, 5, 10, 22$. The predictors are the daily realized volatility, RV_t, the intermediate-frequency (IF) realized volatility, $RV_{t,IF}$, and the low-frequency (LF) realized volatility, $RV_{t,LF}$. We define the IF realized volatility as the average realized volatility from period $t-5$ to period $t-1$, and the LF realized volatility as the average realized volatility from period $t-22$ to period $t-1$, as computed using the matched data.

We emphasize that, in order to avoid non-negativity constraints and to bring the data closer to normality, we use the natural logarithm of RV to estimate the HAR-RV model (and its extension to the HAR-RV-S model, which accounts for potential spillover effects. We evaluate the resulting forecasts, however, in terms of the anti-log of RV.

The variant of the HAR-RV model that accounts for spillover effects, the HAR-RV-S model, is given by the following equation:

$$RV_{t+h,i} = \beta_0+\beta_1 RV_{t,i}+\beta_2 RV_{IF,t,i}+\beta_3 RV_{LF,t,i}+\sum_{j\neq i} (\beta_{4,j} RV_{t,j} + \beta_{5,j} RV_{IF,t,j} + \beta_{6,j} RV_{LF,t,j})+u_{t+h}, \tag{2}$$

where i is the agricultural commodity being studied, and the index j denotes the other agricultural commodities, Hence, we obtain the HAR-RV-S model by adding the daily, intermediate-frequency, and low-frequency realized volatilities of the other agricultural commodities and, thereby, account for potential spillover effects at different time resolutions. We emphasize that the HAR-RV-S model captures direct spillover effects among the RVs, not the dynamics of co-volatilities.

We use the R language and environment for statistical computing (R Core Ream 2023; R version 4.3.1) to estimate our forecasting models and to compute all other results that we lay out in this research. We estimate the forecasting models either by means of a recursively expanding estimation window or by means of a rolling-estimation window. We use 50% of the data to initialize the recursive estimations. Similarly, we use 50% of the data to define the length of a rolling-estimation window. Finally, we use the root-mean-squared forecasting error (RMSFE) and the mean-absolute forecasting error (MAFE) to evaluate the out-of-sample performance of the forecasting models, where we compute the ratio of the RMSFE (MAFE) of the HAR-RV-S model and the HAR-RV model to alleviate the interpretation of our empirical results. Hence, a RMSFE (MAFE) ratio smaller than unity implies that the HAR-RV-S model outperforms the HAR-RV model, while a ratio larger than unity signals that the HAR-RV model is the better forecasting model.

3.2 Stacking Algorithms

Given that our sample comprises 15 agricultural commodities, and we have to consider (leaving the intercept term apart) in total $15 \times 3 = 45$ predictors, we use computationally efficient stacking algorithms to estimate the forecasting model given in Equation (2).

The first stacking algorithm that we use (we call this algorithm henceforth the baseline stacking algorithm) has been studied recently by Rauschenberger and Glaab (2021). This baseline stacking algorithm requires that we treat the forecasting model given in Equation (2) as a base learner. Accordingly, we estimate 15 base learners, one for every agricultural commodity. Given the large number of parameters to be estimated, we estimate the base learners either by means of the Lasso estimator, as an elastic net, or by means of the Ridge regression estimator (see Tibshirani (1996), Zou and Hastie (2005)), where we choose the corresponding shrinkage parameter using 10-fold cross validation (CV). We use the CV-based out-of-fold predictions from the base learners to construct a matrix, $\hat{H}^{CV}$, with 15 columns, one for every agricultural commodity. Finally, we construct a meta learner by estimating a regression model, one for every agricultural commodity, of RV_{t+h}, on all predictors in $\hat{H}^{CV}$. Hence, the baseline stacking algorithm implies that the second-stage meta learners extract the information embedded in the predictors of the base learners in a way such that the forecast of RV_{t+h} combines in a linear way the first-stage estimated effects on all RVs. We use the shrinkage estimator that we apply to the base learners to estimate the meta learners. We use the R add-on package "joinet" (Rauschenberger and Glaab 2021) to implement the baseline stacking algorithm.

In addition, we use a modified stacking algorithm that has been proposed recently by Xing et al. (2020). Specifically, we use their residual stacking algorithm and the corresponding R add-on "MTPS" package. The modified stacking algorithm requires that we fit base learners in the first stage and compute the resulting fitted values of the RVs. One then models the residuals one obtains for agricultural commodity k using the fitted first-stage RVs (excluding the one for k) and obtains a meta learner using the first-stage base learner plus the fitted residual function. As a result, one can combine, for example, a first-stage Lasso estimator with a another Lasso estimator or, in case one suspects that the data feature nonlinear patterns that are worthwhile studying, regression trees (Breiman 1984) to obtain a meta learner. We call the latter a Lasso-RF model because a regression tree represents a general (rather than a linear) residual function.

4 Empirical Results

4.1 Full-Sample Results

We start the discussion of our empirical results by eyeballing the heatmaps we plot in Figure 3, which show the full-sample coefficients of the HAR-RV-S model for the four different forecasting horizons. The results are based on the Lasso version of the baseline stacking algorithm. The upper left heatmap shows that at the short forecasting horizon, $h = 1$, the coefficients of the classic HAR-RV model (that is, the diagonal cells of the map) dominate the scenery. The colors of most of the off-diagonal cells indicate that the spillover coefficients are close to zero and, in some cases, negative. The coefficients of the HAR-RV-S model somewhat gain in prominence as we move on to one of the intermediate forecast horizons, $h = 5, 10$, plotted in the upper right and lower left heatmaps. While there are several positive off-diagonal coefficients, we also observe various negative estimated coefficients, especially when we consider the off-diagonal RV_{LF} coefficients in the upper part of the heatmaps. Finally, at the long-forecast horizon, $h = 20$, it appears that, while there are still some noticeable spillover effects, the own

RV_{LF} coefficients gain somewhat in relative importance again (lower right panel). Taken together, the estimated coefficients indicate that it may be possible to improve in-sample model fit by accounting for spillover effects.

– Figures 3 and 4 about here. –

The results we summarize in Figure 4 illustrate that this is indeed the case. Figure 4 plots in-sample ratios of the RMSFE for a comparison of the HAR-RV-S model with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the HAR-RV-S model produces a smaller in-sample RMSFE than the HAR-RV model. The results, irrespective of whether we study a Lasso estimator, an elastic net, or a Ridge regression estimator, indicate that the in-sample fit of the HAR-RV-S model relative to the classic HAR-RV model tends to improve as we increase the length of the forecast horizon.

4.2 Forecasting Results

In-sample fit, of course, does not necessarily carry over to an out-of-sample analysis. We start our comparison of the out-of-sample performance of the HAR-RV-S model, as estimated by the baseline stacking estimator, with that of the HAR-RV model, estimated by the OLS technique. Figure 5 depicts the resulting RMSFE ratios that we obtain when we consider a recursive-estimation window, while Figure 6 depicts the corresponding RMSFE ratios for a rolling-estimation window. The key result for both types of estimation windows is that the classic HAR-RV model outperforms the HAR-RV-S model for the vast majority of commodities, especially when we increase the length of the forecasting horizon. This key result is not sensitive to the specific choice of the shrinkage estimator (Lasso estimator, elastic net, Ridge regression).

– Figures 5 and 6 about here. –

Figure 7 shows, for the example of a recursive-estimation window, that we observe the superior performance of the HAR-RV model relative to the HAR-RV-S model also when we consider the MAFE as our metric of forecasting accuracy. The MAFE ratios should be less sensitive to large forecasting error in the wake of a sudden outburst of RV (see Figure 1) than the RMSFE ratio, but the results clearly demonstrate that our key result is robust to the change of the metric of forecast accuracy.

– Figure 7 about here. –

Figure 8 (for a recursive-estimation window) and Figure 9 (for a rolling-estimation window) summarize the results we obtain when we study the modified stacking algorithm. For the modified stacking algorithm, we consider four alternative combinations of estimators: we consider a Lasso-Lasso estimator, a Lasso-RF (that is, a general regression-tree-based residual function) estimator, a Ridge-Ridge estimator, and a Ridge-RF estimator. Four all four combinations of estimators, we use the RMSFE ratio to quantify relative forecasting performance. Across the four combinations of estimators, we observe that the HAR-RV-S model does not outperform the classic HAR-RV model. Quite to the contrary, the HAR-RV model exhibits a robust superior performance, especially as the length of the forecasting horizon increases.

– Figures 8 and 9 about here. –

4.3 Robustness Checks

As our first robustness check, we summarize in Figures 10 (recrusive-estimation window) and 11 (rolling-estimation window) results for a multivariate Lasso, a multivariate elastic net, and a multivariate Ridge regression estimator. For estimation of the multivariate shrinkage estimators, we use the R add-on package "glmnet" (Friedman et al. 2010). We observe for the short forecast horizon a few cases for which the HAR-RV-S model performs better than the classic HAR-RV model, but the general message conveyed by the results is in line with the results for the stacking algorithms. The classic HAR-RV model performs well for the majority of agricultural commodities at the short forecast horizon, and it performs better than the HAR-RV-S model at the intermediate and long forecast horizons.

– Figures 10 and 11 about here. –

As another robustness check, we study the relative performance of the HAR-RV-S model along the quantiles of the distribution of the actual realizations of RV during the out-of-sample period. We plot the results in Figures 12 (recursivce-estimation window) and 13 (rolling-estimation window), where we focus on the baseline stacking algorithm for the sake of space. We observe that the HAR-RV-S model slightly outperforms the HAR-RV model at the short forecast horizon for intermediate quantiles close to the median, mainly when we consider a Lasso estimator. For the intermediate and long forecast horizons, in contrast, the HAR-RV model clearly outperforms the HAR-RV-S model. The relative performance of the HAR-RV-S model only starts improving at some of the very upper quantiles, but this improvement is not strong enough to beat the forecasting performance of the HAR-RV model in a robust way.

– Figures 12 and 13 about here. –

As yet another robustness check, we consider the possibility that the strength of spillover effects between the RVs has been varying over time. If so, the performance of the HAR-RV-S model to the classic HAR-RV model may have undergone corresponding changes over time. In order to study this question in some more detail, we plot in Figure 14 rolling-window estimates of the Diebold and Yilmaz (2009, 2023) total dynamic spillover index (implemented using the R add-on package “Spillover”; see Urbina (2023); see Bubák et al. (2011) for a discussion of the link between the multivariate HAR-RV model and the Diebold-Yilmaz index). The estimation results clearly reveal that the strength of the total spillover effects among the RVs of the agricultural commodities in our sample has increased over time. This increase in the strength of the total spillover effects warrants a closer inspection of the relative forecasting performance of the HAR-RV-S and HAR-RV models during subsample periods.

– Figure 14 about here. –

We summarize the results of such a subsample analysis in Figures 15 (recursive-estimation window) and 16 (rolling-estimation window), where we use the first 450 out-of-sample forecasts to define the first subsample period, and the remaining forecasts to define the second subsample period (the exact number of out-of-sample forecasts depends on the forecast horizon). While we find a superior performance of the HAR-RV-S model in terms of the RMSFE ratio for some combinations of agricultural commodities and forecast horizons, the general picture that emerges from the analysis of the subsamples is that the relative forecasting performance of the HAR-RV-S model is not systematically better in the second than in the first subsample. There are a few exceptions from this general picture. For example, in some model configurations, the HAR-RV-S model outperform the HAR-RV model in the second but not in the first subsample when we consider the RVs of CC and KC, but not at all forecast horizons and not for all four combinations of estimators. Moreover, as in the case of the full-sample analysis, the relative performance of the HAR-RV model in general strengthens in the length of the forecast horizon. It is also interesting to observe that, in the first subsample, the HAR-RV-S- model outperforms the HAR-RV model for C and S (in case of the latter only for the rolling-estimation window) when we use the Lasso-Lasso and the Ridge-Ridge estimators, where the good performance of the HAR-RV-S model tends to strengthen in the length of the forecast horizon. Thus, our results for the first subsample partially overlap (that is, for C and S) with the results reported by Marfati et al. (2022), who report results for the sample period 2013−2018 (and Chinese data, so it is clear that the results are not directly comparable).

– Figures 15 and 16 about here. –

4.4 Extension to Energy Commodities and Precious Metals

It is interesting to study whether the out-of-sample results when we extend our HAR-RV-S model for agricultural commodities to include other important commodities. In extending the model in this way, we can account for the potential impact of spillover effects across different classes of commodities, as noted in a number of stdies involving the agriculture, energy and precious metals markets,[2] on out-of-sample forecasting performance. In order to extend the HAR-RV-S model in this way, we consider the RVs of several energy commodities (natural gas: NG; heating oil: HO, and; coal: CL) and precious metals (gold: GC; copper: HG; palladium: PA; platinum: PL, and; silver: SI). The data source is the same as that described in detail in Section 2, so that we can directly match by date the RVs of the agricultural commodities with those of the energy commodities and the precious metals. We plot the RVs of the energy commodities and the precious metals at the end of the paper (Appendix; Figures A2 and A1), where we also report the full-sample correlation matrix for the members of the three commodity groups (Figure A3). We also plot the corresponding total dynamic spillover index (Figure A4), which shows that the spillover effects among the three members of the three commodity groups have increased towards the end of the sample period. As in the agricultural-commodities-only model, the spillover effects are also visible in the full-sample RMSFE ratios (based on the modified stacking algorithm; Figure A5). The RMSFE ratios clearly decrease in the length of the forecast horizon, especially when we combine the shrinkage estimators with regression trees. Finally, the results that we report in Figure A6 (for a recursive-estimation window) and in Figure A7 (for a rolling-estimation window) corroborate the main finding of our empirical research that the HAR-RV-S model, with only few exceptions, does not outperform out-of-sample the forecasting performance of the classic HAR-RV model, especially as the length of the forecasting horizon increases.

[2]See, for example, Nazlioglu et al. (2013), Kang et al. (2017), Mensi et al. (2017), Ji et al. (2018), Luo and Ji (2018), Chang et al. (2019), Lu et al. (2019), Dahl et al. (2020), Luo and Zhang (2020), Yip et al. (2020).

5 Concluding Remarks

Modeling and forecasting realized volatilities of financial asset prices in general and of commodity price fluctuations in particular is of key importance for financial market participants and policymakers. Financial market participants rely on accurate forecasts of realized volatilities when solving portfolio-optimization problems and pricing of derivative securities. Policymakers, in turn, need accurate forecasts of realized volatilities when designing policies to mitigate the potential adverse effects of a rise in economic, and in case of agricultural commodities perhaps even political uncertainty, associated with sudden increases in the volatility of price fluctuations. A natural and important research question, therefore, is whether forecasts of the realized volatilities of commodity price fluctuations can benefit when a forecaster takes into account spillover effects across the realized volatilities of agricultural commodities. The results we have reported in this research clearly demonstrate that such spillover effects exist, that they can be strong, that they may vary over time, and that accounting for spillover effects by means of a simple HAR-RV-S model has beneficial effects in an in-sample analysis. We do not observe, however, systematic out-of-sample forecasting gains relative to a classic HAR-RV model.

In order to obtain out-of-sample forecast the RVs of 15 agricultural commodities (and, in an extended model, three energy commodities and five precious metals), we have used various multi-task stacking algorithms as well as a multivariate shrinkage estimator. The multi-task stacking algorithms in particular have the advantage that it is straightforward to implement them in high dimensional multi-task forecasting problems. Modeling and forecasting the realized volatilities of the various agricultural commodities that we have studied in this research can be interpreted as belonging to this class of problems. While the multivariate shrinkage estimator retains a simple linear structure of the forecasting model, the multi-task stacking algorithm open up the possibility to combine different base and meta learners, where for the latter we also have used regression trees so as to explore potential nonlinear structures in the data. Irrespective of the algorithm or combination of base and meta learners that we have studied, we have obtained the same main finding that spillover effects do not leverage out-of-sample forecast accuracy relative to the classic HAR-RV model. Our main finding implies that the research strategy used by some researchers in recent papers (see, for example, Bonato et al. (2022, 2024, forthcoming)) to forecast RVs of agricultural commodities in an univariate modeling approach is likely to be a good starting point for further analysis, and can be also considered beneficial from the perspective of investors looking for optimal portfolio allocations, and policymakers aiming to stabilize food prices.

This does not mean that a multivariate modeling approach cannot yield important additional and novel insights. In fact, in future research, it is interesting to study whether other algorithms developed in the large and rapidly growing machine-learning literature corroborate our main finding, or whether application of other algorithms brings to the forefront features of the data that the algorithms and estimators we have applied in our research have failed to detect. In technical terms, it is also interesting to explore how the stacking algorithms we have studied in this research can be combined with the type of multivariate HAR-RV-cum-GARCH models discussed in the related earlier literature. Such an extension also would render it possible to compare more directly the results we have reported in this paper with the results that Marfatia et al. (2022) have reported in their recent empirical study of a small set of agricultural commodities (and a shorter sample period, using Chinese data). Furthermore, against the background of the much discussed financialization of commodity markets, it is worthwhile to investigate whether accounting for spillover effects across the realized volatilities of different asset classes (for example, agricultural commodities and stock markets) yields insights that help to improve the accuracy of out-of-sample forecasts of realized volatilities (as in, for example, De Nicola et al. (2016)).

References

Aït-Youcef, C. (2019). How index investment impacts commodities: A story about the financialization of agricultural commodities. Economic Modelling, 80, 23–33.

Andersen T.G., and Bollerslev T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. International Economic Review, 39(4), 885–905.

Asai, M., Gupta, R., and McAleer, M. (2019). The Impact of jumps and leverage in forecasting the co-Volatility of oil and gold futures. Energies, 12(17), 1–17.

Asai, M., Gupta, R., and McAleer, M. (2020). Forecasting volatility and co-volatility of crude oil and gold futures: Effects of leverage, jumps, spillovers, and geopolitical risks. International Journal of Forecasting, 36(3), 933–948.

Asai, M., and McAleer, M. (2017). The impact of jumps and leverage in forecasting co-volatility. Econometric Reviews, 36(6-9), 638–650.

Bahloul, W., Balcilar, M., Cunado, J., and Gupta, R (2018). The role of economic and financial uncertainties in predicting commodity futures returns and volatility: Evidence from a nonparametric causality-in-quantiles test. Journal of Multinational Financial Management, 45, 52–71.

Beckmann, J., and Czudaj, R. (2014). Volatility transmission in agricultural futures markets. Economic Modelling, 36, 541–546.

Bonato, M. (2019). Realized correlations, betas and volatility spillover in the agricultural commodity market: What has changed? Journal of International Financial Markets, Institutions and Money, 62, 184–202.

Bonato, M., Cepni, O. Gupta, R., and Pierdzioch, C. (2022). El Niño, La Niñna, and forecastability of the realized variance of agricultural commodity prices: Evidence from a machine learning approach. Journal of Forecasting, 42(4), 785–801.

Bonato, M., Cepni, O. Gupta, R., and Pierdzioch, C. (2024). Forecasting the realized volatility of agricultural commodity prices: Does sentiment matter? Journal of Forecasting. DOI: https://doi.org/10.1002/for.3106.

Bonato, M., Cepni, O. Gupta, R., and Pierdzioch, C. (Forthcoming). Financial stress and realized volatility: The case of agricultural commodities. Research in International Business and Finance.

Breiman, L. (1984) Classification and regression Trees. Wadsworth, Belmont, CA.

Bubák, V., Kočenda, E., and Žikeš, F. (2011). Volatility transmission in emerging European foreign exchange markets. Journal of Banking and Finance, 35, 2829–2841.

Cagli, E.C., Mandaci, P.E., and Taskin, D. (2023). The volatility connectedness between agricultural commodity and agri businesses: Evidence from time-varying extended joint approach. Finance Research Letters, 52, 103555.

Chang, C.L., Liu, C.P., and McAleer, M. (2019). Volatility spillovers for spot, futures, and ETF prices in agriculture and energy. Energy Economics, 81, 779–792.

Chatziantoniou, I., Degiannakis, S., Filis, G., and Lloyd, T. (2021). Oil price volatility is effective in predicting food price volatility. Or is it? Energy Journal, 42(6), 25–48.

Čech, F., and Baruník, J. (2017). On the modelling and forecasting of multivariate realized volatility: Generalized Heterogeneous Autoregressive (GHAR) Model. Journal of Forecasting, 36(2), 181–206.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. Journal of Financial Econometrics, 7(2), 174–196.

Dahl, R.E., Oglend, A., and Yahya, M. (2020). Dynamics of volatility spillover in commodity markets: linking crude oil to agriculture. Journal of Commodity Markets, 20, 100–111.

Degiannakis, S., Filis, G., Klein, T., and Walther, T. (2022). Forecasting realized volatility of agricultural commodities. International Journal of Forecasting, 38(1), 74–96.

De Nicola, F., De Pace, and Hernandez (2016). Co-movement of major energy, agricultural, and food commodity price returns: A time-series assessment. Energy Economics, 57, 28–41.

Diebold, F.X., and Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. Economic

Journal, 119(534), 158–171.

Diebold, F.X., and Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. International Journal of Forecasting, 28(1), 57–66.

FAO. (2010). Price volatility in agricultural markets: Evidence, impact on food security and policy responses. Available for download from: http://www.fao.org/docrep/013/am053e/am053e00.pdf.

FAO. (2011). Price volatility in food and agricultural markets: Policy responses. Available for download from: http://www.fao.org/fileadmin/templates/est/Volatility/Interagency_Report_to_the_G20_on_Food_Price_Volatility.pdf.

Friedman, J., Tibshirani, R., and Hastie, T. (2010). Regularization paths for generalized linear models via coordinate descent. Journal of Statistical Software, 33(1), 1–22. Available for download from: https://doi.org/10.18637/jss.v033.i01.

Gardebroek, C., Hernandez, M., and Robles, M. (2016). Market interdependence and volatility transmission among major crops. Agricultural Economics. 47(2), 141–155.

Gil-Alana, L.A., Cunado, J., and de Gracia, F.P. (2012). Persistence, long memory, and unit roots in commodity prices. Canadian Journal of Agricultural Economics, 60(4), 451–468.

Greb, F., and Prakash, A. (2015). Has price volatility changed? Food Outlook, Rome, Italy: Food and Agriculture Organization of the United Nations.

Greb, F., and Prakash, A. (2017). Assessing volatility patterns in food crops. In FAO commodity and trade policy research working paper series, Available for download from: http://www.fao.org/3/a-i7066e.pdf.

Gupta, R., Pierdzioch, C. (2023). Climate risk and the volatility of agricultural commodity price fluctuations: A prediction experiment. In: Bourghelle, D., Grandin, P., Jawadi, F., and Rozin, P. (eds.) Behavioral finance and asset prices: The influence of investor's emotions. Contributions to Finance and Accounting, 23–44, Springer, Cham, United States.

Ji, Q., Bahloul, W., Geng, J.-B., and Gupta, R. (2020). Trading behaviour connectedness across commodity markets: evidence from the hedgers' sentiment perspective. Research in International Business and Finance, 52, 101114.

Ji, Q., Bouri, E., Roubaud, D., and Shahzad, S.J.H. (2018). Risk spillover between energy and agricultural commodity markets: A dependence-switching CoVaR-copula model. Energy Economics, 75, 14–27.

Kang, S.H., McIver, R., and Yoon, S.M. (2017). Dynamic spillover effects among crude oil, precious metal, and agricultural commodity futures markets. Energy Economics, 62, 19–32.

Lu, Y., Yang, L., and Liu, L. (2019). Volatility spillovers between crude oil and agricultural commodity markets since the financial crisis. Sustainability, 11, 396.

Luo, J. and Chen, L. (2019). Multivariate realized volatility forecasts of agricultural commodity futures. Journal of Futures Markets, 39(12), 1565−1586.

Luo, J., and Chen, L. (2020a). Modeling and forecasting the multivariate realized volatility of financial markets with time-varying sparsity. Emerging Markets Finance and Trade, 56(2), 392−408.

Luo, J., and Chen, L. (2020b). Realized volatility forecast with the Bayesian random compressed multivariate HAR model. International Journal of Forecasting, 36(3), 781−799.

Luo, J., Cepni, O., Demirer, R., and Gupta, R. (Forthcoming). Forecasting multivariate volatilities with exogenous predictors: An application to industry diversification strategies. Journal of Empirical Finance.

Luo, J., and Ji, Q. (2018). High-frequency volatility connectedness between the US crude oil market and China's agricultural commodity markets. Energy Economics, 76, 424−438.

Luo, J., Klein, T., Ji, Q., and Hou, C. (2019). Forecasting realized volatility of agricultural commodity futures with infinite hidden Markov HAR models. International Journal of Forecasting, 38(1), 51−73.

Luo, J., Marfatia, H.A., Ji, Q., and Klein, T. (2023). Co-volatility and asymmetric transmission of risks between the global oil and China's futures markets. Energy Economics, 117, 106466.

Luo, J., and Zhang, Q. (2020). Risk contagions between global oil markets and China's agricultural commodity markets under structural breaks. Applied Economics, 53(46), 1–22.

Luo, J., and Zhang, Q. (2024). Air pollution, weather factors, and realized volatility forecasts of agricultural commodity futures. Journal of Futures Markets, 44(2), 151–217.

Marfatia, H.A., Luo, J., and Ji, Q. (2022). Forecasting the volatility of agricultural commodity futures: The role of co-volatility and oil volatility. Journal of Forecasting, 41(2), 383–404.

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

Mensi, W., Tiwari, A.K., Bouri, E., Roubaud, D., and Al-Yahyaee, K.H. (2017). The dependence structure across oil, wheat, and corn: A wavelet-based copula approach using implied volatility indexes. Energy Economics, 66, 122–139.

Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., and Pictet, O.V. (1997). Volatilities of different time resolutions – Analyzing the dynamics of market components. Journal of Empirical Finance, 4(2-3), 213–239.

Nazlioglu, S., Erdem, C., and Soytas, U. (2013). Volatility spillover between oil and agricultural commodity markets. Energy Economics, 36, 658–665.

Ordu, B.M., Oran, A., and Soytas, U. (2018). Is food financialized? Yes, but only when liquidity is abundant. Journal of Banking and Finance, 95, 82–96.

Poon, S-H., and Granger, C.W.J. (2003). Forecasting volatility in financial markets: A review. Journal of Economic Literature, 41(2), 478–539.

R Core Team (2023). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. Available for download from: https://www.R-project.org/.

Rapach, D.E., Wohar, M.E., and Strauss, J. (2008). Forecasting stock return volatility in the presence of structural breaks. In: Rapach, D.E., and Wohar, M.E. (eds.) Forecasting in the Presence of Structural Breaks and Model Uncertainty. Frontiers of Economics and Globalization, 381–416, Emerald, Bingley, United Kingdom.

Rauschenberger, A., and Glaab, E. (2021). Predicting correlated outcomes from molecular data. Bioinformatics, 37(21), 3889–3895.

Shiba, S., Aye, G.C., Gupta, R., and Goswami, S. (2022). Forecastability of agricultural commodity futures realised volatility with daily infectious disease-related uncertainty. Journal of Risk and Financial Management, 15(11), 525.

Tian, F., Yang, K., and Chen, L. (2017a). Realized volatility forecasting of agricultural commodity futures using the HAR model with time-varying sparsity. International Journal of Forecasting, 33(1), 132–152.

Tian, F., Yang, K., and Chen, L. (2017b). Realized volatility forecasting of agricultural commodity futures using long memory and regime switching. Journal of Forecasting, 36(4), 421–430.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B: Statistical Methodology, 58(1), 267–288.

Trujillo-Barrera, A., Mallory, M., and Garcia, P. (2012). Volatility spillovers in U.S. crude oil, ethanol, and corn futures markets. Journal of Agricultural Resource Economics. 37(2), 247–262.

Urbina, J. (2023). Spillover: Spillover/Connectedness Index Based on VAR Modelling. R package version 0.1.0.3. Available for download from: https://CRAN.R-project.org/package=Spillover.

Xing, L., Lesperance, M.L., and Zhang, X. (2020). Simultaneous prediction of multiple outcomes using revised stacking algorithms. Bioinformatics, 36(1), 65–72.

Yang, K., Tian, F., Chen, L., and Li, S. (2017). Realized volatility forecast of agricultural futures using the HAR models with bagging and combination approaches. International Review of Economics and Finance, 49, 276–291.

Yip, P.S., Brooks, R., Do, H.X., and Nguyen, D.K. (2020). Dynamic volatility spillover effects between oil and agricultural products. International Review of Financial Analysis, 69, 101465.

Zhao, A.B., and Cheng, T. (2022). Stock return prediction: Stacking a variety of models. Journal of Empirical Finance, 67, 288–317.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society Series B: Statistical

Methodology, 67(2), 301–320.

Živkov, D., Durašković, J., and Gajić-Glamočlija, M. (2022). Multiscale downside risk interdependence between the major agricultural commodities. Agribusiness, 38(4), 990–1011,

Živkov, D., Njegić, J., and Pećanac M. (2019): Multiscale interdependence between the major agricultural commodities. Agricultural Economics – Czech, 65(2), 82–92.

Figure 1: RV of Agricultural Commodities

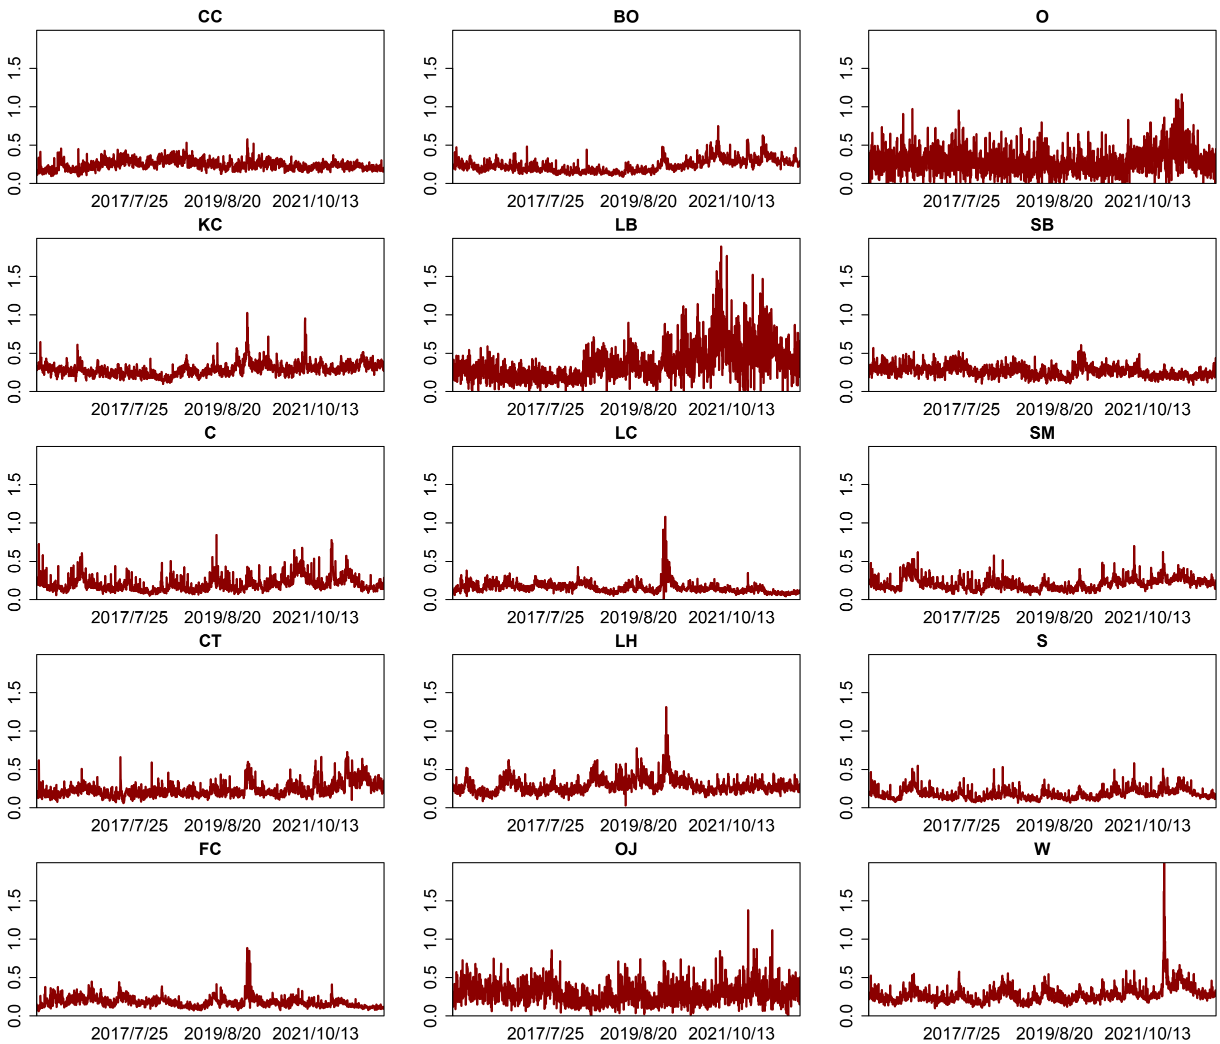

Figure 2: Full-Sample Correlation Matrix

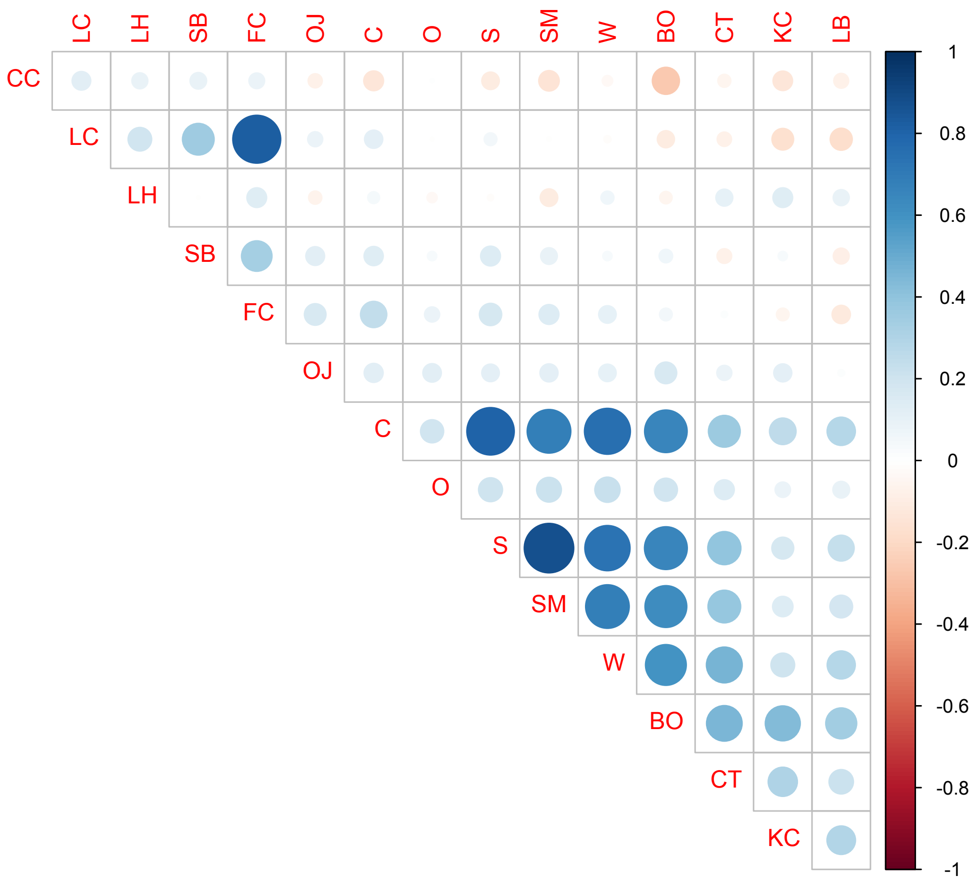

Figure 3: Full-Sample Estimated Coefficients (Baseline Stacking Algorithm)

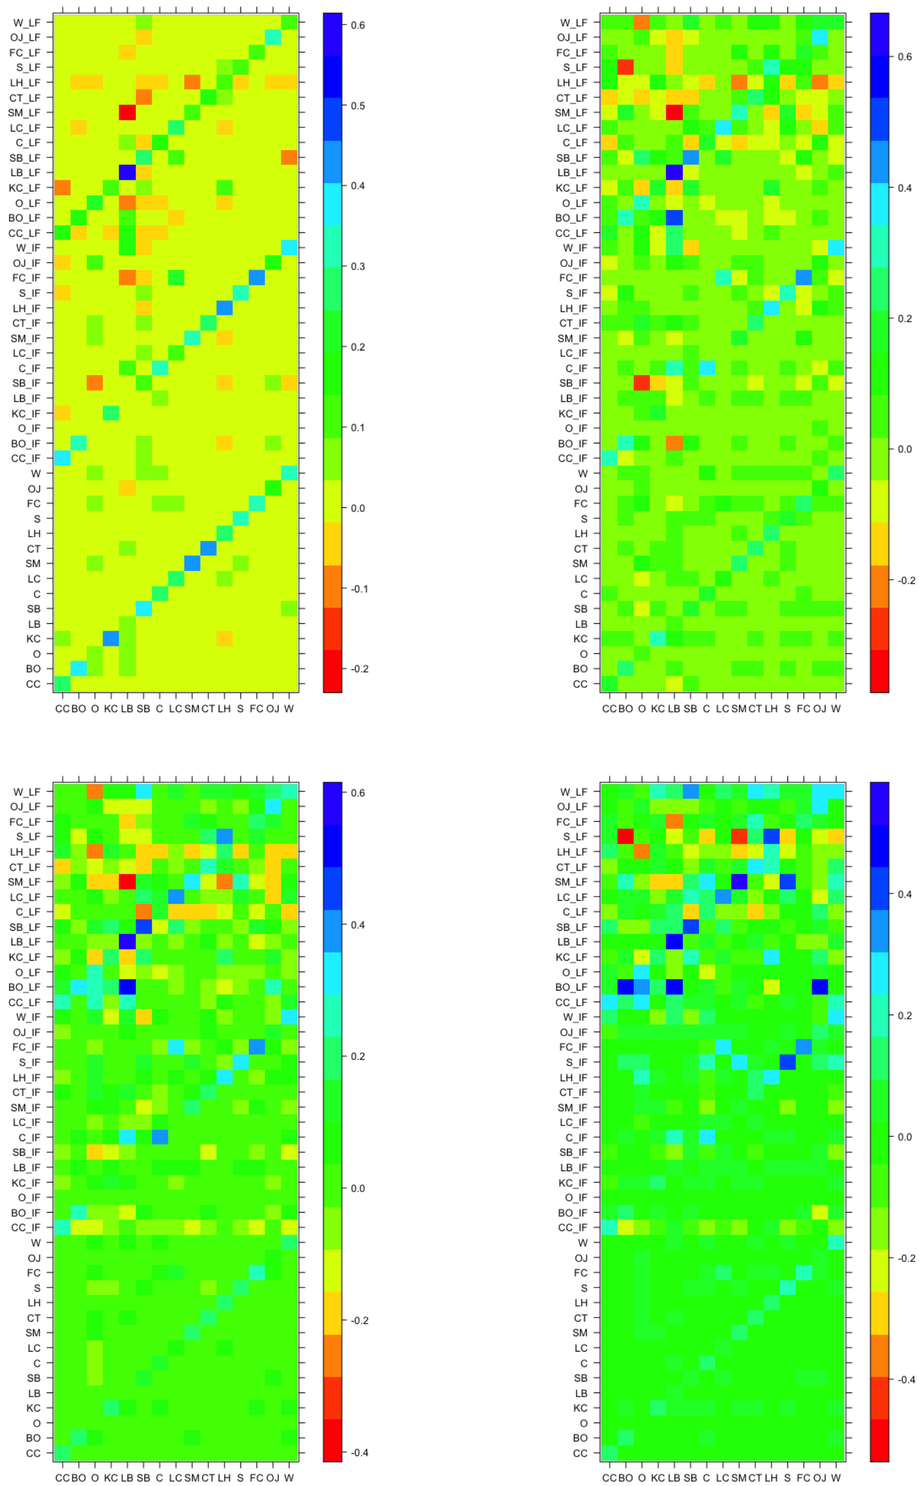

The forecast horizons are $h = 1, 5, 10, 20$ (starting in the upper left panel).

Figure 4: RMSFE Ratios for the Full Sample (Baseline Stacking Algorithm)

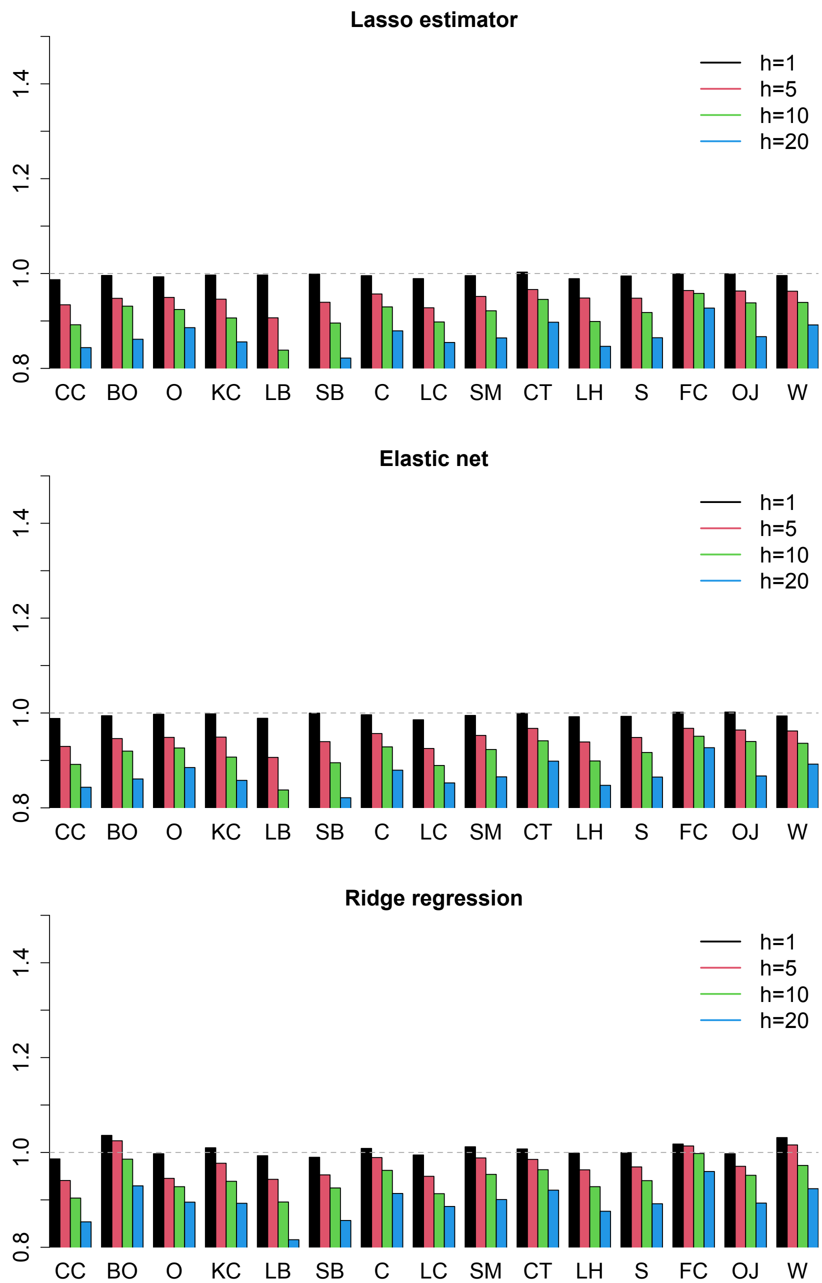

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 5: RMSFE Ratios for a Recursive Window (Baseline Stacking Algorithm)

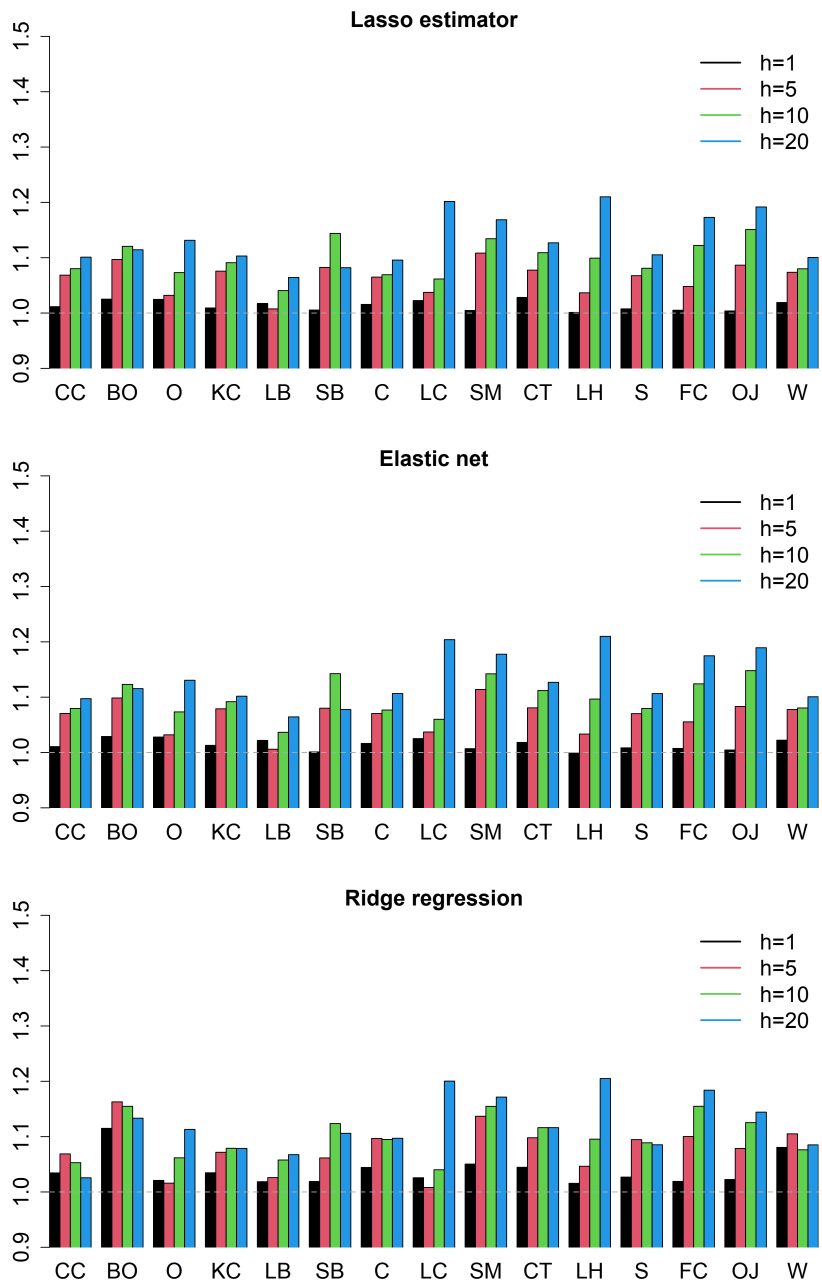

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 6: RMSFE Ratios for a Rolling Window (Baseline Stacking Algorithm)

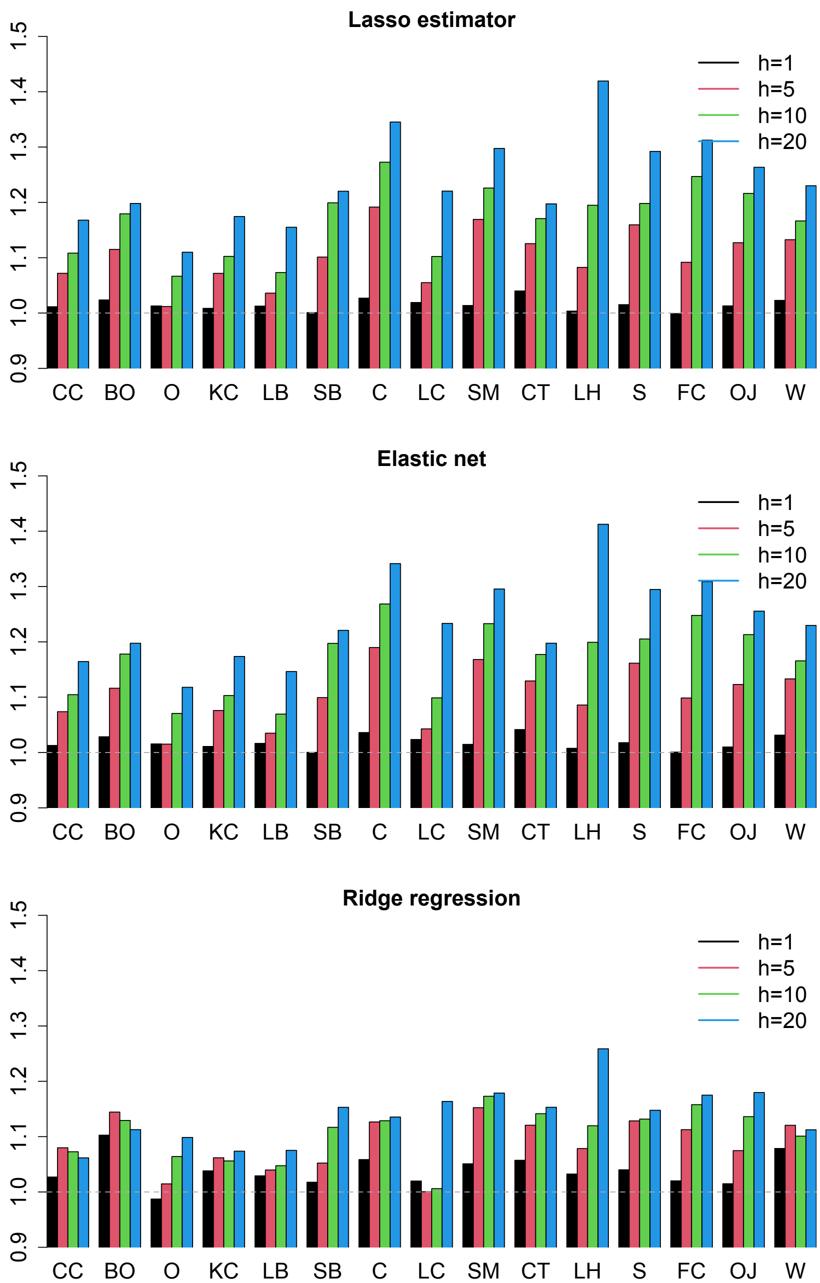

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 7: MAFE Ratios for a Recursive Window (Baseline Stacking Algorithm)

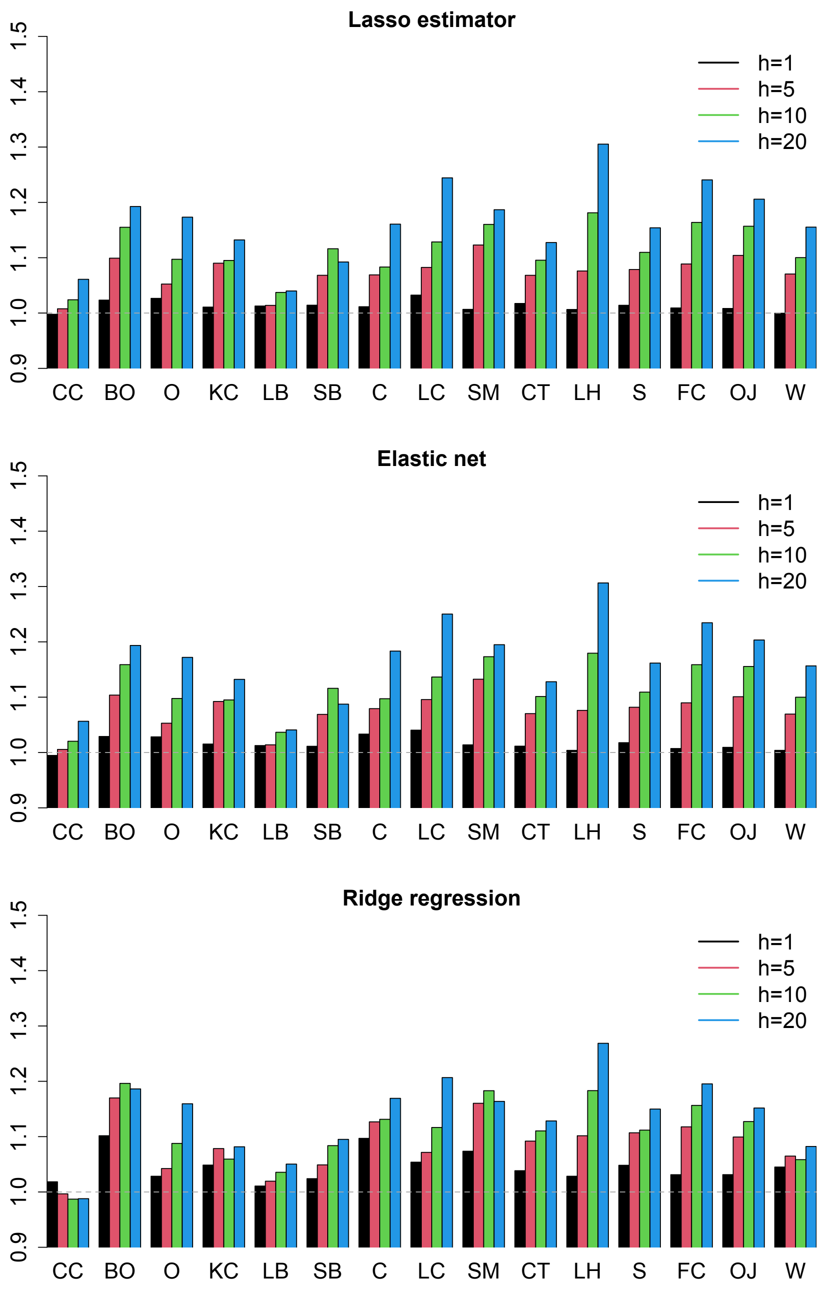

MAFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A MAFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample MAFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 8: RMSFE Ratios for a Recursive Window (Modified Stacking Algorithm)

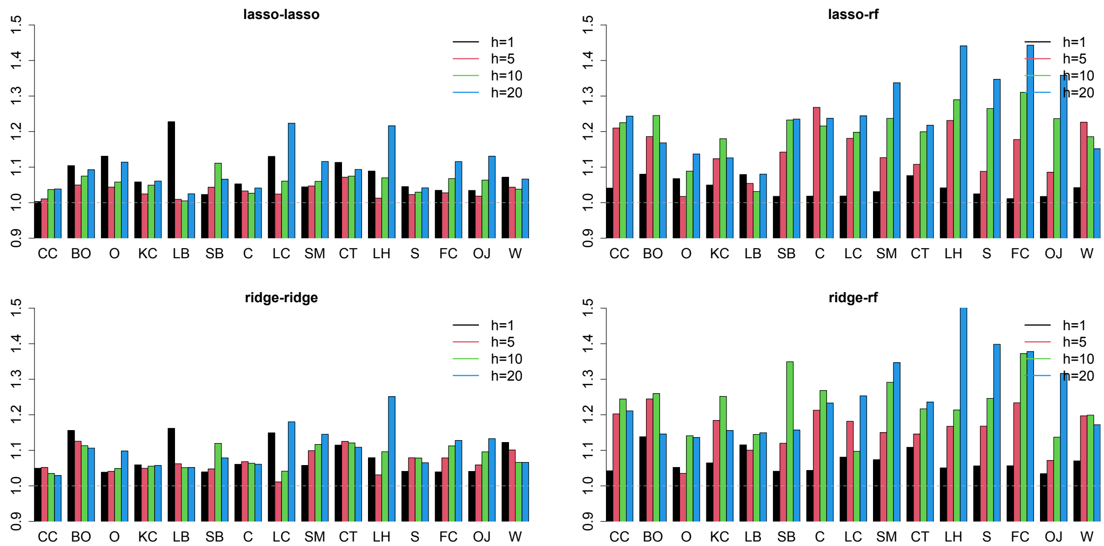

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 9: RMSFE Ratios for a Rolling Window (Modified Stacking Algorithm)

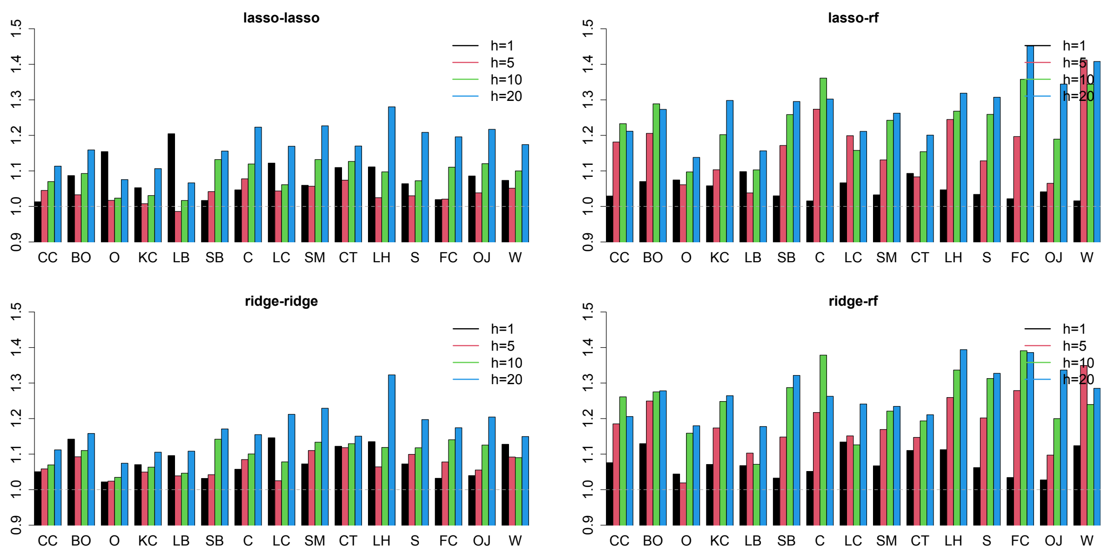

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 10: RMSFE Ratios for a Recursive Window (Multivariate Shrinkage Estimator)

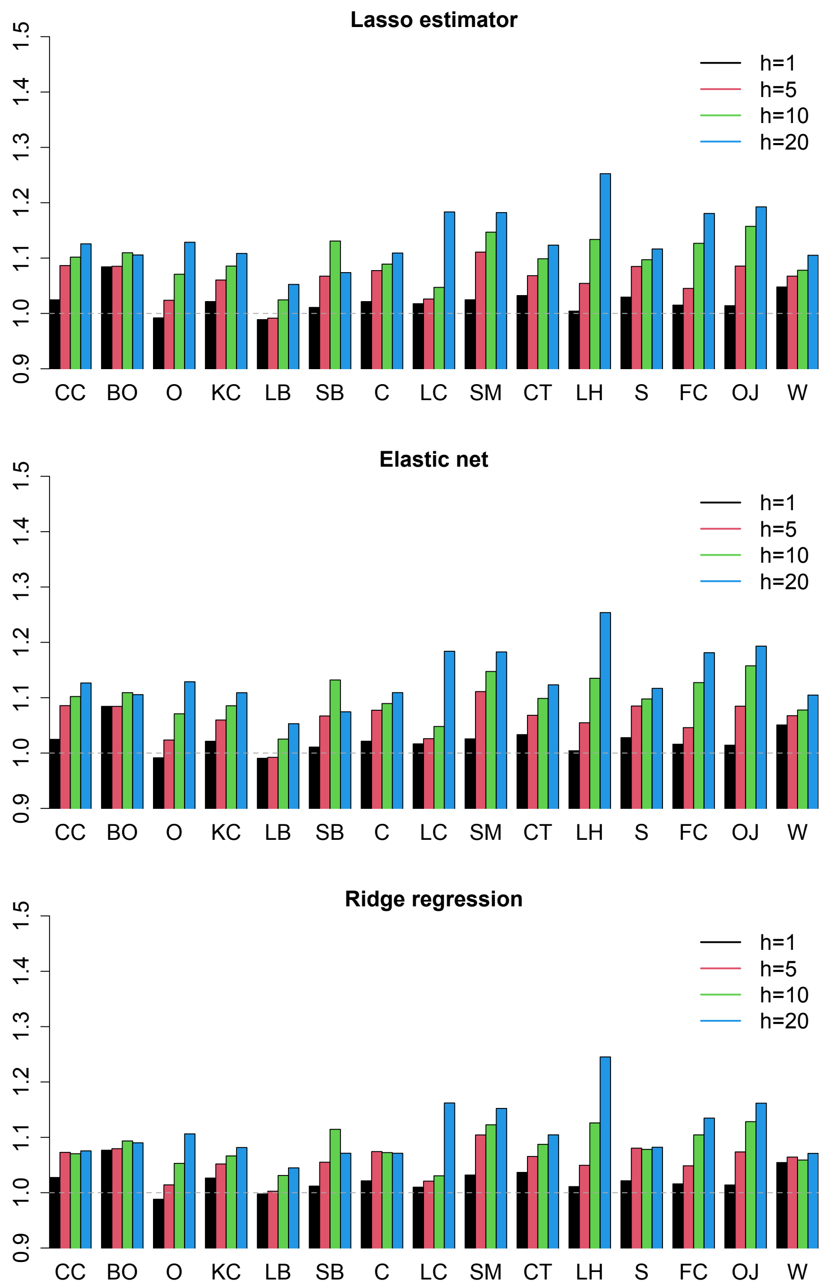

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 11: RMSFE Ratios for a Rolling Window (Multivariate Shrinkage Estimator)

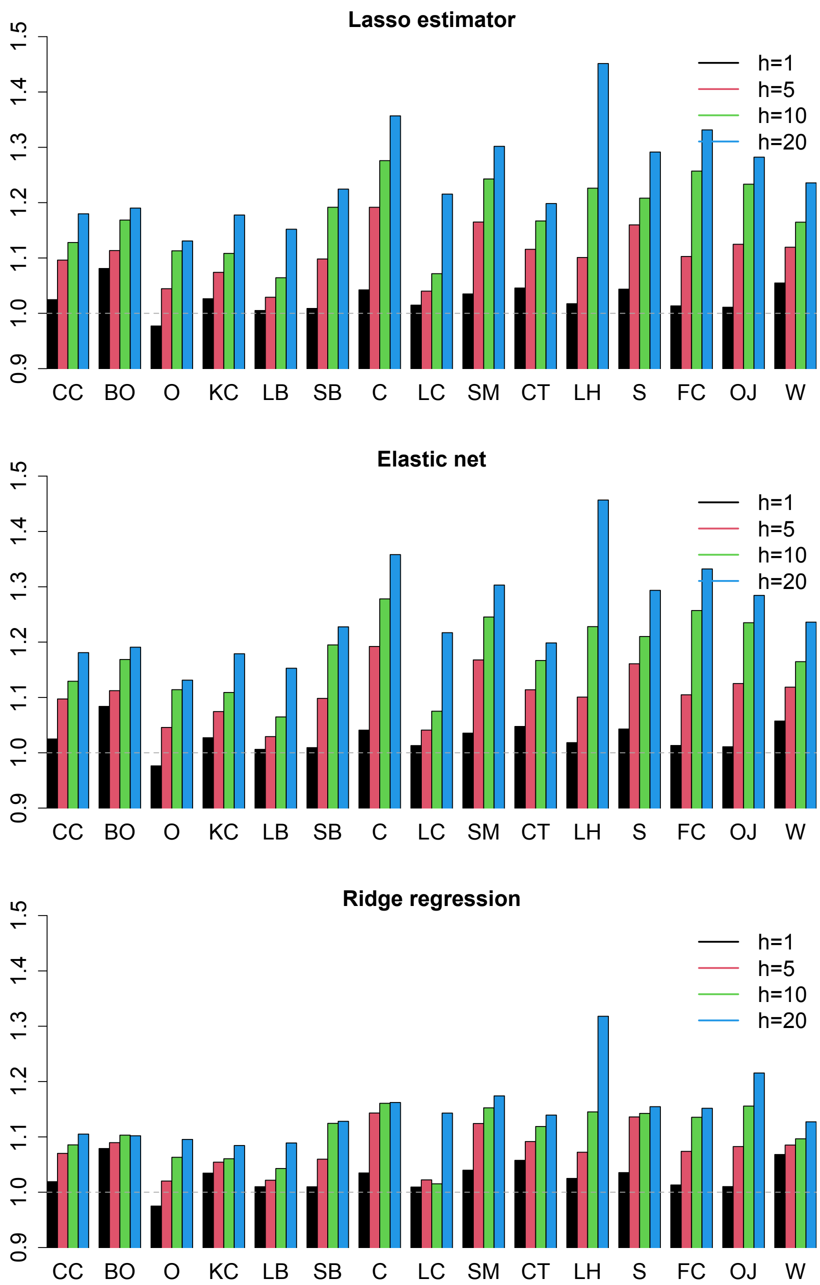

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 12: Quantile-Based RMSFE Ratios for a Recursive Window (Baseline Stacking Algorithm)

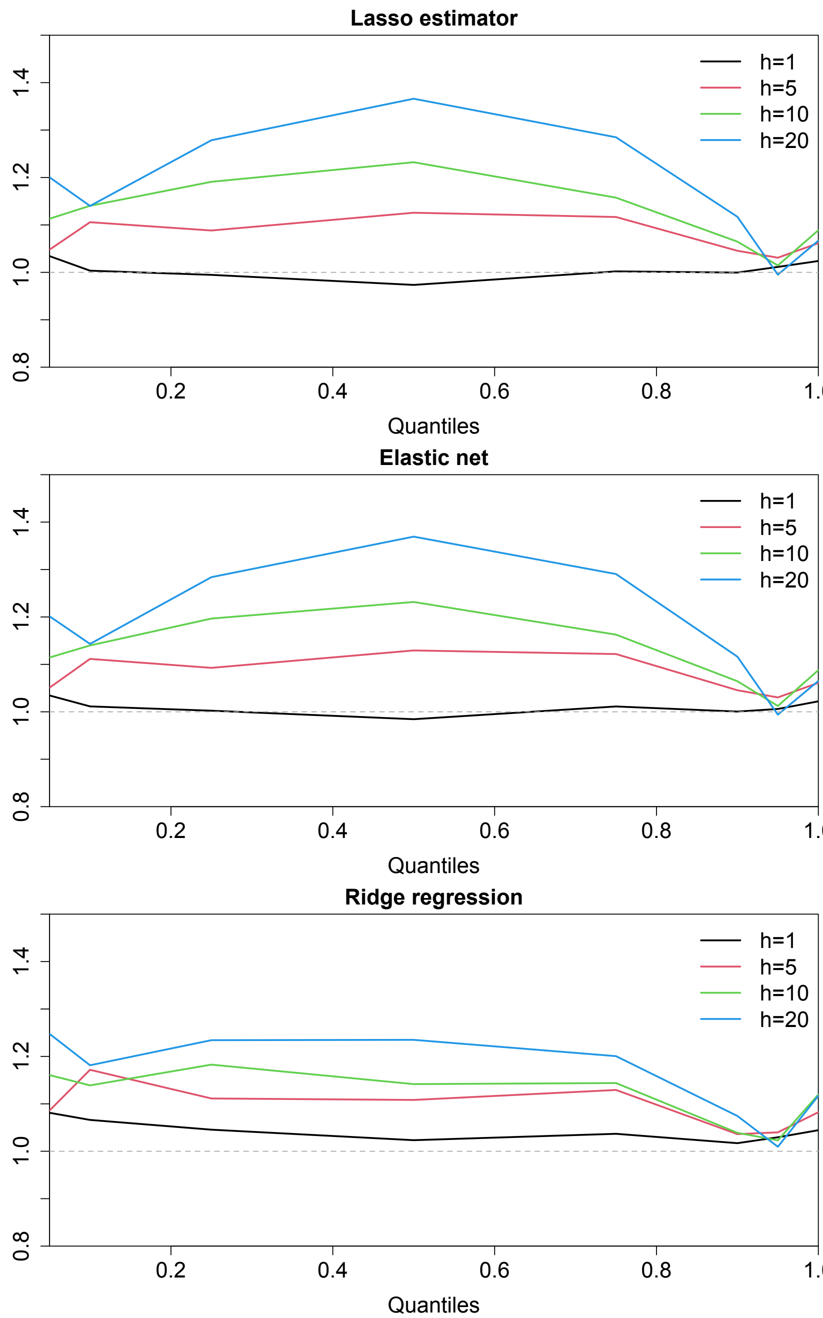

RMSFE ratios for different quantiles of the realizations of RV. RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 13: Quantile-Based RMSFE Ratios for a Rolling Window (Baseline Stacking Algorithm)

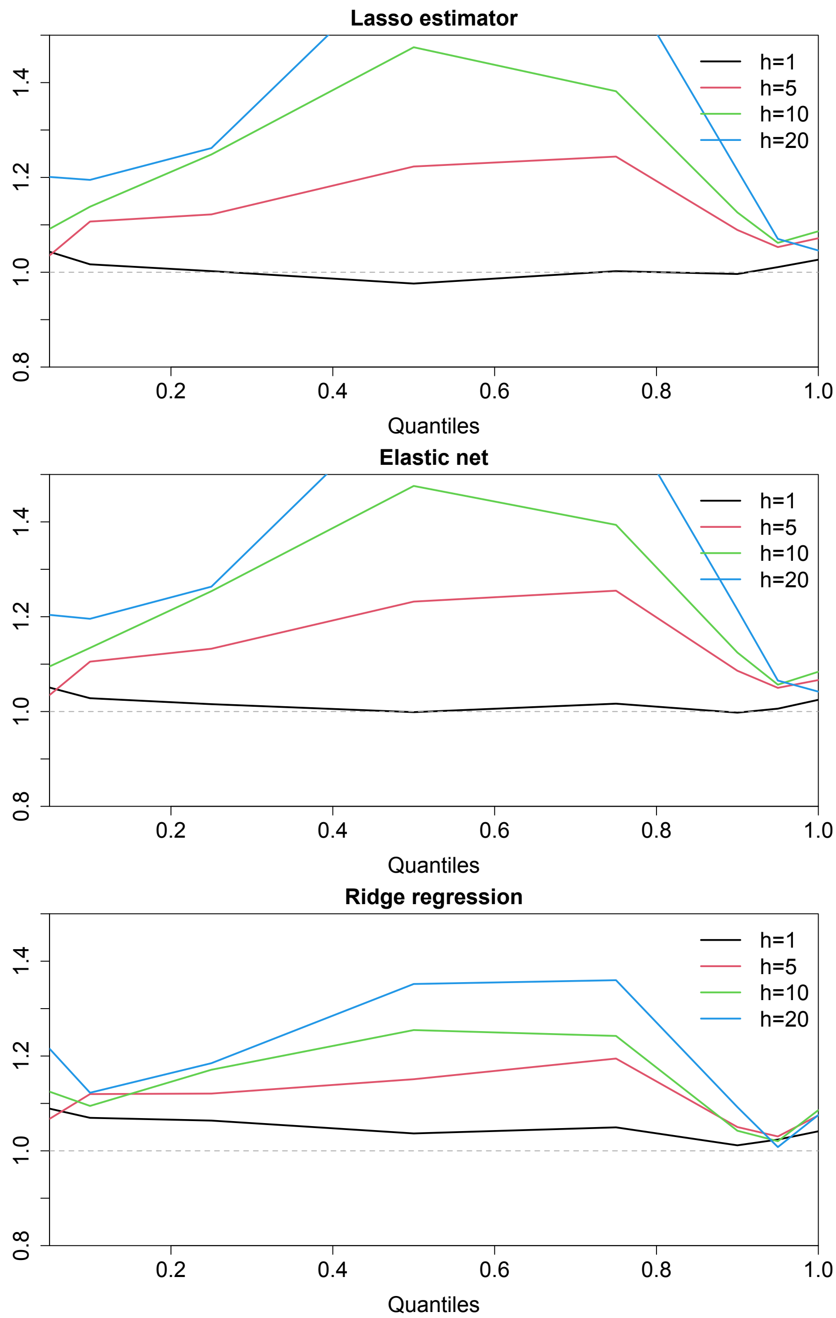

RMSFE ratios for different quantiles of the realizations of RV. RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 14: Rolling-Window Estimates of a Spillover Index

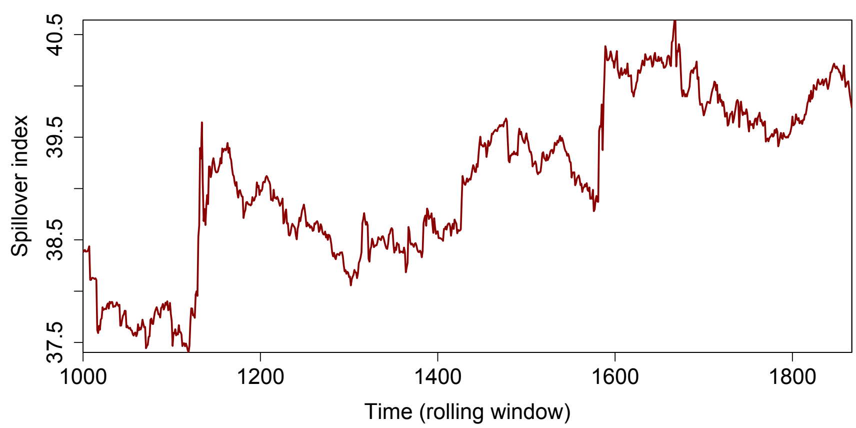

The total dynamic spillover index is derived from a VAR(5) model estimated using a rolling-estimation window of length 1,000 observations and a 10-step-ahead generalized forecast error variance decomposition.

Figure 15: Subsample Analysis for a Recursive Window (Modified Stacking Algorithm)

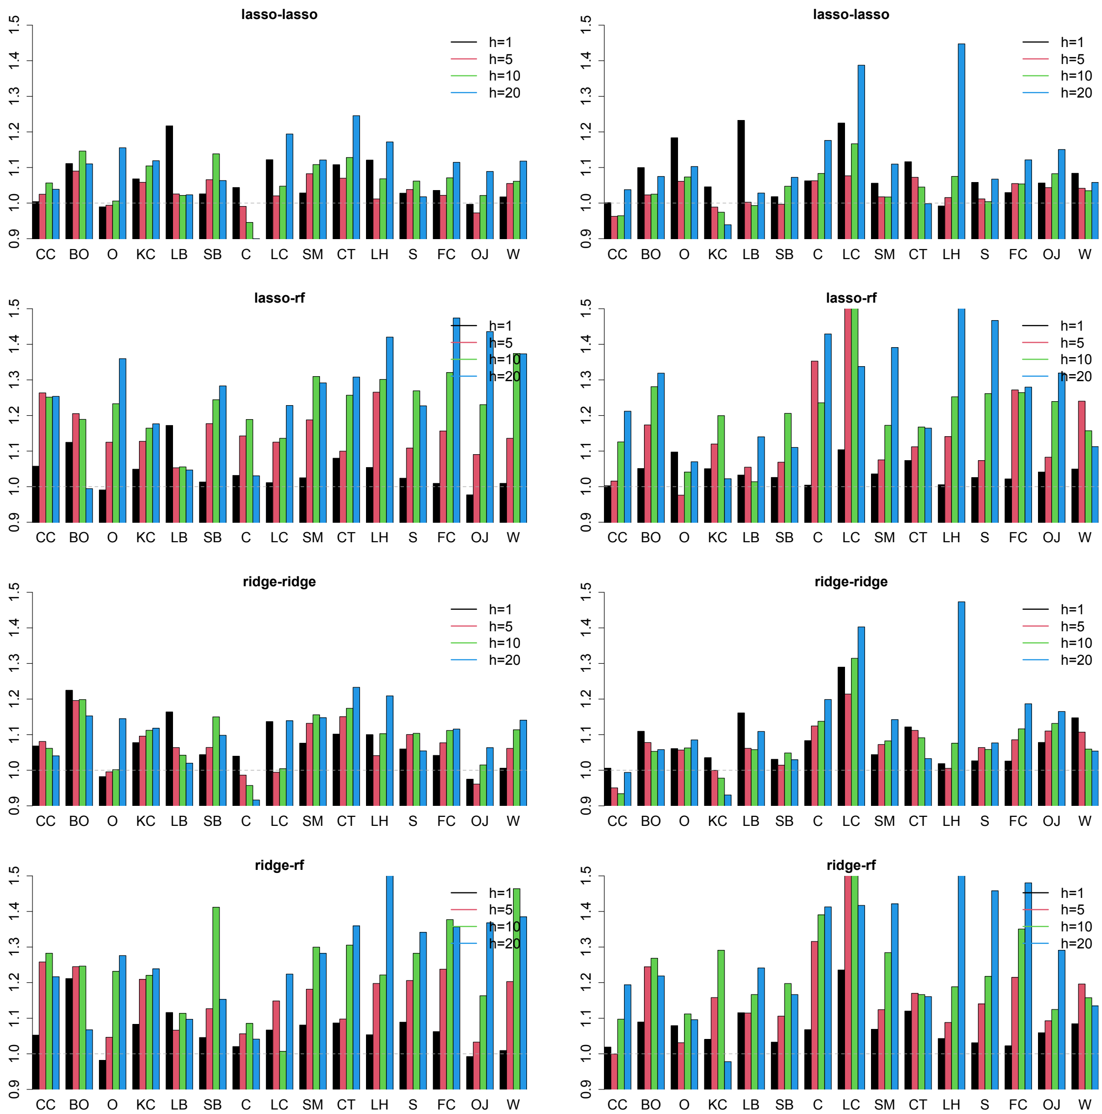

The panels on the left-hand side summarize the results for the first subsample. The panels on the right-hand side summarize the results for the first subsample. The first subsample comprises the first 450 out-of-sample forecasts. The second subsample obtains upon deleting the first 450 out-of-sample forecasts. RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure 16: Subsample Analysis for a Rolling Window (Modified Stacking Algorithm)

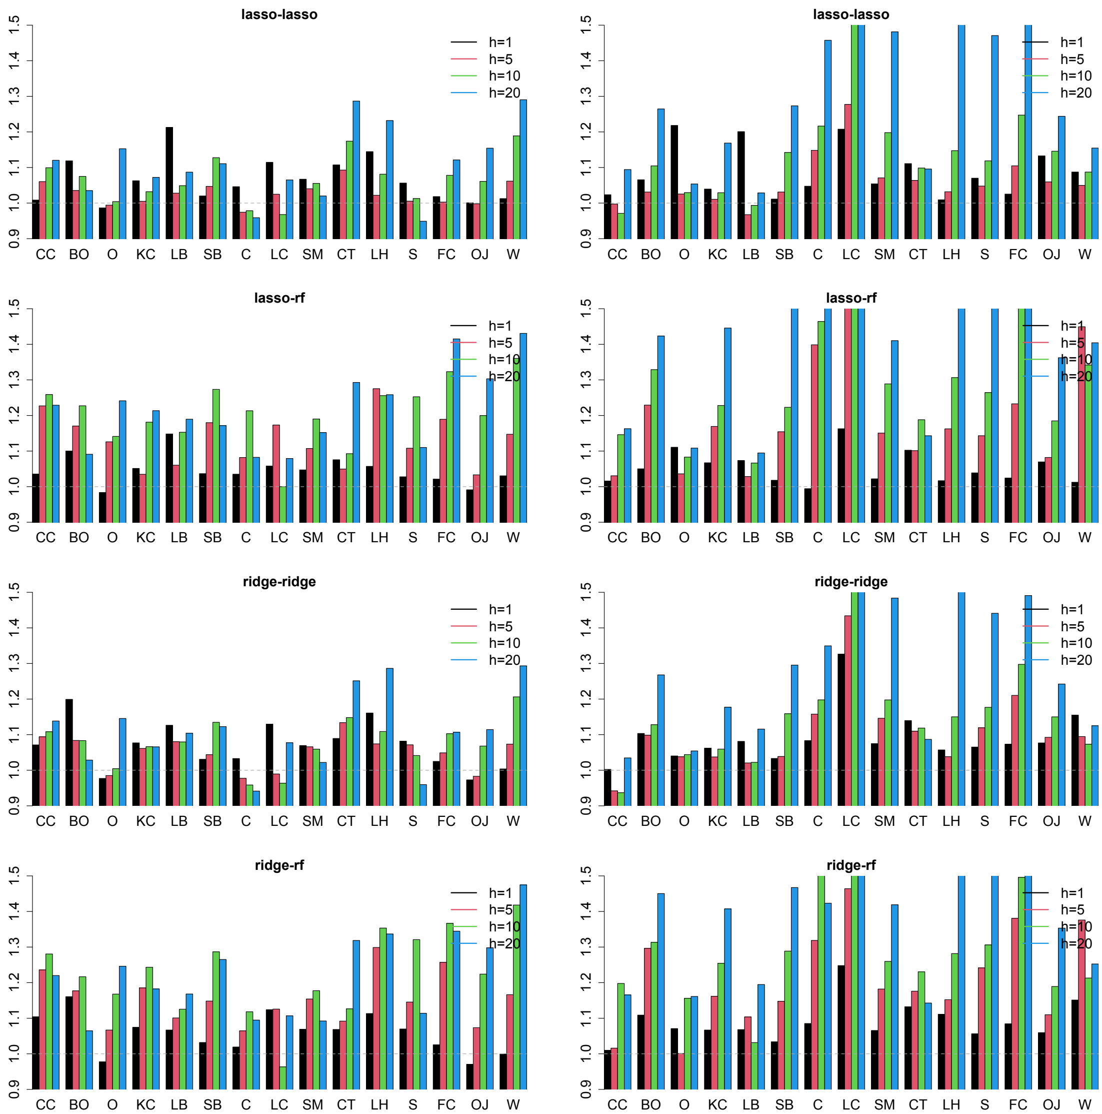

The panels on the left-hand side summarize the results for the first subsample. The panels on the right-hand side summarize the results for the first subsample. The first subsample comprises the first 450 out-of-sample forecasts. The second subsample obtains upon deleting the first 450 out-of-sample forecasts. RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Appendix

Figure A1: RV of Energy Commodities

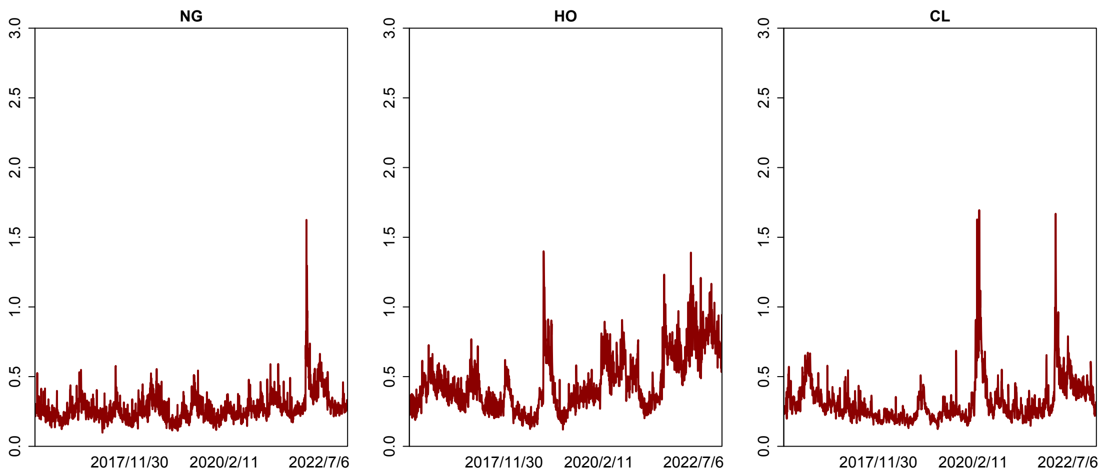

Figure A2: RV of Precious Metals

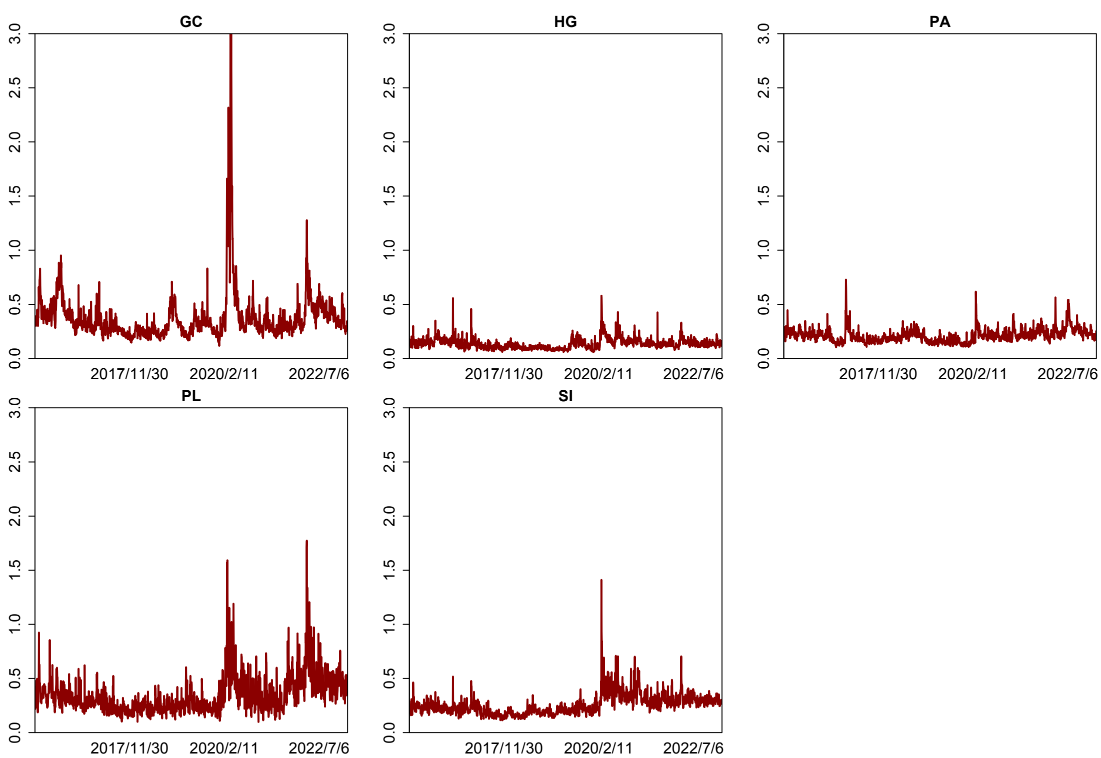

Figure A3: Full-Sample Correlation Matrix (Extended Sample of Commodities)

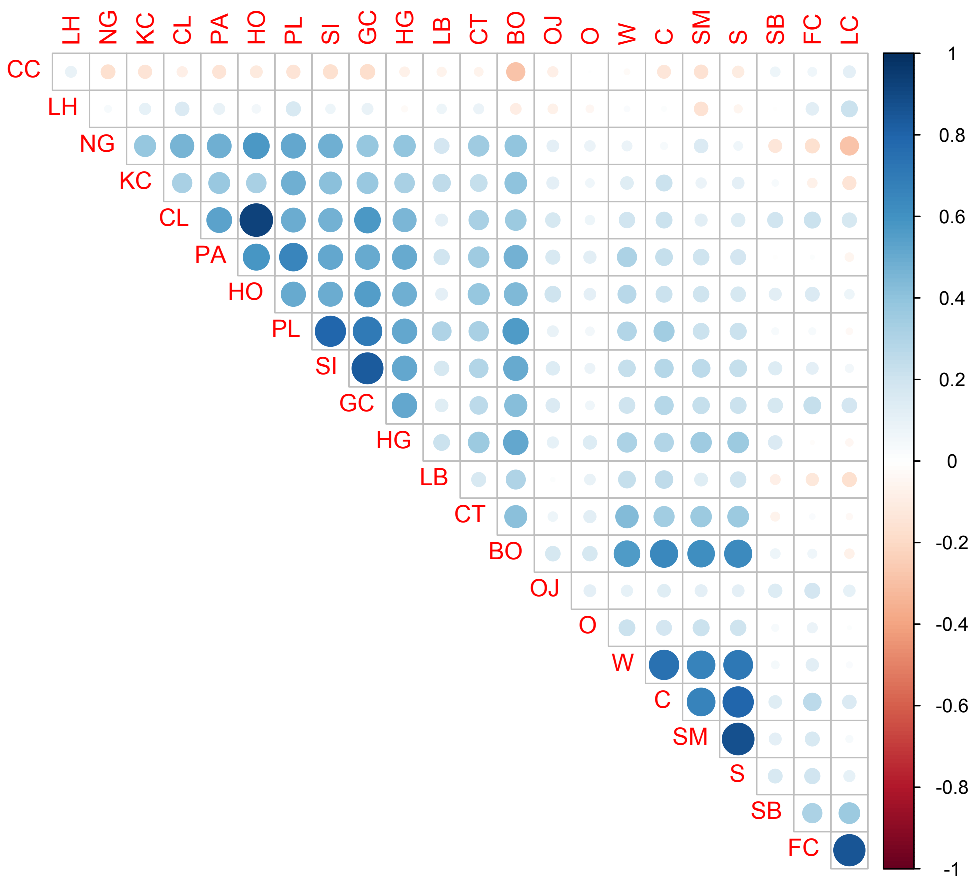

Figure A4: Rolling-Window Estimates of a Spillover Index (Extended Sample of Commodities

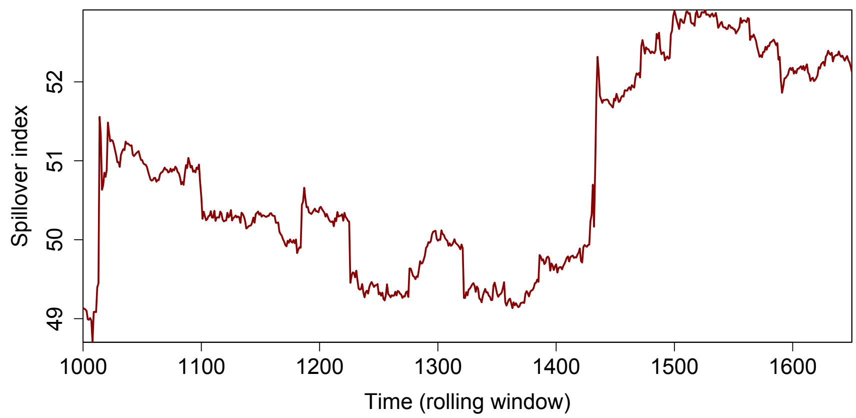

The total dynamic spillover index is derived from a VAR(5) model estimated using a rolling-estimation window of length 1,000 observations and a 10-step-ahead generalized forecast error variance decomposition.

Figure A5: Full Sample RMSFE Ratios (Modified Stacking Estimator)

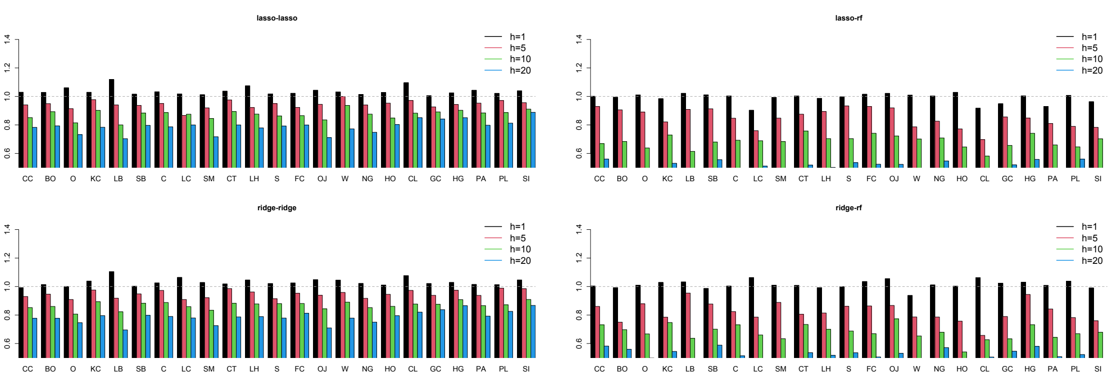

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure A6: RMSFE Ratios for a Recursive Window (Modified Stacking Estimator)

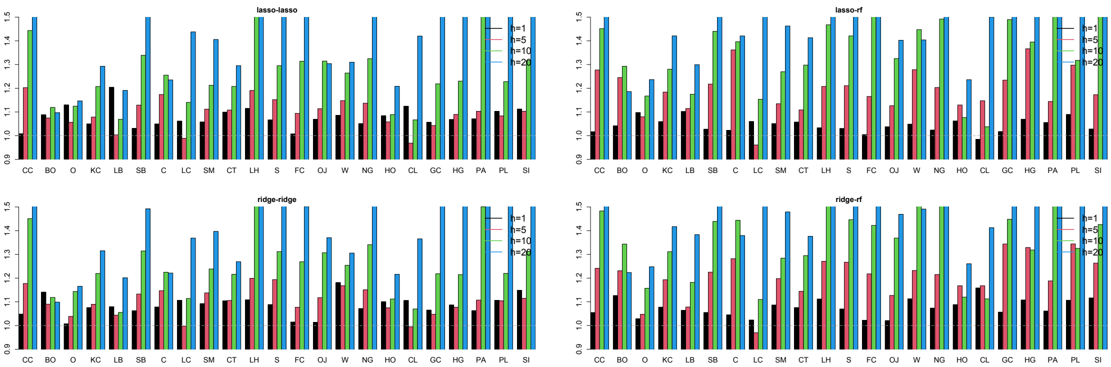

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.

Figure A7: RMSFE Ratios for a Rolling Window (Modified Stacking Estimator)

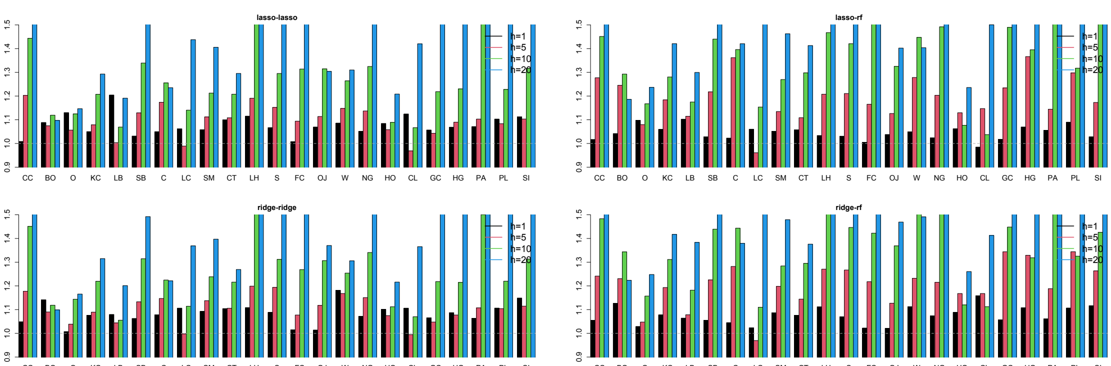

RMSFE ratios for a comparison of the HAR-RV-S model (meta learner) with the HAR-RV model. A RMSFE ratio smaller than unity indicates that the meta learner produces a smaller in-sample RMSFE than the HAR-RV model. The forecast horizons are $h = 1, 5, 10, 20$.