Realizing correlations across asset classes*

Niels S. Grønborg[a†] Asger Lunde[b]
Kasper V. Olesen[c‡] Harry Vander Elst[d§]

[a]*Department of Economics and Business Economics, Aarhus University, CREATES & The Danish Finance Institute, ngroenborg@econ.au.dk*
[b]*Department of Economics and Business Economics, Aarhus University & CREATES, alu@copenhageneconomics.com*
[c]*Department of Economics and Business Economics, Aarhus University & CREATES, mail@kasperolesen.dk*
[d]*RockSling Analytics, harry@rockslinganalytics.com*

*We thank an anonymous referee for several constructive comments and suggestions on an earlier version of the paper. The authors acknowledge support from CREATES - Center for Research in Econometric Analysis of Time Series, and thank the AU Ideas Pilot Center, Stochastic and Econometric Analysis of Commodity Markets, for financial support. This paper was written while Vander Elst was visiting CREATES. The paper has benefitted immensely from comments by members of the PhD committees of Olesen and Vander Elst, seminar participants at Aarhus University, Lund University, University of Zürich, conference participants at CFE 2014 (Pisa), Nordic Econometric Meeting 2015 (Helsinki), SoFiE 2015 (Aarhus), CFE 2015 (London), and Commodity and Emerging Market Risk 2016 (New York), Mikkel Bennedsen, Jonas N. Eriksen, Andrew Patton (discussant), and David Veredas. The usual disclaimers apply. Declaration of interests: none.

†Corresponding author, e-mail: ngroenborg@econ.au.dk, phone: (45) 8716 5285. Fuglesangs Allé 4, 8210 Aarhus V, Denmark.

‡Part of the research was done while Olesen was with Bank of America Merrill Lynch; the opinions and views expressed in this paper are those of the author and not necessarily those of his employer.

§The opinions and views expressed in this paper are those of the author, and not necessarily those of his employer. The author acknowledges financial support from a FRESH grant.

Abstract

We introduce a simple and intuitive composite model for forecasting correlations for use in portfolio optimization. Each element of the composite model is based on a realized volatility model. To test our model, we consider an investor seeking to diversify an equity portfolio by including commodities. In a high-frequency setting, we demonstrate that significant economic gains can be achieved by basing portfolio decisions on our modeling framework. The gains depend on the quality of the chosen volatility model, and for our preferred model, they are economically significant despite the realistic constraints on short selling and portfolio turnover.

Keywords: Commodities, futures markets, portfolio selection, Realized Beta GARCH.

JEL classification: C58, G11, G17

1. Introduction

The traded volume in commodity markets has increased significantly in recent years and these markets provide exciting new possibilities for portfolio diversification. Diversification has been a keyword in the financial literature since the concept of portfolio selection was formalized by Markowitz (1952), but the diversification benefits from including commodities might be different than people think. There is, for example, a widespread narrative in the popular press and among financial pundits that adding gold to a portfolio will dramatically reduce the risk of the portfolio.[1] There are certainly cases where a large reduction in portfolio variance can be achieved by adding gold to the portfolio, but in general the benefits will depend on the correlation between gold and the portfolio, as well as how it is added to the portfolio (i.e. does the investor choose a long or a short position in gold). Furthermore, diversification benefits from adding gold to a portfolio are time-varying as correlations vary over time.

In this paper, we consider an investor who holds an equity portfolio as represented by S&P 500 futures, and we explore the potential diversification benefits from including selected commodity futures contracts in the portfolio. The lesson from Markowitz (1952) is that correlations determine the optimal portfolio composition, therefore it is important for the investor to know the correlations between the returns of all assets in his or her portfolio. In case the correlations vary over time, the investor will need a model to forecast how the correlations change in order to select portfolio weights.[2] We

[1] In an interview with CNBC (December 30, 2014), Jim Cramer, the host of the TV show MAD Money, put it this way: ”I consider gold as an insurance policy.” The article explains that gold is attractive because ”it tends to go up when everything else goes down.”

[2] Gârleanu and Pedersen (2013) investigates how such a model coupled with a sophisticated trading

start out by presenting two simple examples, which serve to illustrate that correlations do indeed change over time.

The top panel of Figure 1 presents three different measures of the correlation between the return of gold futures and the return of the S&P 500 E-mini futures contract over the period from early 2007 to the end of 2014. If it is assumed that the correlation is constant, it can be estimated by the unconditional sample correlation of daily returns of the two contracts. The horizontal, gray, dashed line represents the unconditional correlation. Note, that the unconditional correlation is low and slightly positive over this sample, indicating that on average gold futures, contrary to popular beliefs, actually move in the same direction as the S&P 500 futures and that diversification benefits may be smaller than expected.[3] The assumption of constant correlation may, however, be too restrictive. The light gray, dashed-dotted line represents the estimates of the correlation obtained from a rolling window analysis, where the correlation is estimated by the sample correlation of 60 observations on daily frequency. The resulting correlation estimate varies a lot over time. This analysis suggests that the correlations are not only time-varying, but also, that correlations change sign multiple times over the sample period. We observe that the correlation can change from positive to negative several times within a single year and that changes can be quite dramatic. In 2008-2009, the correlation changes from -0.5 to 0.5 in a matter of months. This has huge implications for portfolio choice and thus diversification benefits. Finally, the dark gray, solid line represents the realized correlation of the two contracts. High-frequency data are available for many commodity futures contracts and based on in-

strategy might benefit investors.

[3]In a simple portfolio with long positions in two assets, the portfolio variance is increasing in the correlation between the assets. If the investor expects a negative correlation and the actual correlation is positive, the portfolio variance will be larger than expected.

tradaily observations, we can obtain a very accurate estimate of the true, unknown, correlation between the gold and S&P 500 futures contracts. The correlation estimated from high-frequency data is much more volatile than the correlation estimated from daily data. The question is whether the more accurate estimate will this translate into larger diversification benefits for the investor?

[Insert Figure 1 About Here]

Gold is often mentioned because of its diversification potential, but other commodities are equally interesting. The price of crude oil is often used as an indicator of the state of the world economy and could represent a big potential for diversification. As a second example, the bottom panel of Figure 1 presents the three different measures of correlation between the prices of crude oil futures and the S&P 500 E-mini futures prices. First, we note that the unconditional correlation is much higher than for gold, indicating that crude oil is in general more highly correlated with the S&P 500 than gold is. Second, there is significant time-variation in the correlation, both when estimated on a rolling window and when we consider the realized correlation. Finally, as it was the case with gold, the realized correlation has the potential to be more informative than the estimate based on the daily observations.

In this paper, we go beyond the two commodities presented in Figure 1 and consider a portfolio of S&P 500 E-mini futures and six different commodities futures in order to mimic an investor who seeks diversification benefits from the inclusion of commodities in an equity portfolio.

Empirically, correlations are not constant and portfolio selection requires models for forecasting the time-varying correlations. The availability of high-frequency data has spurred a large interest in the academic literature of modeling and forecasting financial time series. On the one hand, these forecasting models might prove valuable for investors seeking to choose portfolio weights to maximize expected return in a manner that suits their respective risk preferences, but on the other hand, the high-frequency models might require investors to rebalance their portfolios too frequently. In the real world, investors face transaction costs, as well as potential short-selling and turnover constraints in this process and might not benefit from the use of high-frequency models. Additionally, if many assets are considered, it can result in very complex problems with a large number of correlations, which can make forecasting very challenging.

In this paper, we use a composite modeling approach, which produces forecasts of a covariance matrix for many assets by combining forecasts of volatilities of individual assets and a large number of pairwise correlations. The practice of building a multivariate model from a set of bivariate models was pioneered by Ledoit et al. (2003). We will refer to this modeling style as MacGyver style modeling (see Engle, 2009). The quality of the forecast of the covariance matrix obtained from the modeling approach depends on the quality of the forecasts from the volatility models used for the individual components. We use the realized beta GARCH (RBG) model of Hansen et al. (2014) to incorporate the information from high-frequency return data.

To test the performance of the RBG model in a MacGyver style setup, we consider a realistic portfolio exercise in which an investor diversifies an all equity position by including commodity futures into the portfolio. We document impressive diversification benefits for this investor. When forming the minimum variance portfolio based on our model framework, the volatility is reduced by more than 20% compared to a portfolio consisting only of S&P 500 futures. It is a reasonable question to ask, whether this effect stems purely from including more assets in the portfolio. We demonstrate that this is not the case. This diversification benefit highly depends on the ability to produce good forecasts of the covariance matrix. Another question is whether the

RBG model is a good choice for the volatility model. We compare the performance of the portfolio based on the RBG model to the performances of portfolios based on other volatility models. We consider among others the dynamic conditional correlation GARCH (DCC) of Engle (2002) and the multivariate high-frequency-based volatility (HEAVY) models of Noureldin et al. (2012). Among the models we consider, we find that the RBG-based portfolio has the lowest variance. These results are robust to different assumptions regarding constraints on short selling and portfolio turnover.

We also investigate the portfolio performance when the investor considers a momentum strategy. We show that the RBG model in a MacGyver style model results in a portfolio, which is more desirable to the investor, than portfolios based on any of the other models, which we consider. This result holds both in terms of returns-based measures and in terms of utility-based measures. The RBG model's ability to capture volatility spillover is shown to be particularly important, when we consider the momentum strategy. Finally, we document that the MacGyver style modeling alone is not responsible for the good results as the chosen bivariate model greatly impacts the results.

The idea that commodities can offer diversification benefits to an equity investor is by no means new. Bodie and Rosansky (1980) document diversification gains for an equity investor that changes to a 60/40 position in equities and commodity futures.

Our paper is based on data for individual futures contracts while several previous studies are based on indices. Based on the index constructed by Gorton and Rouwenhorst (2006), Cheung and Miu (2010) document statistically significant diversification gains. Belousova and Dorfleitner (2012) document time variation in diversification benefits and that benefits vary across commodities. More recently, Daskalaki et al. (2017) document diversification benefits both in- and out-of-sample.

You and Daigler (2012) document diversification benefits from individual commodity futures in a relatively low frequency setting by studying weekly returns. In this paper, we rely on high frequency intraday information.

Gârleanu and Pedersen (2013) develop and apply a complex dynamic trading strategy to a sample of daily returns for individual commodity futures. The focus of that paper is on the dynamic trading strategy, whereas our focus is the development of a simple, but powerful forecasting method for the covariance matrix.

Previous studies demonstrate that features from equity markets are also present in commodity futures. Frazzini and Pedersen (2014) show that betting against beta is a profitable strategy for a portfolio consisting of, among other assets, commodity futures. Asness et al. (2013) document value and momentum patterns in commodity futures data. Finally, and most closely related to the present paper, Bollerslev et al. (2018) document strong patterns in realized volatility across asset classes and economic benefits from this information. Our paper contributes in a different direction as we show how economic benefits can be obtained by an investor using the simple MacGyver style modeling and how the benefits depend on the quality of the individual volatility models.

The remainder of this paper is organized as follows. In Section 2, we present the data set used in the analysis. In Section 3, we introduce the RBG model and the MacGyver style modeling approach. In Section 4, we present a comprehensive and realistic analysis of the economic gains from using our proposed framework. We conclude in Section 5.

2. High-frequency futures data

Our high-frequency futures data are from Tick Data Inc., and includes three metal commodity futures, three energy commodities, and one equity index. Time series of futures prices are constructed using software provided by Tick Data. As futures contracts eventually expire, each time series consists of data for multiple futures contracts. The roll-over from one contract to the next is set to occur when the next contract becomes the most traded. Gârleanu and Pedersen (2013) argue that, because this roll leaves the investor with the same net exposure as before the roll, it is reasonable to abstract from transaction costs in association with the roll. We adopt the same strategy. For more details on the futures data, see Christoffersen et al. (2019). The sample period starts January 3, 2007 and ends December 31, 2014.

We consider S&P 500 E-minis, copper, gold, silver, heating oil, light crude oil, and natural gas futures in this analysis. Table 1 presents contract-specific information for the futures we consider. All seven contracts have traded for a long period of time. Silver started trading in 1983, whereas the S&P 500 E-minis is the most recently introduced contract and started trading in 1997. All contracts are traded on the exchanges in New York and prices are in U.S. dollars. On each trading day there is a break in the trading. For the seven contracts, the breaks are relatively short and completely overlapping. Table 1 presents the average numbers of daily trades for each of the futures calculated for four different years. We focus only on trades in the most active contract. S&P 500 E-minis were traded heavily throughout the sample. In 2005, there were on average more than 1,000 daily trades in each of the commodity contracts, and in 2007 this number was at least doubled for all the contracts. Crude oil in particular experienced a remarkable capital inflow and the average number of daily trades increased from 3,962 to 50,425.[4] The average numbers of daily trades increased again for all commodities from 2007 to 2010, and from 2010 to 2013, the numbers increased again for five of the six commodities. The average number of trades for crude oil was

[4]2006 was the inception year for many large energy ETFs, such as the United States Oil Fund, which explains a large part of this development.

approximately 130,000 for both 2010 and 2013. These numbers indicate the massive inflow of capital into the seven futures contracts during our sample period.[5]

[Insert Table 1 About Here]

Selection of the six commodities is based on two criteria. First, relatively short and overlapping trading breaks, which means that we have many trades for the different contracts occurring at roughly the same time. This is very useful when considering the intraday correlations.[6] Second, a high number of trades that allows for very precise estimation of the volatilities and the correlations between the contracts. We use 2007 as the starting point for our analysis as a high number of intradaily observations are available for all contracts. We assessed the performance of high-frequency-based models relative to models using daily data. To make this comparison, it was important that the contracts we selected were traded on a frequency, which is truly high, and therefore we exclude contracts with low trading volume or long trading breaks, even though a larger number of contracts would be desirable from a diversification perspective.

We followed the process of Barndorff-Nielsen et al. (2009) to clean the daily transaction data for the selected commodities. Further details are in Appendix A.

3. Modeling and forecasting correlations

In this section, we first present the realized measures used in the analysis. Next, we present the RBG model and describe how to construct the composite covariance matrix. Finally, we briefly introduce a number of alternative volatility models.

[5] We do not report trading volumes, but the information is available from Tick Data Inc.

[6] The fact that the trading breaks are short and overlapping means that we can ignore them when constructing the realized measures.

3.1 *Realized volatility measures*

Correlations and covariances are key elements in portfolio theory, therefore realized covariances are of immense importance in this paper. Realized covariances are studied in great detail in Barndorff-Nielsen and Shephard (2004). We follow the original implementation of the estimators closely and refer to the original papers for details.

The estimation techniques we employ require prices to be observed on a homogeneous grid of time coordinates. High-frequency prices are not observed on such a grid, which implies that synchronization is required. We use previous-tick interpolation to obtain a homogeneous grid of evenly spaced prices (see Dacorogna et al., 2001).

Microstructure effects can potentially affect the properties of our estimators (see Epps, 1979). We use the multivariate realized kernel introduced by Barndorff-Nielsen et al. (2011), which is a class of estimators that are robust to measurement errors and microstructure effects induced by asynchronous trading. Importantly, non-flat-top realized kernels are used to ensure positive semi-definiteness. Additionally, we apply sparse sampling, which is shown to effectively remove the influence of microstructure effects in equity markets (e.g., Barndorff-Nielsen and Shephard, 2007). Prices are sampled at a 15-minute frequency with sub-sampling every 15 seconds. Further details on the realized measures are in Appendix A.

All the realized measures are computed using the transactions data described in the previous section. We do not consider jump robust estimators in this study. For more details on this, see, for example, Christensen et al. (2014) and Vander Elst and Veredas (2017).

3.2 Realized beta GARCH

The MacGyver modeling approach requires forecasts of all pairwise correlations and volatilities. We use the RBG model of Hansen et al. (2014) to obtain these forecasts. The RBG is a dynamic model that models the conditional covariance matrix of returns similar to other GARCH models. However, the information set is richer and includes both the realized volatilities of assets i and j on day t, denoted $x_{i,t}$ and $x_{j,t}$, respectively, and the realized correlation between the two on day t, denoted $y_{i,j,t}$. As the marginal models that are used for assets i and j are identical, we provide details for asset i only. Let $h_{i,t}$ be the conditional variance of asset i, and let $\tilde{h}_{i,t} = \log h_{i,t}$, and $\tilde{x}_{i,t} = \log x_{i,t}$. Returns for asset i, $r_{i,t}$ are modeled with the following univariate realized GARCH model:

$$\begin{aligned}
r_{i,t} &= \mu_i + e^{\tilde{h}_{i,t}/2} z_{i,t}, \\
\tilde{h}_{i,t} &= a_i + b_i \tilde{h}_{i,t-1} + c_i \tilde{x}_{i,t-1} + d_i \tilde{h}_{j,t} + \tau_{i,1} z_{i,t-1} + \tau_{i,2} \left(z_{i,t-1}^2 - 1 \right) \\
\tilde{x}_{i,t} &= \xi_i + \varphi_i \tilde{h}_{i,t} + \delta_{i,1} z_{i,t} + \delta_{i,2} \left(z_{i,t}^2 - 1 \right) + u_{i,t},
\end{aligned} \tag{1}$$

where $z_{i,t} \sim i.i.d.N\,(0,1)$ and $u_{i,t} \sim i.i.d.N\,(0, \sigma_{u_i})$ are mutually independent, and $\theta_i = (\mu_i, a_i, b_i, c_i, d_i\ , \tau_{i,1}, \tau_{i,2}, \xi_i, \varphi_i, \delta_{i,1}, \delta_{i,2}, h_{i,1})'$ is the vector of parameters in the model. In the following, we also consider a restricted version of the RBG model, where $d_i = 0$ for all assets. We refer to this model as the restricted realized beta GARCH, RBG.r. For details on the univariate realized GARCH, see Hansen et al. (2012).

To model the dynamics of the correlations we employ the following:

$$\begin{aligned}
F\left(\rho_{j,i,t}\right) &= a_{ji} + b_{ji} F\left(\rho_{j,i,t-1}\right) + c_{ji} F\left(y_{j,i,t-1}\right) \\
F\left(y_{j,i,t}\right) &= \xi_{ji} + \varphi_{ji} F\left(\rho_{j,i,t}\right) + \nu_{j,t},
\end{aligned} \tag{2}$$

where $F(\rho) = \frac{1}{2}\log\frac{1+\rho}{1-\rho}$ denotes the Fisher transform, $\theta_{j,i} = (a_{ji}, b_{ji}, c_{ji}, \xi_{ji}, \varphi_{ji}, \rho_{j,i,1})'$ is the vector of parameters in the joint model, and $v_{j,t} \sim i.i.d\left(0, \sigma_{v_j}\right)$. The last two equations in (1) and (2) are measurement equations required for the specification of the conditional density $f\left(x_{j,t}, y_{j,i,t} \left| r_{j,t}, r_{i,t}, x_{i,t}, \mathcal{F}_{t-1}\right.\right)$. The measurement errors $u_{j,t}$ and $v_{j,t}$ are assumed independent of $z_{i,t}$ and $z_{j,t}$ but allowed to be mutually correlated, as follows:

$$\mathbf{\Sigma} = \text{Var}\begin{bmatrix} u_{i,t} \\ u_{j,t} \\ v_{j,t} \end{bmatrix} = \begin{bmatrix} \sigma^2_{u_i} & \sigma_{u_i,u_j} & \sigma_{u_i,v_j} \\ \bullet & \sigma^2_{u_j} & \sigma_{u_j,v_j} \\ \bullet & \bullet & \sigma^2_{v_j} \end{bmatrix}. \tag{3}$$

We obtain estimates using the maximum likelihood estimation approach outlined in Hansen et al. (2014). Forecasting is also outlined in Hansen et al. (2014) and we can obtain k-step ahead forecasts of conditional variances of assets i and j and their conditional correlation. These are denoted $h_{i,t+k|t}$, $h_{j,t+k|t}$, and $\rho_{i,j,t+k|t}$, respectively. Details for estimation and forecasting are in Appendix B.

3.3 *MacGyver style model for the covariance matrix*

We combine the forecasts of all the pairwise correlations into a forecast of the full covariance matrix, $H_{t+k|t}$, for the n assets. We construct this matrix using the techniques presented in Lunde et al. (2016). This composite covariance matrix has the typical element:

$$H_{i,j,t+k|t} = \rho_{i,j,t+k|t}\sqrt{h_{i,t+k|t}h_{j,t+k|t}}.$$

Several methods can be applied to ensure that $H_{i,j,t+k|t}$ is positive definite. We rely on eigenvalue cleaning as explained in Hautsch et al. (2012) and Lunde et al. (2016) to ensure positive definiteness. An alternative would be to solve the problem by applying a short-selling constraint, as discussed in Jagannathan and Ma (2003).[7]

The elementwise construction of the covariance matrix helps us avoid the challenges arising from the estimation of the full covariance matrix (e.g., Hayashi and Yoshida, 2005; Aït-Sahalia et al., 2010; Christensen et al., 2010; Zhang, 2011; Fan et al., 2012; Bibinger et al., 2014; Engle et al., 2019).

3.4 Volatility models

We produce forecasts of the covariance matrix based on the RBG model and we are interested in the model's implications for portfolio selection. We compare the performance of the resulting portfolio to the performances of portfolios based on nine other forecasting models of the covariance matrix.

The competing models can be classified into two groups, where models in the first group rely on daily information and models in the second group, like the RBG model, rely on daily information and realized measures.

3.4.1 Models with daily information only

The first model is very simple and heavily used in the industry. The forecast of the covariance matrix is simply the covariance matrix of the daily returns calculated based on a moving 60-day window. We call the model RW.r^2. In the RW.r^2 model, we place the same weight on observations from the previous day as on observations from 60 days ago. This might be too restrictive. Therefore, we also consider a 60-day rolling window model, where the observations are weighted according to the RiskMetrics model (see Mina and Xiao, 2001). We call this model RM.r^2 and choose an exponential decay rate of 0.97. The third model is the dynamic conditional correlation GARCH of Engle (2002), which we refer to as the DCC model. We also consider a version of the

[7]This can potentially be problematic if a scenario with no short selling is analyzed

DCC model, which is based on the MacGyver approach, where the covariance matrix is constructed from covariance forecasts from bivariate models. We refer to this as the DCC.2 model.

3.4.2 Models with daily information and realized measures

The first two models we consider are very simple and based exclusively on the multivariate realized kernel. The first, RW.RK, is based on the realized kernel based on a 60-day rolling window. Likewise, the RM.RK is based on a 60-day rolling window, where observations are weighted with an exponentially decaying function. We also consider the scalar and the diagonal versions of the HEAVY models of Noureldin et al. (2012). We denote the scalar version HEAVY.sc, and the diagonal version HEAVY.dg.

As mentioned in Subsection 3.2, we also consider a restricted version of the RBG model, which we call RBG.r. The two HEAVY models, HEAVY.sc and HEAVY.dg, rely on the same methodology to construct a composite covariance matrix from pairwise correlations as is used for the RBG and RBG.r models.

4. Performance evaluation and economic gains

The true test of our forecasts of the covariance matrix is whether this forecast will allow an investor to make better decisions with respect to portfolio selection. We consider an investor who chooses $n = 7$ portfolio weights for a portfolio consisting of six different commodity futures and the S&P 500 E-mini futures. The investor follows a dynamic trading strategy, where a forecast of the covariance matrix for the next day, $H_{t+1|t}$, is obtained daily and the portfolio weights, w_t, are adjusted accordingly.[8] We

[8]Weekly or monthly rebalancing could also be considered but the daily evaluation allows for the longest evaluation period. Results for weekly and monthly rebalancing are in the Internet appendix.

assume that the investor has to obey very general restrictions on short selling and is limited to a certain amount of turnover. The investor faces the following problem:

$$\min_{w_t} w_t' H_{t+1|t} w_t$$

$$s.t. \quad w_t' \iota = 1, \tag{4}$$

$$\left| w_t' \right| \iota \leq 1 + 2s, \tag{5}$$

$$|w_{kt}| \leq \bar{w}, \quad k = 1, \ldots, n \tag{6}$$

$$w_t' \mu_t \geq \mu_0, \tag{7}$$

$$TO_t \leq \delta, \tag{8}$$

where ι is an $n \times 1$ vector of ones. Equation (4) ensures that portfolio weights sum to one. Equation (5) is a short-selling constraint, where s determines the percentage of short positions allowed as presented in Lunde et al. (2016). If $s = 0$, no short selling is allowed. Equation (6) ensures that no single asset gets too large a weight in the portfolio. Equation (7) specifies a minimum required return of the portfolio, given the vector of expected returns, μ_t. Equation (8) constrains the turnover in the portfolio, where turnover is defined as:

$$TO_t = \sum_{i=1}^{n} \left| w_{i,t+1} - w_{i,t} \left(\frac{1 + r_{i,t+1}}{1 + r_{p,t+1}} \right) \right|. \tag{9}$$

Alternatively, the investor might choose an equal-weighted portfolio, also called a $1/n$ portfolio. This portfolio has been documented to have good performance (DeMiguel et al., 2009), and we include it as a benchmark in our analysis.

The portfolio exercise is repeated following a rolling window scheme, where the models are estimated based on an estimation window and the portfolio is formed out-of-sample. Estimation is based on 750 days, meaning that the first portfolio return is realized on December 17, 2009 and the last on December 31, 2014. This leaves us with 1,294 daily portfolio returns in our evaluation period.

4.1 The minimum variance portfolio

The global minimum variance portfolio is often used for economic evaluation in the literature. As argued by Ledoit and Wolf (2018) and Engle et al. (2019), it presents a clean problem, in the sense that the performance of a model for the covariance matrix is not influenced by the estimation of expected returns. In addition to the portfolio volatility, the out-of-sample Sharpe ratio of the minimum variance portfolio has been analysed by Haugen and Baker (1991), Jagannathan and Ma (2003), Nielsen and Aylursubraminian (2008), Ledoit and Wolf (2018), and Engle et al. (2019). Finally, Ledoit and Wolf (2018) and Engle et al. (2019) note that mutual funds are now offering global minimum variance products. In this analysis, we do not consider the global minimum variance portfolio in its strictest sense, as we impose various specifications of the short-selling constraint in Equation (5).[9] As the investor in the classical minimum variance problem is not constrained in terms of turnover, we choose a very high δ to make sure that Equation (8) is non-binding. We ignore the constraint in Equation (7).

For each day in the out-of-sample evaluation period, portfolio weights, w_t, are constructed based on the covariance forecasts from each model. We obtain time series of portfolio returns for all the models, and we base the evaluation of the models on the realized volatility, $\sigma_{p,t+1} = \sqrt{w_t' RC_{t+1} w_t}$, of the returns over the evaluation period. RC_{t+1} is the realized covariance for the seven assets for the period from t to $t+1$ and defined formally in Appendix A. We present the average realized portfolio volatility, $\bar{\sigma}_p$, as well as the ratio of the average of the relalized portfolio volatility to the aver-

[9]We also deviate from the pure minimum variance portfolio by choosing $\bar{w} = 0.5$ and thereby the investor might face a binding condition in Equation (6).

age of the realized volatility of a portfolio consisting exclusively of S&P 500 futures, $VR = \bar{\sigma}_p / \bar{\sigma}_{S\&P}$.

Comparing the means of realized volatilities for eleven different models is problematic as it involves a multiple comparison problem. We use the model confidence set (MCS) of Hansen et al. (2011) to compare the time series of realized volatilities from the different models. The MCS provides a set of models, which includes the best model (the lowest realized volatility) with a given probability. The size of this set is data dependent and it could contain all the models or any other (non-empty) subset of the models, even just a single model. The MCS results are in the form of a set of p-values, one for each model. Low p-values, say below 0.1, indicate that the corresponding model can be excluded from the set of the best models. We denote the p-values estimated based on realized volatilities as $p_{MCS}(\sigma_{p,t})$.

Table 2 presents the average annualized realized portfolio volatilities $\bar{\sigma}_p$, along with the MCS p-values, $p_{MCS}(\sigma_{p,t})$, for the eleven models. The results are presented for three different short-selling constraints: $s = 0\%$, $s = 25\%$, and $s = 50\%$. These portfolios are not based on any turnover constraints. The investor's sole objective is to minimize volatility. Other features of the portfolio, such as average return or Sharpe ratio (SR), are therefore incidental. Such measures do, however, help the reader to understand the portfolio better. Average return and SR are included in Table 2 for completeness, but we focus the discussion on the volatility measures.

The results in Table 2 show that the RBG and RBG.r models perform very well. The two models lead to the lowest realized portfolio volatilities and for no short-selling ($s = 0\%$) they are both included in the MCS as the only models. When short selling is allowed ($s = 25\%$ and $s = 50\%$), the RBG.r model is the only model in the MCS based on a 10% level. For all models, except the $1/n$ portfolio and for all three levels of short selling, the variance ratios are below one, indicating that there are benefits from diversification compared to a portfolio consisting only of S&P 500 futures contracts. The DCC model is the best performer of the models that rely on daily information only. Finally, we see that the $1/n$ portfolio leads to the worst performance in this case, where we only consider the variance. The variance ratio is higher than one, which illustrates that simply adding commodities to a portfolio is not enough to obtain diversification benefits. In order to obtain diversification benefits, the covariance matrix must be estimated. The results also highlight that the MacGyver style modeling alone is not driving the results. The HEAVY.sc, HEAVY.dg, and DCC.2 models all use the MacGyver style modeling approach, but they underperform compared to the RBG and RBG.r models. The fact that the RBG and RBG.r models deliver similar results indicate that modeling the volatility spillover is not very important in this particular exercise.

[Insert Table 2 About Here]

4.1.1 *Short selling and turnover*

The results in Table 2 are not subject to any turnover constraints.[10] In Table 3, we present the results for different combinations of s and δ to investigate effects of short selling and turnover.

Table 3 shows that the RBG model performs very well when the investor has the flexibility to act according to the model's predictions. Portfolio volatility is generally decreasing in both the amount of turnover and short selling allowed. It is worth noting that for extremely strict turnover restrictions, we see very high performance in terms of average return. We believe that this is a coincidence and that it is influenced heavily by the initial position in the different assets. The initial position will play a very important role when the investor cannot change her portfolio very much.[11]

[10] $\delta = 100$ such that Equation (8) is non-binding.

[11] Note, that the initial position is chosen based on data from 2007 to 2009, which is a peculiar period

[Insert Table 3 About Here]

4.1.2 *Time variation of relative performance of different models*

In Goyal and Welch (2008), plots of cumulative errors are used to analyze performance of different return prediction models and to pin point periods in which a particular model performs very well or very poorly. We use a very similar graphical device to compare the performance of portfolios based on the forecasts from different models. We use this approach to compare the performance of a portfolio, when considering different short-selling constraints. We plot the cumulative realized volatility of a portfolio, $\sigma_{p,t,cum} = \sum_{\tau=1}^{t} \sigma_{p,\tau}$, over time. The cumulative realized portfolio volatility is not an intuitive number in its own right, but it allows us to compare the performance of different portfolios. We augment the graphical approach of Goyal and Welch (2008) with a second plot, with which we analyze the performance within a particular year.

In Figure 2, we present the cumulated volatility for the minimum variance portfolio. The figure presents results for the minimum variance portfolio based on three different models, the RBG model, the DCC model, and the HEAVY.dg model. The bottom panel additionally includes a portfolio consisting exclusively of S&P 500 futures contracts. We choose the DCC model as it results in the lowest volatility of all the models using daily data. We choose the HEAVY.dg model as it results in the lowest volatility of the high-frequency models based on the MacGyver approach. The portfolios are formed without short-selling constraints and are not subject to any turnover constraints. The bottom plot highlights the relative performance of the portfolios within the individual years by resetting the cumulative volatility every year.

The top panel of Figure 2 shows that a lower portfolio volatility can be obtained by using the RBG model. The cumulated volatility of the portfolio based on the RBG

for financial markets.

model increases more slowly than the volatilities of the portfolios based on the competing models throughout the sample. The bottom panel shows that the good performance of the RBG model is not caused by any particular subperiod with good performance, but that it consistently leads to lower volatility year after year. We can also see that including commodities in the portfolio leads to a lower volatility than for a portfolio consisting only of S&P 500 futures. The diversification benefits are clear in every year in our evaluation sample; in 2010 and 2011, the improvements are quite dramatic.

[Insert Figure 2 About Here]

The results in Figure 2 are based on an assumption of no limits to short selling. In Figure 3, we investigate the effects of short-selling constraints on the volatility of the portfolio based on the RBG model. The top panel of Figure 3 shows that the diversification benefits are largest when the investor has some flexibility in terms of taking short positions. Interestingly, we see that the portfolio volatility is essentially identical for $s = 25\%$ and $s = 50\%$, indicating that the constraint is non-binding for a large part of the evaluation period. In the bottom panel, we see that the effects of allowing short selling are largest in the first three years of the evaluation period, whereas the three volatilities are essentially the same in the last two years.

[Insert Figure 3 About Here]

We conclude, that if the objective of the investor is to minimize the volatility of the portfolio, then commodities should be included in the portfolio. The covariance matrix, which is used to find the portfolio weights, should be based on either the RBG or the RBG.r model. These results are robust to various constraints regarding short selling and turnover.

4.1.3 Squared returns and investment horizon

The results in Table 2, Table 3, Figure 2, and Figure 3 are all based on realized portfolio volatility, $\sigma_{p,t+1} = \sqrt{\omega_t' RC_{t+1} \omega_t}$. In the Internet appendix, we present the results of Table 2, where the portfolio variance is estimated by daily squared returns. The only difference is how we measure the volatility of the resulting portfolio.

In general, we find that the models using daily data perform best when portfolio volatility is estimated based on daily returns, and that the high-frequency models perform best when portfolio volatility is based on the more precise estimate from high-frequency returns.

In our analysis, we assume that the portfolio is rebalanced daily. In the Internet appendix we include two additional versions of Table 2. One version rebalances weekly and one rebalances monthly. For the monthly investment horizon, we see that most of the high-frequency-based models are included in the MCS.

4.2 Momentum strategy

To implement any strategy, where Equation (7) is considered, one must specify a model for the expected return vector, μ_t, and specify a target return μ_0. We follow Ledoit and Wolf (2018) and Engle et al. (2019) and use the momentum factor of Jagadeesh and Titman (1993) to specify μ_t and choose a value for μ_0.[12] We choose the target return to be the average of the momentum for all the assets, $\bar{\mu}$, if this quantity is positive, that is $\mu_0 = \max(0, \bar{\mu})$.

[12] At time t, we define momentum for asset i as the geometric average of the previous 12 monthly returns on the asset, but exclude the most recent month.

4.2.1 Financial evaluation

For each model, we solve the portfolio problem for different values of s and δ and obtain daily returns for the momentum portfolio. Following Ledoit and Wolf (2018) and Engle et al. (2019), we present, for each model, the out-of-sample Sharpe ratio, SR, calculated as the ratio of the the average out-of-sample return and the out-of-sample standard deviation of the returns. We also compute the Sharpe ratio of a portfolio consisting exclusively of S&P 500 futures, $SR_{S\&P}$, and present the ratio of $SR/SR_{S\&P}$. The limitations of the Sharpe ratio, as presented in, for example, Marquering and Verbeek (2004), Colacito and Engle (2006), and King et al. (2010) are well understood and an economic evaluation of model performance cannot be based on this ratio alone.

In order to quantify the difference to the $1/n$ portfolio, we present the mean absolute deviation from equal-weighted portfolio weights, MAD, which is computed as:

$$MAD = \frac{1}{n}\sum_{i=1}^{n}\left|w_i - \frac{1}{n}\right|.$$

To assess the economic significance of our proposed framework, we follow the tradition of West et al. (1993), Fleming et al. (2001), Rime et al. (2010), and Karstanje et al. (2013) and assume quadratic utility, which allows us to further evaluate the economic performance of each model. Assume that the wealth, W, of the investor evolves according to the following:

$$W_{t+1} = W_t\left(1 + r_{p,t+1}\right),$$

then, under the assumption of quadratic utility, the utility of the investor at the end of period $t+1$ is:

$$U\left(W_{t+1}\right) = W_{t+1} - \frac{\rho}{2}W_{t+1}^2 = W_t\left(1 + r_{p,t+1}\right) - \frac{\rho}{2}W_t\left(1 + r_{p,t+1}\right)^2,$$

where ρ is a parameter controlling the risk preferences of the investor. The amount of wealth invested each period is assumed to be constant, meaning the ρ relates to the relative risk aversion, γ, as $\gamma = \frac{\rho W}{1-\rho W}$. The initial wealth, W_0, is set to 1. For this investor, the average realized utility is:

$$U = \frac{1}{T}\sum_{t=0}^{T-1} U_{t+1} = \frac{1}{T}\sum_{t=0}^{T-1}\left(1 + r_{p,t+1} - \frac{\gamma}{2(1+\gamma)}\left(1 + r_{p,t+1}\right)^2\right).$$

For a discussion of average utility and expected utility, see West et al. (1993).

Methods based on realized measures often lead to very high turnover in the portfolio. For each forecasting method, we present the average turnover, TO, defined as the average over time of the expression in Equation (9). While turnover is interesting in its own right, it is also very important in the context of transaction costs. Transaction costs have been analyzed by Marquering and Verbeek (2004), Han (2006), King et al. (2010), and Karstanje et al. (2013). In the presence of fixed proportional transaction costs, τ, the average realized utility can be expressed as:

$$\frac{1}{T}\sum_{t=0}^{T-1}\left(1 + r_{p,t+1} - \tau TO_{t+1} - \frac{\gamma}{2(1+\gamma)}\left(1 + r_{p,t+1} - \tau TO_{t+1}\right)^2\right).$$

Specifically, we can define τ^{be}, which is the break-even transaction cost. That is, τ^{be} is the proportional cost per period, which eliminates any utility from using a specific model to construct portfolio weights. If transaction costs are higher than τ^{be}, the strategy is too expensive and should not be undertaken.

We compare the realized utility from the different forecasting models by applying the MCS. We estimate the MCS based on the time series of realized utility from each model, and denote the corresponding p-values, $p_{MCS}(U_t)$.

In Table 4, we present average return, average volatility, SR, $SR/SR_{S\&P}$, MAD, TO, τ^{be}, and $p_{MCS}(U_t)$ for momentum strategies based on the eleven different models.

When forming the portfolios, we allow for some degree of short selling, $s = 0.5$, and we do not impose any turnover restrictions.

The results in the column labelled SR in Table 4 show that using the RBG model leads to a much higher annualized Sharpe ratio than any of the other models. The results in the column labelled $\bar{\sigma}_p$ show that only the RBG model leads to higher Sharpe ratios than a portfolio consisting of only S&P 500 futures. The equal-weighted portfolio performs worst of all the models. The large difference in performance between the RBG model and the RBG.r model indicates that the volatility spillover effect is particularly important in this context. The question is, whether the good performance comes at the price of too high turnover in the portfolio. The results in the column labelled TO show that turnover is indeed highest for the RBG model. The results in the column labelled τ^{be} indicate, however, that the average realized utility from using the RBG model will remain positive for relatively high levels of trading costs. The break-even transaction costs are generally very high compared to the transaction costs considered by Marquering and Verbeek (2004) for equities. Finally, the results in the column labelled $p_{MCS}(U_t)$ show that the RBG model leads to the highest realized utility. When realized utilities are compared using the MCS, the RBG model is the only model in the set.

[Insert Table 4 About Here]

4.2.2 Short selling and turnover

The results in Table 4 are based on the assumption that $s = 0.5$, indicating that some degree of short selling is allowed. The results are based on the further assume that the turnover constraint is not binding, δ is high. In Table 5, we investigate the effects of these assumptions and consider the out-of-sample Sharpe ratio of the momentum portfolio based on the RBG model for different combinations of s and δ.

The results in Table 5 show that to benefit from a better model for the covariance matrix, the investor must be allowed a certain turnover in the portfolio and also some degree of short selling. Not surprisingly, we find the best performance when turnover and short-selling constraints are quite liberal. Interestingly the RBG model can perform decently, outperforming the S&P 500 portfolio (average annual SR of 0.82) even when no short selling is allowed, as long as limits to turnover are not too strict. In fact, turnover restrictions seem to matter more than short selling restrictions as the model is struggling to perform for high values of s, when δ is low. Note, that for very low values of δ, we actually see good performance. As discussed above, we believe this performance to be incidental and highly influenced by the investor's initial position in the different assets. If the investor cannot change her portfolio very much, then the initial position becomes very influential.

[Insert Table 5 About Here]

4.2.3 Performance fee

Finally, we can use the average realized utility to compare two models, say m_1 and m_2. This approach is also considered by Karstanje et al. (2013). We include fixed costs in every period equal to ϕ in the average realized utility for the returns based on m_2. By equating the average realized utility of the two models, we can evaluate how much the investor is willing to pay to switch from m_1 to m_2. We interpret ϕ as a performance fee and choose it to ensure that:

$$\sum_{t=0}^{T-1} \left(1 + r_{p,t+1}^{m_1} - \frac{\gamma}{2(1+\gamma)} \left(1 + r_{p,t+1}^{m_1}\right)^2\right) = \sum_{t=0}^{T-1} \left(1 + r_{p,t+1}^{m_2} - \phi - \frac{\gamma}{2(1+\gamma)} \left(1 + r_{p,t+1}^{m_2} - \phi\right)^2\right), \tag{10}$$

where $r^{m_1}_{p,t+1}$ and $r^{m_2}_{p,t+1}$ denote the returns at time $t+1$ from the portfolio based on forecasts from models m_1 and m_2, respectively. A positive value of ϕ will mean that an investor will pay to use m_2 in stead of m_1. A negative value of ϕ indicates that an investor requires compensation to use m_2 in stead of m_1.

Table 6 presents the performance fees, which an investor, with preferences specified as above, would be willing to pay to switch from a model in the rows to a model in the columns. These results are based on the momentum strategy where some short selling is allowed, $s = 0.5$, and there are no turnover constraints.

The first row of Table 6 contains only negative numbers. This means that an investor with the stated preferences would have to be compensated to use any other model than the RBG model. The compensation ranges from 4 to 8 bps per day and thus represents a significant quantity in an economic sense. Again we find evidence that the $1/n$ model is inferior to all other models as the last column contains only negative numbers.

[Insert Table 6 About Here]

4.2.4 *Portfolio value*

The results for the momentum portfolio presented above are largely based on time averages over the evaluation period, but it is also interesting to investigate how the value of the portfolio has evolved over time. Based on the vector of daily returns for the futures contracts, r_{t+1}, and the portfolio weights at time t, ω_t, we can calculate the portfolio return at time $t+1$ as $r_{p,t+1} = \omega_t' r_{t+1}$. In Figure 4, we plot the cumulated portfolio value, $r_{p,t,cum} = \sum_{\tau=1}^{t} r_{p,\tau}$. The sub-plots illustrate the relative performance within a given year. The plot includes values of portfolios based on three different models, the RBG, the DCC model, and the HEAVY.dg model, along with a portfolio consisting of S&P 500 futures. The DCC and HEAVY.dg models are chosen as they are the best performing alternative models based on daily returns and high-frequency returns, respectively.

The top panel of Figure 4 shows that the RBG model outperforms the existing models during our entire evaluation period. Interestingly, the HEAVY.dg and DCC models struggle to outperform the S&P 500 portfolio, indicating that good covariance forecasts are extremely important for managing portfolios and obtaining benefits from diversification. In the bottom panel, the performance of the RBG model as compared to the S&P 500 portfolio is particularly strong in 2011. There is evidence of diversification benefits in all years except 2013, where an investor would have been better off by holding S&P 500 futures than by including commodity futures in her portfolio using our approach.

[Insert Figure 4 About Here]

4.2.5 Squared returns and investment horizon

In the Internet appendix present results for the case, where portfolio volatility is measured based on squared returns. In this case, the RBG model also outperforms all other models. The Internet appendix also presents results for weekly and monthly rebalancing. Including commodities in this portfolio generally does not improve the SR in this scenario. We believe that this is because the performance signal in Equation (7) performs poorly multiple steps ahead. A simple strategy, where the investor invests only in S&P 500 futures does not require the investor to model expected returns and is therefor unaffected by the poor quality of such a model. Another explanation is that the models require more frequent rebalancing than allowed here to perform well. We do not believe that this is the problem, as the results for the minimum variance portfolio were not too affected by this.

5. Conclusion

In this paper, we employ a MacGyver style modeling approach based on the realized beta GARCH model to forecast a covariance matrix between six commodity futures contracts and a S&P 500 futures contract. We demonstrate that an investor can obtain significant diversification benefits from including commodities in an equity portfolio if forecasts of the covariance matrix are based on our methodology. The results are robust to various specifications of short-selling and turnover constraints. If the investor considers a momentum strategy, we show that using our method leads to the most valuable portfolio as long as the investor is allowed a certain turnover in the portfolio. The results are robust to the introduction of trading costs. Interestingly, our approach clearly outperforms an equal-weighted portfolio in all the main exercises we consider here. This means that diversification benefits are not simply a results of including more assets, but are also heavily affected by our ability to forecast covariances. The realized beta GARCH model outperforms other models using the MacGyver style modeling approach, meaning that this approach alone is not sufficient to achieve the economically meaningful gains, which we document. Finally, we find that it is very important to allow for volatility spillover between assets when considering a momentum strategy, but less important when the only objective is to minimize portfolio volatility.

References

Aït-Sahalia, Yacine, Jianqing Fan, and Dacheng Xiu, 2010, High-frequency covariance estimates with noisy and asynchronous financial data, *Journal of the American Statistical Association* 105, 1504–1517.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579–625.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *The Journal of Finance* 58, 929–985.

Barndorff-Nielsen, Ole E., Peter R. Hansen, Asger Lunde, and Neil Shephard, 2009, Realised kernels in practice: Trades and quotes, *Econometrics Journal* 12, C1–C32.

Barndorff-Nielsen, Ole E., Peter R. Hansen, Asger Lunde, and Neil Shephard, 2011, Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading, *Journal of Econometrics* 162, 149–169.

Barndorff-Nielsen, Ole E., and Neil Shephard, 2004, Econometric analysis of realised covariation: High-frequency based covariance, regression, and correlation in financial economics, *Econometrica* 72, 885–925.

Barndorff-Nielsen, Ole E., and Neil Shephard, 2007, Variation, jumps, and high-frequency data in financial econometrics, in Richard Blundell, Torsten Persson, and Whitney K Newey, eds., *Advances in Economics and Econometrics. Theory and Applications, Ninth World Congress*, Econometric Society Monographs, 328–372 (Cambridge University Press).

Belousova, Julia, and Gregor Dorfleitner, 2012, On the diversification benefits of commodities from the perspective of euro investors, *Journal of Banking and Finance* 36, 2455–2472.

Bibinger, Markus, Nikolaus Hautsch, Peter Malec, and Markus Reiss, 2014, Estimating the quadratic covariation matrix from noisy observations: Local method of moments and efficiency, *Annals of Statistics* 42, 1312–1346.

Bodie, Z., and V.I. Rosansky, 1980, Risk and return in commodity futures, *Financial Analyst Journal* 36, 27–39.

Bollerslev, Tim, Benjamin Hood, John Huss, and Lasse Heje Pedersen, 2018, Risk everywhere: Modeling and managing volatility, *The Review of Financial Studies* 31, 2730–2773.

Cheung, C. Sherman, and Peter Miu, 2010, Diversification benefits of commodity futures, *Journal of International Financial Markets, Institutions & Money* 20, 451–474.

Christensen, Kim, Silja Kinnebrock, and Mark Podolskij, 2010, Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data, *Journal of Econometrics* 159, 116–133.

Christensen, Kim, Roel Oomen, and Mark Podolskij, 2014, Fact or friction: Jumps at ultra high-frequency, *Journal of Financial Economics* 114, 576–599.

Christoffersen, Peter, Asger Lunde, and Kasper Olesen, 2019, Factor structure in commodity futures return and volatility, *Journal of Financial and Quantitative Analysis* 54, 1083–1115.

Colacito, Riccardo, and Robert F. Engle, 2006, Testing and valuing dynamic correlations for asset allocation, *Journal of Business & Economic Statistics* 24, 238–253.

Dacorogna, Michel M., Ramazan Gencay, Ulrich Müller, Richard B. Olsen, and Olivier V. Pictet, 2001, *An introduction to high-frequency finance* (Academic Press, London).

Daskalaki, Charoula, George Skiadopoulos, and Nikolas Topaloglou, 2017, Diversification benefits of commodities: A stochastic dominance efficiency approach, *Journal of Empirical Finance* 44, 250–269.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal, 2009, Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?, *Review of Financial Studies* 22, 1915–1953.

Engle, Robert F., 2002, Dynamic conditional correlation - A simple class of multivariate GARCH models, *Journal of Business & Economic Statistics* 20, 339–350.

Engle, Robert F., 2009, High dimensional dynamic correlations., in J.L. Castle, and N. Shephard, eds., *The Methodology and Practice of Econometrics: Papers in Honour of David F Hendry*, 122–148 (Oxford University Press).

Engle, Robert F., Olivier Ledoit, and Michael Wolf, 2019, Large dynamic covariance matrices, *Journal of Business & Economic Statistics* 37, 363–375.

Epps, Thomas W., 1979, Comovements in stock prices in the very short run, *Journal of the American Statistical Association* 74, 291–298.

Fan, Jianqing, Yingying Li, and Ke Yu, 2012, Vast volatility matrix estimation using high-frequency data for portfolio selection, *Journal of the American Statistical Association* 107, 412–428.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2001, The economic value of volatility timing, *Journal of Finance* 56, 329–352.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gârleanu, Nicolae, and Lasse Heje Pedersen, 2013, Dynamic trading with predictable returns and transaction costs, *The Journal of Finance* 58, 2309–2340.

Gorton, Gary, and K. Geert Rouwenhorst, 2006, Facts and fantasies about commodity futures, *Financial Analyst Journal* 62, 47–68.

Goyal, Amit, and Iwo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *The Review of Financial Studies* 21, 1455–1508.

Han, Yufeng, 2006, Asset allocation with a high dimensional latent factor stochastic volatility model, *Review of Financial Studies* 19, 237–271.

Hansen, Peter R., Zhuo Huang, and Howard H. Shek, 2012, Realized GARCH: A joint model of returns and realized measures of volatility, *Journal of Applied Econometrics* 27, 877–906.

Hansen, Peter R., Asger Lunde, and James M. Nason, 2011, The model confidence set, *Econometrica* 79, 456–497.

Hansen, Peter R., Asger Lunde, and Valeri Voev, 2014, Realized beta GARCH: A multivariate GARCH model with realized measures of volatility, *Journal of Applied Econometrics* 29, 774–799.

Haugen, Robert A., and Nardin L. Baker, 1991, The efficient market inefficiency of capitalization-weighted stock portfolios, *The Journal of Portfolio Management* 17, 35–40.

Hautsch, Nikolaus, Lada M. Kyj, and Roel C. A. Oomen, 2012, A blocking and regularization approach to high-dimensional realized covariance estimation, *Journal of Applied Econometrics* 27, 625–645.

Hayashi, Takaki, and Nakahiro Yoshida, 2005, On covariance estimation of non-synchronously observed diffusion processes, *Bernoulli* 11, 359–379.

Jagadeesh, Narasimhan, and Sheridan Titman, 1993, Return to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65–91.

Jagannathan, Ravi, and Tongshu Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints help, *Journal of Finance* 54, 1651–1684.

Karstanje, Dennis, Elvira Sojli, Wing Wah Tham, and Michel van der Wel, 2013, Economic valuation of liquidity timing, *Journal of Banking & Finance* 37, 5073–5087.

King, Michael, Lucio Sarno, and Elvira Sojli, 2010, Timing exchange rates using order flow: the case of the loonie, *Journal of Banking & Finance* 34, 2917–2928.

Ledoit, Olivier, Pedro Santa-Clara, and Michael Wolf, 2003, Flexible multivariate GARCH modeling with an application to international stock markets, *Review of Economics & Statistics* 85, 735–747.

Ledoit, Olivier, and Michael Wolf, 2018, Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks, *The Review of Financial Studies* 30, 4349–4388.

Lunde, Asger, Neil Shephard, and Kevin Sheppard, 2016, Econometric analysis of vast covariance matrices using composite realized kernels and their application to portfolio choice, *Journal of Business and Economics Statistics* 34, 504–518.

Markowitz, Harry, 1952, Portfolio selection, *Journal of Finance* 7, 77–91.

Marquering, Wessel, and Marno Verbeek, 2004, The economic value of predicting stock index and volatility, *Journal of Financial and Quantitative Analysis* 39, 407–429.

Mina, Jorge, and Jerry Y. Xiao, 2001, Return to Riskmetrics: The evolution of a standard, *Technical Report, Risk Metrics Group* .

Nielsen, Frank, and Raman Aylursubraminian, 2008, Far from the madding crowd - volatility efficient indices, Research Insights, MSCI Barre.

Noureldin, Diaa, Neil Shephard, and Kevin Sheppard, 2012, Multivariate high-frequency-based volatility (HEAVY) models, *Journal of Applied Econometrics* 27, 907–933.

Rime, Dagfinn, Lucio Sarno, and Elvira Sojli, 2010, Exchange rate forecasting, order flow and macroeconomic information, *Journal of International Economics* 80, 72–88.

Vander Elst, Harry, and David Veredas, 2017, Smoothing it out: Empirical and simulation results for disentangled realized covariances, *Journal of Financial Econometrics* 15, 106–138.

West, Kenneth D., Hali J. Edison, and Dongchul Cho, 1993, A utility based comparison of some models for exchange rate volatility, *Journal of International Economics* 35, 23–45.

You, Leyuan, and Robert T. Daigler, 2012, A Markowitz optimisation of commodity futures portfolios, *Journal of Futures Markets* 33, 343–368.

Zhang, Lan, 2011, Estimating covariation: Epps effect and microstructure noise, *Journal of Econometrics* 160, 33–47.

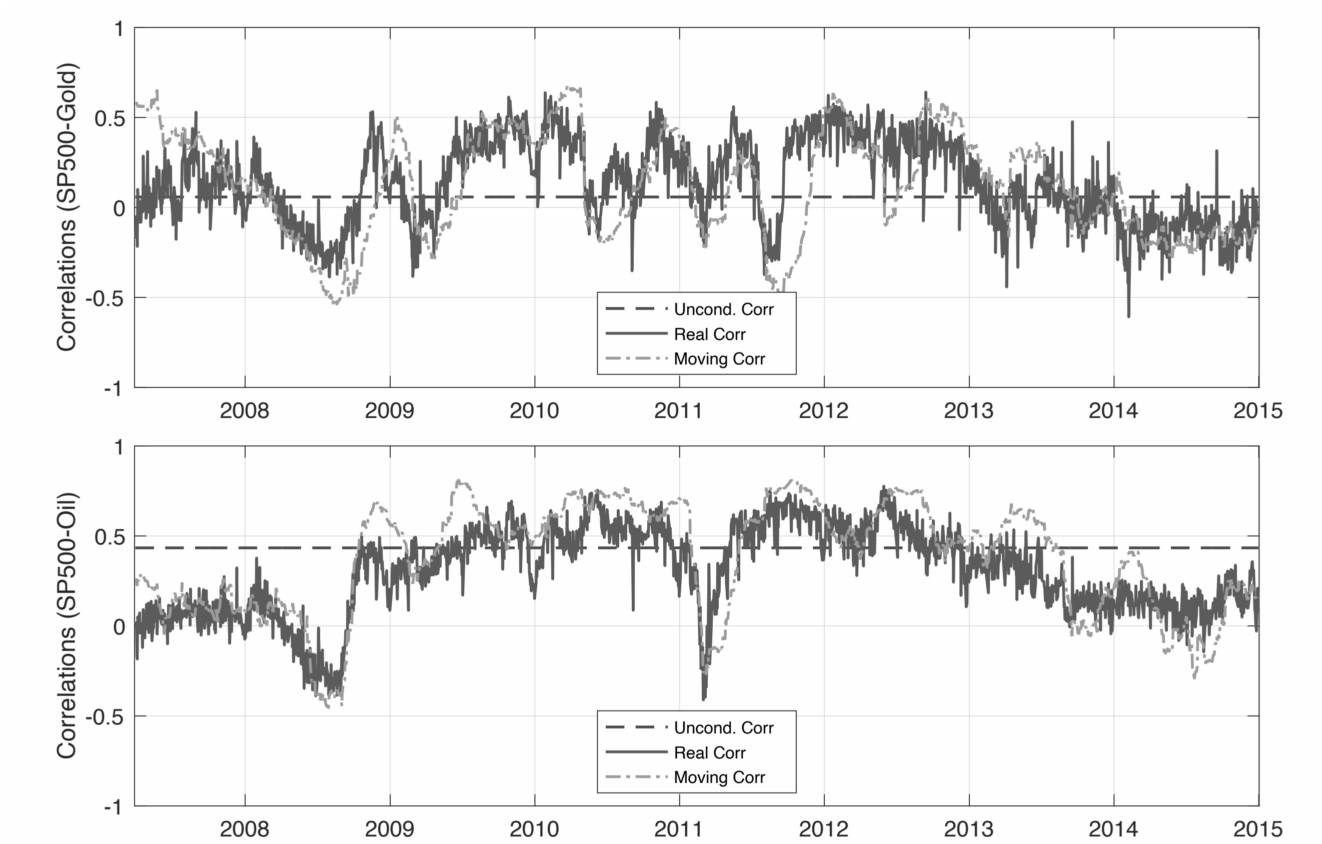

Figure 1: Time-varying correlations. The figure contains plots of correlations between the return on S&P 500 E-mini futures and the returns for two different commodities futures, gold in the top panel and oil in the bottom panel. Each panel presents three different measures of the correlation between the returns on S&P 500 futures and the returns on commodity futures. The horizontal, dashed line represents the sample correlation based on the full sample of daily observations. The light gray line represents the sample correlation calculated on a moving window of 60 daily observations. Finally, the solid, dark gray line represents the realized correlation based on the realized kernel and relying on the full sample of high-frequency data.

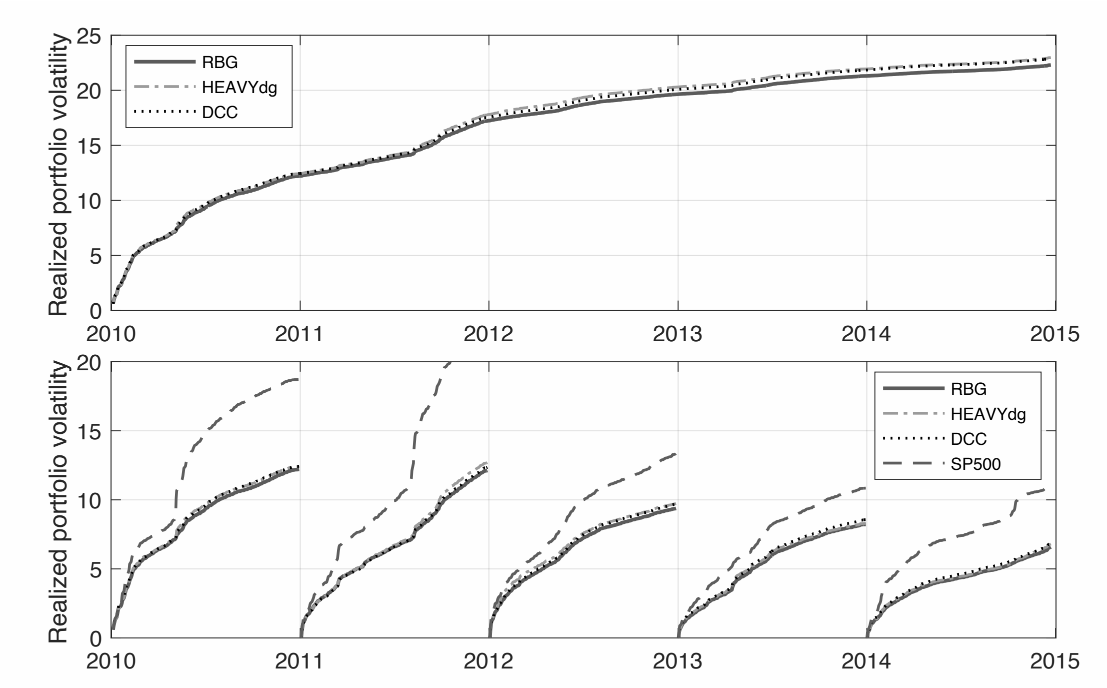

Figure 2: Realized portfolio volatility. The figure contains a comparison of the cumulated volatilities of the minimum variance portfolios based on forecasts from three different models. The cumulated volatility, $\sigma_{p,t,cum} = \sum_{\tau=1}^{t} \sigma_{p,\tau}$, is based on the realized portfolio volatility, $\sigma_{p,t+1} = \sqrt{\omega_t' RC_{t+1} \omega_t}$. Portfolio volatilities are based on forecasts for the RBG model, the DCC model, and the HEAVY.dg model. The figure also presents the cumulated volatility for a portfolio consisting of S&P 500 futures. The portfolios are not subject to any short-selling or turnover constraints. The bottom panel presents results of the analysis of the main plot, when carried out on a yearly basis.

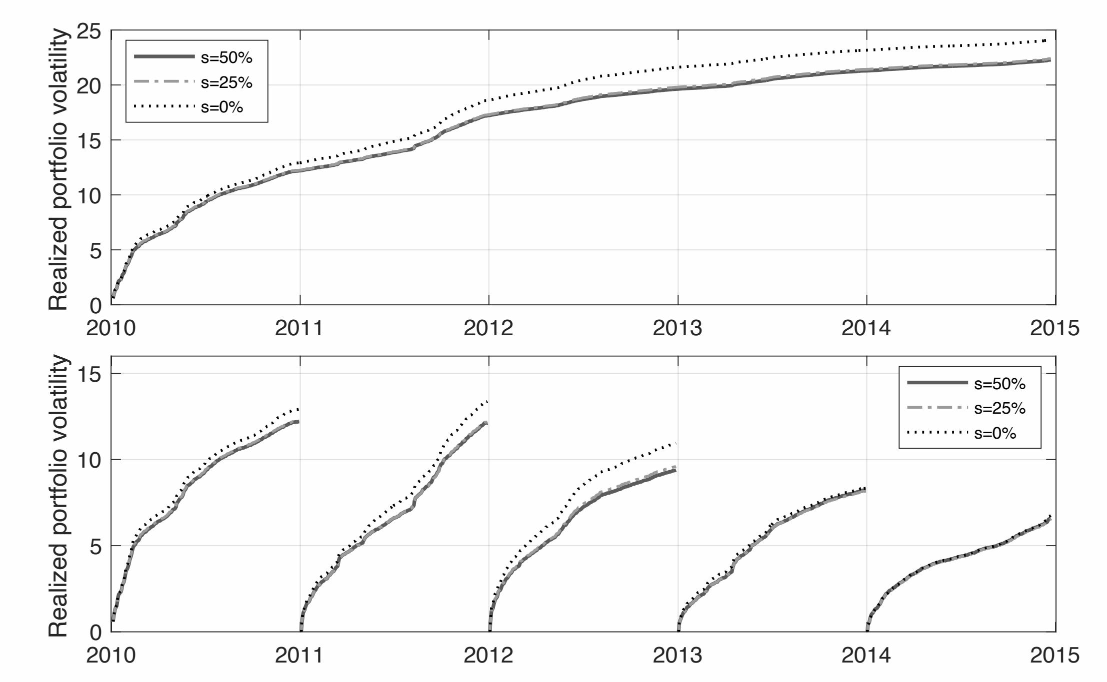

Figure 3: Effects of short-selling constraints. The figure shows the cumulated volatility of the minimum variance portfolios based on forecasts from the RBG model for different short-selling constraints. The cumulated volatility, $\sigma_{p,t,cum} = \sum_{\tau=1}^{t} \sigma_{p,\tau}$, is based on the realized portfolio volatility, $\sigma_{p,t+1} = \sqrt{\omega_t' RC_{t+1} \omega_t}$. Portfolio volatilities are based on forecasts for the RBG model, where the different lines represent different short-selling constraints, and the results are not subject to any turnover constraints. The bottom panel presents results of the analysis of the main plot, when carried out on a yearly basis.

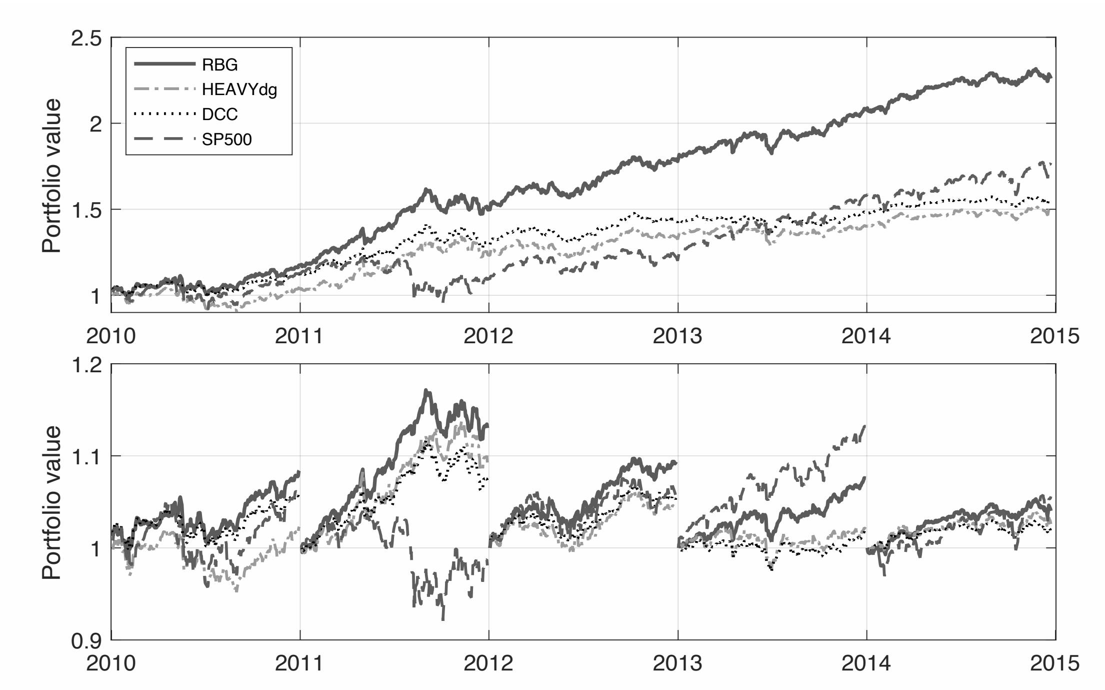

Figure 4: Portfolio value. The figure shows the cumulated portfolio value of the momentum portfolio based on forecasts from three different models. The cumulated portfolio value, $r_{p,t,cum} = \sum_{\tau=1}^{t} r_{p,\tau}$, is based on daily portfolio value, $r_{p,t+1} = \omega_t' r_{t+1}$. Portfolio weights are based on the forecasts for the RBG model, the DCC model, and the HEAVY.dg model. The figure also presents the cumulated value for a portfolio consisting of S&P 500 futures. A certain degree of short-selling is allowed, $s = 0.5$, but the results are not subject to any turnover constraints. The sub-plots presents results of the analysis of the main plot, when carried out on a yearly basis.

Table 1: Futures contracts.

Contract	First Trade	Exchange	Break	Average Daily Trades				Realized Covariance
				2005	2007	2010	2013	S&P 500
Equity Index								
S&P 500 (ES)	1997	COMEX	5:15-6:00pm	100,289	112,620	447,970	394,755	1.00
Metal								
Copper (HC)	1989	COMEX	5:15-6:00pm	1,313	4,097	19,513	34,511	0.86
Gold (GC)	1984	COMEX	5:15-6:00pm	4,687	19,892	65,384	103,725	0.21
Silver (SI)	1983	COMEX	5:15-6:00pm	2,900	6,817	22,400	32,341	0.53
Energy								
Heating Oil (HO)	1984	NYMEX	5:15-6:00pm	1,525	10,691	17,401	21,576	0.69
Light Crude (CL)	1987	NYMEX	5:15-6:00pm	3,962	50,425	135,008	129,535	0.98
Natural Gas (NG)	1993	NYMEX	5:15-6:00pm	2,290	15,500	41,064	54,539	0.31

The table presents contract information for seven futures contracts. The column labeled Contract divides the contracts into three groups and specify the underlying asset and the corresponding symbol in parenthesis. NYMEX and COMEX are both part of the CME Group. All times are in Eastern Standard Time. Average daily trade numbers are calculated over the respective calendar year from all intraday trades. Average realized covariances with S&P 500 E-minis are calculated based on realized close-close returns for the entire sample period.

Table 2: Realized volatility of the minimum variance portfolio and MCS p-values for all models, and different short-selling constraints

	$s = 0\%$					$s = 25\%$					$s = 50\%$				
	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	VR	$p(\sigma_{p,t})$	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	VR	$p(\sigma_{p,t})$	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	VR	$p(\sigma_{p,t})$
RBG	12.9	**10.74**	1.20	0.77	1.000	14.8	10.04	1.48	0.72	0.088	15.3	10.03	1.52	0.72	0.030
RBG.r	7.35	**10.74**	0.68	0.77	0.999	7.30	**10.01**	0.73	0.72	1.000	7.36	**10.00**	0.74	0.72	1.000
HEAVY.sc	6.78	11.49	0.59	0.82	0.000	7.48	10.70	0.70	0.77	0.000	7.44	10.72	0.69	0.77	0.000
HEAVY.dg	6.79	11.42	0.59	0.82	0.000	7.41	10.61	0.70	0.76	0.000	7.40	10.62	0.70	0.76	0.000
MW.RK	4.68	11.23	0.42	0.80	0.000	4.26	10.54	0.40	0.75	0.000	4.26	10.53	0.40	0.75	0.000
RM.RK	5.02	11.11	0.45	0.79	0.000	4.55	10.39	0.44	0.74	0.000	4.60	10.38	0.44	0.74	0.000
DCC	6.28	11.06	0.57	0.79	0.000	6.70	10.39	0.64	0.74	0.000	6.94	10.37	0.67	0.74	0.000
DCC.2	5.28	11.09	0.48	0.79	0.000	5.50	10.46	0.53	0.75	0.000	5.66	10.46	0.54	0.75	0.000
MW.r^2	5.62	11.38	0.49	0.81	0.000	7.70	10.78	0.71	0.77	0.000	7.67	10.77	0.71	0.77	0.000
RM.r^2	6.16	11.24	0.55	0.80	0.000	7.89	10.61	0.74	0.76	0.000	7.99	10.60	0.75	0.76	0.000
1/n	3.16	15.43	0.20	1.10	0.000	3.16	15.43	0.20	1.10	0.000	3.16	15.43	0.20	1.10	0.000

The table presents the average volatility of the minimum variance portfolio based on forecasts for 11 different models. The volatility estimate at time $t+1$, $\sigma_{p,t+1}$, is based on the portfolio weights obtained from the forecast of the covariance matrix for each model at time t and the realized covariance matrix of futures prices, RC_{t+1}, such that $\sigma_{p,t+1} = \sqrt{\omega_t' RC_{t+1} \omega_t}$. We present the average of $\sigma_{p,t}$ taken over the full sample, and denoted $\bar{\sigma}_p$. We also present the average return, $\bar{\mu}_p$ and the Sharpe ratio, SR. VR denotes the ratio of $\bar{\sigma}_p$ to the average realized volatility of a portfolio consisting exclusively of S&P 500 futures. The table also presents the p-values, $p_{MCS}(\sigma_{p,t})$, of the model confidence set, which is based on the time series of realized portfolio volatility. The columns present results for three different short-selling constraints. None of the results are subject to turnover constraints.

Table 3: Performance measures for the minimum variance portfolio based on the RBG model for different short-selling and turnover constraints

	$s = 0\%$			$s = 25\%$			$s = 50\%$		
δ	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR
0.0001	8.22	12.12	0.68	7.99	11.53	0.69	8.22	12.62	0.65
0.0025	6.50	11.88	0.55	6.48	11.33	0.57	5.74	11.59	0.49
0.005	6.19	11.62	0.53	5.68	11.04	0.51	5.28	11.17	0.47
0.01	6.27	11.42	0.55	5.83	10.81	0.54	5.46	10.84	0.50
0.05	5.43	11.07	0.49	3.82	10.38	0.37	3.48	10.37	0.34
0.1	7.56	10.90	0.69	6.62	10.20	0.65	6.18	10.20	0.61
1	7.83	11.33	0.69	6.71	10.70	0.63	6.30	10.72	0.59
10	9.53	10.85	0.88	8.68	10.15	0.85	8.59	10.15	0.85
100	12.9	10.74	1.20	14.8	10.04	1.48	15.3	10.03	1.52

The table presents performance measures for the minimum variance portfolio based on the forecasts of the RBG model. The columns present results for three different short-selling constraints while the rows contain different turnover constraints. For each specification, we present average return, realized volatility, and the Sharpe ratio.

Table 4: Financial performance measures for the momentum portfolio for all models, $s = 0.5$, $\delta = 100$

	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	$SR/SR_{S\&P}$	MAD	TO	τ^{be}	$p_{MCS}(U_t)$
RBG	17.03	10.60	1.606	1.752	0.191	0.157	0.336	1.000
RBG.r	8.16	10.61	0.769	0.840	0.190	0.156	0.132	0.000
HEAVY.sc	7.26	11.43	0.635	0.693	0.209	0.085	0.179	0.000
HEAVY.dg	7.04	11.40	0.617	0.673	0.208	0.091	0.159	0.000
MW.RK	4.50	11.23	0.400	0.437	0.171	0.046	0.099	0.000
RM.RK	4.95	11.03	0.449	0.490	0.173	0.047	0.142	0.000
DCC	8.13	10.99	0.740	0.807	0.191	0.125	0.179	0.000
DCC.2	7.17	11.05	0.649	0.708	0.193	0.152	0.128	0.000
MW.r^2	8.50	11.48	0.741	0.808	0.183	0.091	0.236	0.000
RM.r^2	9.58	11.34	0.845	0.922	0.185	0.095	0.282	0.000
1/n	3.16	15.43	0.205	0.223	0.000	0.009	-1.510	0.000

The table presents financial measures of the performances of the momentum portfolios, where portfolio weights are based on the forecasts for 11 different models. The column labeled SR contains the out-of-sample Sharpe ratio of the portfolio, calculated as the ratio of average daily portfolio returns and the corresponding standard deviation. The column labeled $SR/SR_{S\&P}$ presents the ratio of the Sharpe ratio of a portfolio to the Sharpe ratio of a portfolio consisting exclusively of S&P 500 futures. MAD is the mean absolute deviation of the portfolio weights from a $1/n$ portfolio. The column labeled τ^{be} presents the minimum average proportional trading costs, which will eliminate any gains from forming a portfolio based on forecasts from a particular model. Finally, $p_{MCS}(U_t)$ presents the p-values of the model confidence set estimated based on the time series of realized utility from the momentum portfolios based on the different models. A certain degree of short selling is allowed, $s = 0.5$, and the results are not subject to any turnover constraints.

Table 5: Performance measures for the momentum portfolio based on the RBG model for different short-selling and turnover constraints

		$s = 0\%$			$s = 25\%$			$s = 50\%$	
δ	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR	$\bar{\mu}_p$	$\bar{\sigma}_p$	SR
0.0001	8.21	12.14	0.68	7.90	12.28	0.64	8.04	13.61	0.59
0.0025	7.68	12.08	0.64	6.24	11.68	0.53	5.49	12.44	0.44
0.005	7.37	12.25	0.60	5.96	11.40	0.52	5.32	12.04	0.44
0.01	6.20	12.38	0.50	5.69	11.19	0.51	5.52	11.29	0.49
0.05	5.97	12.22	0.49	4.60	11.07	0.42	3.85	11.09	0.35
0.1	7.03	12.12	0.58	6.41	10.89	0.59	6.05	10.85	0.56
1	6.95	12.31	0.56	6.51	11.11	0.59	6.28	11.21	0.56
10	7.78	12.08	0.64	8.59	10.83	0.79	8.24	10.80	0.76
100	10.9	11.91	0.91	16.3	10.63	1.54	17.0	10.60	1.60

The table presents performance measures for the momentum portfolio based on the forecasts of the RBG model. The columns present results for three different short-selling constraints while the rows contain different turnover constraints. For each specification, we present average return, realized volatility, and the Sharpe ratio.

Table 6: Performance fee for the momentum portfolio for all models, $s = 0.5$, $\delta = 100$

	m_2	m_3	m_4	m_5	m_6	m_7	m_8	m_9	m_{10}	m_{11}
m_1, RBG	-0.053	-0.055	-0.067	-0.064	-0.054	-0.049	-0.048	-0.041	-0.061	-0.079
m_2, RBG.r		-0.002	-0.014	-0.011	-0.001	0.004	0.005	0.012	-0.008	-0.026
m_3, HEAVY.sc			-0.013	-0.010	0.001	0.006	0.007	0.014	-0.006	-0.024
m_4, HEAVY.dg				0.003	0.013	0.018	0.019	0.026	0.006	-0.012
m_5, RW.RK					0.011	0.016	0.016	0.023	0.003	-0.015
m_6, RM.RK						0.005	0.006	0.013	-0.007	-0.025
m_7, DCC							0.001	0.008	-0.012	-0.030
m_8, DCC.2								0.007	-0.013	-0.031
m_9, RW.r^2									-0.020	-0.038
m_{10}, RM.r^2										-0.018
m_{11}, 1/n										

The table presents the average daily performance fee (multiplied by 100) that eliminates the utility gains of one model compared to another model. It is the amount an investor would be willing to pay to switch from the model in the row to the model in the column. A certain degree of short selling is allowed, $s = 0.5$, and the results are not subject to any turnover constraints.

Appendix

A. Realized measures

Let $p^v_{t,i}$ be the intraday arbitrage-free log-price recorded at time t on day v for asset i and define the corresponding intraday log-returns as:

$$r^v_{t_m,i} = p^v_{t_m,i} - p^v_{t_{m-1},i}, \qquad m = 1, ..., n^v_i. \tag{A.1}$$

Returns for asset i are based on the n^v_i intraday prices recorded on a trading period ranging from after 5:15pm on day $v-1$ until just before 5:15pm on day v, the period of trading considered. Time coordinates are scaled to evolve in the unit interval $[0,1]$ associated with one period of trading, and for each asset, we define a partition, $\pi_{n^v_i} = [0 = t^i_0 < t^i_1 < ... < t^i_{n^v_i} = 1]$. If $\forall i, j = 1, ..., d : \pi_{n^v_i} = \pi_{n^v_j}$, for the d assets, the baseline realized covariance among the elements of $r_{v,t} = \{r^v_{t,i}\}_{i=1,...,d}$ over $[0,1]$ is defined as:

$$RC_v := \sum_{m=1}^{n^v} r_{v,t_m} r'_{v,t_m}, \qquad \text{where } n^v = n^v_1 = n^v_2 = ... = n^v_d. \tag{A.2}$$

To compute realized covariances, a homogeneous grid of evenly spaced prices is created using previous-tick interpolations as introduced in Dacorogna et al. (2001). Denoting a fixed grid of times containing n_δ points by $G_\delta = [\delta < ... < m\delta < ... < n_\delta \delta \leq 1]$, where δ denotes the calendar time sampling frequency, synchronous prices are constructed along the n_δ points of the grid as $p^v_{m\delta,i} = p^v_{t_j,i}$, where $j = \max(j' \mid t_{j'} \leq m\delta)$ and $t_j \leq m\delta < t_{j+1}$. Realized covariances are then computed using the returns obtained from the synchronized grid of high-frequency prices and are denoted by $RC^{(\delta)}_v$.

Sparse sampling of the realized covariances, $\delta = 15$ minutes, is applied with sub-sampling every 15 seconds. Evidence from the equity markets suggests that the impact of microstructure noise on realized measures is immaterial at this frequency [see, e.g., Barndorff-Nielsen and Shephard (2007) for details]. The choice of 15 minutes realized covariance, is likely to lead to a be a noisy but unbiased measure of the integrated covariance matrix.

We use the multivariate realized kernel (MRK) introduced by Barndorff-Nielsen et al. (2011). MRK is a class of estimators that is robust to measurement errors and microstructure effects induced by asynchronous trading. The MRK is based on a homogeneous series of high-frequency prices that are constructed using refresh time [see Barndorff-Nielsen et al. (2011) for details]. Specifically, non-flat-top realized kernels are used to ensure positive semi-definiteness,

$$K_v(p) = \sum_{h=-H}^{H} k\left(\frac{h}{H+1}\right)\Gamma_{v,h}, \qquad \Gamma_{v,h} = \sum_{j=|h|+1}^{n^v} r_{v,t_j} r'_{v,t_{j-h}}, \tag{A.3}$$

where $k(x)$ is the non-stochastic Parzen kernel function, $\Gamma_{v,h}$ is the realized auto covariance function, $H = c^{\star}\xi^{4/5}n^{3/5}$, $c^{\star} = \left\{\frac{k''(0)^2}{k_{\bullet}^{0,0}}\right\}^{1/5}$, $k_{\bullet}^{0,0} = \int_0^{\infty} k(x)^2 dx$, and $\xi^2 = \frac{\omega^2}{\sqrt{T\int_0^1 \sigma_u^4 du}}$, which is estimated as in Barndorff-Nielsen et al. (2011).

B. Estimation and forecasting

B.1 Estimation

Following Hansen et al. (2014), we adopt Gaussian specifications for the marginal and conditional densities implying that the maximum likelihood estimators of the variance-covariance parameters are given by:

$$\hat{\sigma}_{u_i}^2 = \frac{1}{T}\sum_{t=1}^{T}\hat{u}_{i,t}^2, \quad \hat{\sigma}_{u_j,u_i} = \frac{1}{T}\sum_{t=1}^{T}\hat{u}_{i,t}\hat{u}_{j,t}, \quad \hat{\sigma}_{v_j,u_i} = \frac{1}{T}\sum_{t=1}^{T}\hat{u}_{i,t}\hat{v}_{j,t}, \tag{B.1}$$

and

$$\hat{\Omega} = \frac{1}{T}\sum_{t=1}^{T}\hat{U}_{j,t}\hat{U}_{j,t}', \qquad \text{where} \quad \hat{U}_{j,t} = \begin{pmatrix} \hat{u}_{j,t} \\ \hat{v}_{j,t} \end{pmatrix} - \begin{pmatrix} \hat{\sigma}_{u_j,u_i}/\hat{\sigma}_{u_i}^2 \\ \hat{\sigma}_{v_j,u_i}/\hat{\sigma}_{u_i}^2 \end{pmatrix}\hat{u}_{i,t}. \tag{B.2}$$

The parameters are then estimated by maximizing

$$\ell(\theta) = -\frac{1}{2}\left(\ell_{z_i}(\theta_i) + \ell_{u_i}(\theta_i) + \ell_{z_j|z_i}(\theta) + \ell_{u_j,v_j|u_i}(\theta)\right), \tag{B.3}$$

where $\theta = \left(\theta_i', \theta_j', \theta_{j,i}'\right)'$ and $\ell_{z_i}(\theta_i) = \sum_{t=1}^{T}\left[\tilde{h}_{i,t}(\theta_i) + z_{i,t}^2(\theta_i)\right]$, $\ell_{u_i}(\theta_i) = T\left[\log\hat{\sigma}_{u_i}^2(\theta_i) + 1\right]$, $\ell_{u_j,v_j|u_i}(\theta) = T\left[\log\det\hat{\Omega}(\theta) + 2\right]$, and

$$\ell_{z_j|z_i}(\theta) = \left(\sum_{t=1}^{T}\log\left\{\left[1-\rho_{j,i,t}^2(\theta)\right]h_{j,t}(\theta)\right\} + \frac{\left(z_{j,t}(\theta) - \rho_{j,i,t}(\theta)\,z_{i,t}(\theta)\right)^2}{1-\rho_{j,i,t}^2(\theta)}\right), \tag{B.4}$$

as we can compute $\rho_{j,i,t} = F^{-1}\left\{a_{ji} + b_{ji}F\left(\rho_{j,i,t-1}\right) + c_{ji}F\left(y_{j,i,t-1}\right)\right\}$ independently of $h_{j,t}$ and $z_{j,t}$ recursively for $t = 2, \ldots, T$. For details on the multivariate realized beta GARCH, see Hansen et al. (2014).

B.2 Forecasts

The exposition in the main text presents only one-step ahead forecasts, but point forecasts are readily computable for different steps in the framework of the RBG model. The system of equations follows directly from Section 3.

$$\begin{aligned}\tilde{h}_{i,t+1} &= a_i + c_i\xi_i + (b_i + c_i\varphi_i)\tilde{h}_{i,t} + c_i\delta_i(z_{i,t}) + \tau_i(z_{i,t}) + c_i u_{i,t} \\ \tilde{h}_{j,t+1} &= a_j + c_j\xi_j + (b_j + c_j\varphi_j)\tilde{h}_{j,t} + c_j\delta_j(z_{j,t}) + \tau_j(z_{j,t}) + c_j u_{j,t} \\ \tilde{\rho}_{j,i,t+1} &= a_{ji} + c_{ji}\xi_{ji} + (b_{ji} + c_{ji}\varphi_{ji})\tilde{\rho}_{j,i,t} + c_{ji}v_{j,t},\end{aligned} \tag{B.5}$$

where $\tilde{\rho}_{j,i,t} := F(\rho_{j,i,t})$, $\delta_k(z_{k,t}) := \delta_{k,1}z_{k,t} + \delta_{k,2}\big(z_{k,t}^2 - 1\big)$, $\tau_k(z_{k,t}) := \tau_{k,1}z_{k,t} + \tau_{k,2}\big(z_{k,t}^2 - 1\big)$ for $k = i, j$. As pointed out by Hansen et al. (2014), the system of equations in (B.5) is seen to have a VARMA(1,1) representation by writing $\tilde{V}_{t+1} = \big(\tilde{h}_{i,t}, \tilde{h}_{j,t}, \tilde{\rho}_{j,i,t}\big)'$. This implies that:

$$\tilde{V}_{t+1} = C + A\tilde{V}_t + B\epsilon_t, \tag{B.6}$$

where $\epsilon_t = (\delta_i(z_{i,t}), \tau_i(z_{i,t}), \delta_j(z_{j,t}), \tau_j(z_{j,t}), u_{i,t}, u_{j,t}, v_{j,t})'$ and

$$C = \begin{bmatrix} a_i + c_i\xi_i \\ a_j + c_j\xi_j \\ a_{ji} + c_{ji}\xi_{ji} \end{bmatrix}, \quad A = \begin{bmatrix} b_i + c_i\varphi_i & 0 & 0 \\ 0 & b_j + c_j\varphi_j & 0 \\ 0 & 0 & b_{ji} + c_{ji}\varphi_{ji} \end{bmatrix}$$

$$B = \begin{bmatrix} c_i & 1 & 0 & 0 & c_i & 0 & 0 \\ 0 & 0 & c_j & 1 & 0 & c_j & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & c_{ji} \end{bmatrix}.$$

The k-step ahead forecasts can be computed as $\mathrm{E}(\tilde{V}_{t+k}|\mathcal{F}_t) = \tilde{V}_{t+k|t} = A^k\tilde{V}_t + \sum_{j=0}^{k-1} A^j C$. The RBG model is expressed using non-linear functions of the objects of interest implying for instance that $\mathrm{E}[F^{-1}(\rho_{j,i,t+k})|\mathcal{F}_t] \neq F^{-1}(\mathrm{E}[\rho_{j,i,t+k}|\mathcal{F}_t])$. This requires us to base forecasts of volatilities and correlations on simulation methods or alternatively on a bootstrapping procedure. We apply the simulation-based approach on which more details are provided in the section below.

B.2.1 Simulation approach

Let $V_t = \left(h_{j,t}, h_{i,t}, \rho_{i,j,t}\right)'$ denote the vector of non-transformed variables, the function

$f : \mathbb{R}^2 \times [0,1] \curvearrowright \mathbb{R}^3$ such that $f^{-1}(\tilde{V}_t) = V_t$, and start from the VARMA(1,1) specification in Eq. (B.6). From this, one can recursively construct point forecasts as

$$\tilde{V}_{t+k} = C + A\tilde{V}_{t+k-1} + B\epsilon_{t+k-1}. \tag{B.7}$$

The one-step ahead forecast $V_{t+1|t}$ does not require simulation since it is $\mathcal{F}_t$-measurable. For $k > 1$, $V_{t+k|t}$ is computed based on simulations as $\frac{1}{S}\sum_{s=1}^{S} f^{-1}(\tilde{V}^s_{t+k|t})$, where $\tilde{V}^s_{t+k|t}$ is obtained from $\tilde{V}^s_{t+k} = C + A\tilde{V}_{t+k-1} + B\epsilon^s_{t+k-1}$. ϵ^s_{t+k-1} is constructed by sampling the residuals of the RBG model from a conditional Gaussian distribution

$$\zeta_{t+k} := \begin{pmatrix} z_{i,t+k} \\ z_{j,t+k} \\ u_{i,t+k} \\ u_{j,t+k} \\ \nu_{j,t+k} \end{pmatrix} \sim N_5 \left(0, \begin{bmatrix} I_2 & 0 \\ 0 & \Sigma \end{bmatrix} \right), \qquad t = 1, ..., N. \tag{B.8}$$

The simulation is performed using an estimate of the matrix Σ. Residuals for commodity j, which are correlated with commodity i, can be computed using the law of motion of correlation in Equation (2), together with the simulated values for $\nu_{j,t+k}$ and defining:

$$z_{j,t+k} := \rho_{j,i,t+k} z_{i,t+k} + \sqrt{1-\rho^2_{j,i,t+k}}\breve{z}_{j,t+k}. \tag{B.9}$$

The assumptions regarding the distribution of ζ_t might be called into question and a parametric bootstrap may be preferable in some instances. The empirical properties of financial returns on equities standardized by realized measures do, however, have an empirical density close to normal [see, e.g., Andersen et al. (2003)]. We assume that this is also the case for commodities.