

2012 | 24

Working Paper

Norges Bank Research

Oil price density forecasts: exploring the linkages with stock markets

Marco J. Lombardi and Francesco Ravazzolo

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:

servicesenter@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatterens regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail:

servicesenter@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

Oil price density forecasts: exploring the linkages with stock markets*

Marco J. Lombardi[†] Francesco Ravazzolo[‡]

December 5, 2012

Abstract

In the recent years several commentators hinted at an increase of the correlation between equity and commodity prices, and blamed investment in commodity-related products for this. First, this paper investigates such claims by looking at various measures of correlation. Next, we assess to what extent correlations between oil and equity prices can be exploited for asset allocation. We develop a time-varying Bayesian Dynamic Conditional Correlation model for volatilities and correlations and find that joint modelling of oil and equity prices produces more accurate point and density forecasts for oil which lead to substantial benefits in portfolio wealth.

Keywords: Oil price, stock price, density forecasting, correlation, Bayesian DCC.

JEL Classification : C11, C15, C53, E17, G17.

*We thank seminar participants at ECB-NB “Modeling and forecasting oil prices”, MMF “Understanding Oil and Commodity Prices”, and CAMP “Forecasting and Analysing Oil Prices” workshops for helpful comments. This paper is part of the research activities at the Centre for Applied Macro and Petroleum economics (CAMP) at BI Norwegian Business School. CAMP is supported by Statoil’s research program in petroleum economics. The views expressed in this paper are our own and do not necessarily reflect those of Bank for International Settlements and/or Norges Bank.

[†]Bank for International Settlements, Centralbahnplatz 2, CH-4002 Basel (CH). Email: marco.lombardi@bis.org

[‡]Norges Bank and BI Norwegian Business School, Bankplassen 2, NO-0107 Oslo (Norway). Email: francesco.ravazzolo@norges-bank.no

1 Introduction

The past decade has witnessed a broad-based surge in commodity prices, with oil a frontrunner. The upward trend in prices has been ascribed to booming demand at the global level, but fluctuations around it have been substantial, especially after the onset of the Great Recession. Investing in commodities has generated hefty returns and has become increasingly popular, in spite of the high risks associated to this type of investment due to inherent volatility of commodity prices. Indeed, most fund managers have started advising their customers to devote a quota of their portfolios to commodity-related products as part of long-term diversification strategy, due to the fact the commodity prices are believed to possess lower correlation with other asset classes, most notably with stock markets.

At the same time, substantial inflows into commodity-related investment products have led many commentators to speculate on whether commodities are increasingly behaving as an asset class. The empirical evidence on a lasting impact of financial investment on commodity prices is, at best, scant (see Fattouh et al. [2012] for a recent survey). Yet, financial investors in commodity futures markets may have less commodity-specific knowledge and a different attitude compared to commercial players, and hence exit or enter trades based on their overall macroeconomic risk perceptions rather than market-specific factors. Such flows could, in principle, lead to higher correlation with equity prices.

In this paper, we try to shed some light on the issue. If one looks at the most recent years, commodity and equity prices have appeared to be increasingly correlated, being both extremely sensitive to news concerning the global macroeconomic environment rather than to idiosyncratic and market-specific shocks. We first provide a complete characterization of this phenomenon, by computing correlations using different model-based and model-free methodologies and trying to identify relevant turning points. Our results suggest that correlation, after having hovered around zero for more than a decade,

has indeed increased markedly, starting in 2008.

The connection between equity and oil prices has already been examined in the literature: Kilian and Park [2009] report that the response of equity prices to oil price shocks depends on their nature; Cassassus and Higuera [2011] show that oil price changes are good predictors of equity prices; Chang et al. [2011] report evidence of volatility spillovers between oil and equity prices. To the best of our knowledge, however, no studies have examined the predictability of oil and equity prices in a density forecasting framework. This raises to question of to what extent co-movements between oil and equity prices could be exploited to improve forecasts in either direction. We use Bayesian constant parameter univariate and bivariate models for oil and stock prices and derive a time-varying Bayesian Dynamic Conditional Correlation (DCC) model which can account for the changes in their relationship observed after 2008, see Della Corte et al. [2010]. We conduct a density forecasting exercise and find that the DCC provides statistically superior density forecasts compared to a no-change random walk model.

We assess the economic value of such forecasting gains by considering an oil and equity prices asset-allocation framework. The use of a Bayesian framework in estimating our models allows us to explicitly account for the fact that higher order moments and the full predictive densities of oil and equity prices are uncertain which can result in suboptimal allocation. The estimated predictive densities depend on the data and the prior and integrate estimation risk into the portfolio allocation, see Kandel and Stambaugh [1996], Barberis [2000], Avramov [2002], Cremers [2002], Kan and Zhou [2007] and Jacquier and Polson [2012]. Ravazzolo et al. [2007], Guidolin and Na [2007] and Pettenuzzo and Timmermann [2011] extend the analysis to account for instability uncertainty. Della Corte et al. [2011] and Della Corte et al. [2010] assess the economic value of volatility and correlation timing respectively. We extend this literature by including in the asset allocation problem a new class of assets, commodity prices and specifically oil prices. Our results

highlight that a time-varying jointly modeling of oil and equity prices leads to economic gains relative to passive strategies and constant parameter models, especially at times of large swings.

The remainder of the paper is organized as follows. Section 2 describes the data and documents changes in correlation over our sample. Section 3 presents bivariate models for oil and stock prices, investigates their point and density forecast accuracy and applies them to active asset investment strategies. Section 4 concludes.

2 Oil prices and stock markets

The idea that oil and stock prices should display a negative correlation finds root into expressing stock prices as the discounted value of future dividends. In case of an increase in energy prices, firms will see their profits shrink, and will therefore distribute less dividend. In a sense, however, this hints at oil prices being driven exogenously. It is now widely acknowledged that this is not the case. Oil price increases come on the back of booming economic activity due to increases in demand. Therefore, the sign of the correlation becomes less clear-cut: both stock and oil prices should increase when positive news on the global macroeconomic outlook materialize. Kilian and Vega [2011] indeed investigate the impact of various macroeconomic announcements on the oil price over a sample running from 1983 to 2008 and fail to find a systematic relationship.

If one takes a long-run perspective, it appears that oil and stock prices at times moved in the same direction and at times diverged. So it seems that their correlation pattern has broadly varied over time. In this section, we will take a stance on this empirical finding by gauging correlations between oil and stock prices using a number of different methods. To start with, we collected weekly returns generated by next-month Brent crude oil futures and the SP500 index, starting from April 1985 until September 2011,

see Figure 1. Oil prices are substantially more volatile as Figure 2 indicates. Figure 3 reports sample correlations computed on moving windows of different length; this is indeed what most commentators refer to when discussing comovements of oil and stock prices. It emerges clearly that correlations have been positive and negative at times, but have overall hovered around zero. The recent swarm of positive correlations appears instead to have been stronger and more persistent compared to past ones: it started with the burst of the financial crisis in 2008, and although correlations have been recently declining, they are still positive. Of course, using different windows somehow alters the intensity of the results, but the main pattern seems to remain.

Using rolling windows of sizeable length may spread over time the influence of extreme episodes or periods of market turbulence. To avoid this, one can use model-based approaches to estimating correlations. A popular model in this direction is to use Dynamic Conditional Correlations (DCC) as proposed by Engle [2002]:

$$\begin{aligned}
y_t &= \Phi(L)y_{t-1} + v_t \\
v_t &= H_t^{0.5}\varepsilon_t, \quad \varepsilon_t \sim N(0, I_N), \quad H_t = D_t R_t D_t \\
D_t^2 &= \text{diag}\{\omega_i\} + \text{diag}\{\kappa_i\}v_{t-1}v'_{t-1} + \text{diag}\{\lambda_i\}D_{t-1}^2, \quad i = 1, \dots, N \\
Q_t &= S(u' - A - B) + A\varepsilon_{t-1}\varepsilon'_{t-1} + BQ_{t-1} \\
R_t &= \text{diag}\{Q_t\}^{-1}Q_t\text{diag}\{Q_t\}^{-1}
\end{aligned} \tag{1}$$

where y_t is a $(N \times 1)$ vector of dependent variables, S is the unconditional correlation matrix of ε_t , A , B and $S(u' - A - B)$ are positive semidefinite matrices; see Appendix for model details. In a nutshell, the DCC is a multivariate GARCH model in which correlations are time-varying according to an autoregressive specification. As such, the DCC accounts for both the time-varying features of volatilities and correlations. The DCC correlations are reported in Figure 4. Although correlations appear to be smaller compared to those estimated using rolling windows, they still look persistently positive

towards the end of the sample.

Finally, we also tried to exploit ultra-high-frequency intradaily data to construct a model-free instantaneous measure of realized correlation, as proposed by Barndorff-Nielsen and Shephard [2004],¹ employing hourly data from October 2002. The sample is shorter, correlations are more volatile and erratic, see Figure 5, but one can still see that, after mid-2008, they have tended to cluster in positive territory.

3 Forecasting oil prices using stock prices

3.1 Point and density forecasts

The debate on the predictability of stock prices is still an open issue in empirical research, see for example Welch and Goyal [2008]. Market efficiency theories imply not predictability, whether market friction theories imply predictability. Evidence of predictability of oil prices has, on contrary, been subject to a break in recent years: from middle of 90's research evidence suggests not predictability where future prices contain all the relevant information and alternative models cannot improve forecast accuracy. However, very recent studies such as Baumeister and Kilian [2012] find that several (Bayesian) reduced form Vector Autoregressive models outperform forecasts based on future prices in real-time analysis. Baumeister and Kilian's models apply macroeconomic data to forecast oil prices, but do not explore the linkage with stock prices. Furthermore, their analysis refers mainly to point forecasting. Kandel and Stambaugh [1996] and Barberis [2000], among others, discuss the role of parameter uncertainty and Bayesian analysis as tool to cope with for return predictability and for asset allocation.

We produce weekly point and density forecasts for oil and stock returns over the sample from 2005W1 to 2011W40 for a total of 353 weeks. We compute $h = 1, 2, \dots, 24$

¹See Bicchetti and Maystre [2012] for a high-frequency intradaily analysis of other energy markets.

steps ahead forecast for each vintage using a bivariate Bayesian Vector Autoregressive model with Minnesota type prior (VAR), see Clark and Ravazzolo [2012] for details, and a bivariate Bayesian DCC model (DCC), see Appendix for details and Della Corte et al. [2010] for an application to assess the economic value of time-varying correlation timing. We compare these forecasts to random walk (RW) and Bayesian autoregressive (AR) forecasts. Bayesian inference on the listed models allows to derive complete predictive densities which we evaluate in density forecasting exercise, the main research question of this paper, and to take into account higher moments, key ingredients in portfolio decisions, in the asset allocation exercise. In particular, to shed light on the predictive ability of individual models, we consider several evaluation statistics for point and density forecasts previously proposed in the literature, see Billio et al. [2012] for a recent application. We evaluate oil and stock forecasts separately in this section and use marginal densities from bivariate models.

We compare point forecasts in terms of Mean Square Prediction Errors (MSPE)

$$MSPE_k = \frac{1}{t^*} \sum_{t=\underline{t}}^{\bar{t}} e_{k,t+1}^2$$

where $t^* = \bar{t} - \underline{t} + 1$ and $e_{k,t+1}^2$ is the square prediction error of model k . We evaluate the predictive densities using two relative measures. Firstly, we consider a Kullback Leibler Information Criterion (KLIC) based measure; see for example Kitamura [2002], Mitchell and Hall [2005], Hall and Mitchell [2007], Amisano and Giacomini [2007], Kascha and Ravazzolo [2010]. The KLIC distance between the true density $p(y_{t+1}|y_{1:t})$ of a random variable y_{t+1} and some candidate density $p(\tilde{y}_{k,t+1}|y_{1:t})$ obtained from model k is defined

as

$$\begin{aligned} \text{KLIC}_{k,t+1} &= \int p(y_{t+1}|y_{1:t}) \ln \frac{p(y_{t+1}|y_{1:t})}{p(\tilde{y}_{k,t+1}|y_{1:t})} dy_{t+1}, \\ &= \mathbb{E}_t[\ln p(y_{t+1}|y_{1:t}) - \ln p(\tilde{y}_{k,t+1}|y_{1:t})]. \end{aligned} \quad (2)$$

where $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot|\mathcal{F}_t)$ is the conditional expectation given information set \mathcal{F}_t at time t . An estimate can be obtained from the average of the sample information, $y_{\underline{t}+1}, \dots, y_{\bar{t}+1}$, on $p(y_{t+1}|y_{1:t})$ and $p(\tilde{y}_{k,t+1}|y_{1:t})$:

$$\overline{\text{KLIC}}_{k+1} = \frac{1}{t^*} \sum_{t=\underline{t}}^{\bar{t}} [\ln p(y_{t+1}|y_{1:t}) - \ln p(\tilde{y}_{k,t+1}|y_{1:t})]. \quad (3)$$

The KLIC chooses the model which on average gives higher probability to events that have actually occurred. In reality we do not know the true density, but for the comparison of two competing models, it is sufficient to consider the Logarithmic Score (LS), which corresponds to the latter term in the above sum,

$$LS_k = -\frac{1}{t^*} \sum_{t=\underline{t}}^{\bar{t}} \ln p(\tilde{y}_{k,t+1}|y_{1:t}), \quad (4)$$

for all k and to choose the model for which the expression in (4) is minimal, or as we report in our tables, the opposite of the expression in (4) is maximal.

Secondly, we also evaluate density forecasts based on the continuous rank probability score (CRPS). The CRPS for the model k measures the average absolute distance between the empirical cumulative distribution function (CDF) of y_{t+h} , which is simply a step function in y_{t+h} , and the empirical CDF that is associated with model k 's predictive

density:

$$\text{CRPS}_{k,t+1} = \int (F(z) - \mathbf{I}\{y_{t+1} \leq z\})^2 dz \quad (5)$$

$$= \mathbf{E}_t |\tilde{Y}_{t+1,k} - y_{t+1}| - \frac{1}{2} \mathbf{E}_t |\tilde{Y}_{t+1,k} - Y'_{t+1,k}|, \quad (6)$$

where F is the CDF from the predictive density $p(\tilde{y}_{k,t+1}|y_{1:t})$ of model k , $\mathbf{I}(\cdot)$ takes a value 1 if $y_{t+1} \leq z$ and equals 0 otherwise, and $\tilde{Y}_{t+1,k}$ and $\tilde{Y}'_{t+1,k}$ are independent random variables with common sampling density equal to the posterior predictive density $p(\tilde{y}_{k,t+1}|y_{1:t})$. The CRPS circumvents some of the drawbacks of the LS, as the latter does not reward values from the predictive density that are close but not equal to the realization (see, e.g., Gneiting and Raftery [2007]) and it is sensitive to outliers; see Gneiting and Ranjan [2011], Groen et al. [2012] and Ravazzolo and Vahey [2012]. Smaller CRPS implies higher precisions.

Since the distribution properties of a statistical test to compare density accuracy performances, both measured in terms of LS and CRPS, are not derived when working with nested models and expanding data window for parameter updating, as is our case, we follow Groen et al. [2012] and test the null of equal finite sample forecast accuracy, based on either a LS and CRPS measures, *versus* the alternative that a model outperformed the RW benchmark using the Harvey et al. [1997] small sample correction of the Diebold and Mariano [1995] and West [1996] statistic to standard normal critical values.² Following evidence in Clark and McCracken [2012] for point forecasting, we apply the same test to investigate superiority in square prediction errors.

Table 1 reports point and density forecast results. Absolute predictability for oil prices is substantially lower than absolute predictability for stock returns: MSPEs are higher, LS lower and CRPS higher for all horizons. Data characteristics discussed in

²We use the left tail p-values for the CRPS based test since we minimize CRPS and right tail for the LS based test since we maximize LS.

section 2 can explain the result. However, the evidence is somewhat different for relative predictability. The VAR and DCC models give lower MSPE and CRPS and higher LS for most of the horizons in forecasting oil prices, but not in forecasting stock returns. Therefore, stock prices seem to contain relevant information to forecast oil prices, whether the opposite is not supported by our analysis. Ferraro et al. [2012] find opposite evidence when investigating exchange rate and oil price predictability: oil prices forecast exchange rates, but exchange rates do not forecast oil prices. Improvements are, however, often very small.

Moreover, the DCC model gives more accurate forecasts relative to the RW benchmark for all the horizons up to 24 weeks and improvements in density forecasting are in several cases statistically significant. A time-varying covariance matrix which can model instability in volatility and correlations between the two variables discussed in section 2 is an important ingredient to predict higher moments of the joint oil and stock predictive density.

3.2 Economic gain analysis

Investors are mainly interested in the economic value of a forecasting model; the marginally higher statistical accuracy documented in the previous section can indeed give rise to substantial economic gains. To assess this, we develop an active short-term investment exercise. The investor's portfolio consists of the stock index, of the Brent oil index and risk free bonds only.³

At the end of each week t , the investor decides upon the fraction $\alpha_{s,t+h}$ of her portfolio to be held in stocks, $\alpha_{o,t+h}$ in the oil and the remaining part in the risk free asset for the period $t + h$, based on the forecast of the oil and stock returns. We constrain $\alpha_{s,t+h}$,

³The risk free asset is approximated by using the weekly federal fund rate. We collect the federal fund rate from the Fred database at the Federal Reserve Bank of St Louis.

$\alpha_{o,t+h}$ to be in the $[0, 1]$ interval and to sum maximum to 1, therefore not allowing for short-sales or leveraging (see Barberis [2000]).⁴ We assume that the investor maximize a power utility function:

$$u(R_{t+h}) = \frac{R_{t+h}^{1-\gamma}}{1-\gamma}, \quad \gamma > 1, \quad (7)$$

where γ is the coefficient of relative risk aversion and R_{t+h} is the wealth at time $t + h$, which is equal to

$$R_{t+h} = R_t ((1 - \alpha_{s,t+h} - \alpha_{o,t+h}) \exp(y_{f,t+h}) + \alpha_{t+h} \exp(y_{f,t+h} + \tilde{y}_{t+h})), \quad (8)$$

where R_t denotes initial wealth, $\alpha_{t+h} = (\alpha_{s,t+h}, \alpha_{o,t+h})$, $y_{f,t+h}$ the h-step ahead risk free rate and \tilde{y}_{t+h} the h-step ahead bivariate forecast of the oil and stock returns in excess of the risk free made at time t . Our investor does not rebalance the portfolios in the period from t to $t + h$, but keeps positions on the three assets constant.⁵

When the initial wealth is set equal to one, i.e. $R_0 = 1$, the investor solves the following problem:

$$\max_{\alpha_{t+h} \in [0,1]^2 \sum \alpha_{t+h} \leq 1} \mathbb{E}_t \left(\frac{(((1 - \alpha_{s,t+h} - \alpha_{o,t+h}) \exp(y_{f,t+h}) + \alpha_{t+h} \exp(y_{f,t+h} + \tilde{y}_{t+h})))^{1-\gamma}}{1-\gamma} \right),$$

The expectation $\mathbb{E}_t(\cdot)$ depends on the predictive density for the oil and stock excess returns, \tilde{y}_{t+h} and the problem can be rewritten as:

$$\max_{\alpha_{t+h} \in [0,1]^2 \sum \alpha_{t+h} \leq 1} \int u(R_{t+h}) p(\tilde{y}_{t+h} | y_{1:t}) d\tilde{y}_{t+h}. \quad (9)$$

⁴We have also investigated exercises with $\alpha_{o,t+h}$ restricted in the $[0,0.3]$ and $[0,0.5]$ intervals. Results are qualitatively similar, economic gains are marginally lower.

⁵In the case of dynamic asset allocation the long-run investor is allowed to rebalance her portfolio during the investment period, adjusting the portfolio weights to reflect new information that arrives. Solving the resulting dynamic programming problem is complicated due to the large number of state variables that enter the problem in a highly nonlinear way, see Barberis [2000] and Guidolin and Timmermann [2007].

where $p(\tilde{y}_{t+h}|y_{1:t})$ is the predictive density for oil and stock excess returns. We approximate the integral in (9) by generating G independent draws from the predictive density $p(\tilde{y}_{t+h}|y_{1:t})$, and then use a numerical optimization method to find:

$$\max_{\alpha_{t+h} \in [0,1]} \frac{1}{G} \sum_{g=1}^G \left(\frac{((1 - \alpha_{t+h}) \exp(y_{f,t+h}) + \alpha_{t+h} \exp(y_{f,t+h} + \tilde{y}_{t+h}^g))^{1-\gamma}}{1 - \gamma} \right) \quad (10)$$

We consider an investor who can choose between different forecast densities of the (excess) oil and stock returns y_{t+h} to solve the optimal allocation problem described above. We include three cases in the empirical analysis below and assume the investor uses alternatively the density from the RW and AR univariate models for each series, the bivariate BVAR and the bivariate BVAR-DCC. Moreover, since the portfolio weights in the active investment strategies change every period, we include transaction costs of $c = 0.05\%$, 5 basis points.

We evaluate the different investment strategies for a risk investor with $\gamma = 2$ by computing the *ex post* annualized mean portfolio return, the annualized standard deviation and the Sharpe ratio. Results are qualitatively similar for $\gamma = 2, 6$. We compare the wealth provided at time $t+h$ by two resulting portfolios by determining the value of multiplication factor of wealth Δ which equates their average utilities. For example, suppose we compare two strategies A and B.

$$\sum_{t=\underline{t}}^{\bar{t}} u(R_{A,t+h}) = \sum_{t=\underline{t}}^{\bar{t}} u(R_{B,t+h} / \exp(r)). \quad (11)$$

where $u(R_{A,t+h})$ and $u(R_{B,t+h})$ are the wealth provided at time $T+h$ by the two resulting portfolios A and B, respectively. Following West et al. [1993], we interpret Δ as the maximum performance fee the investor would be willing to pay to switch from strategy A to strategy B.⁶ We infer the added value of strategies based on individual models and the

⁶See, for example, Fleming et al. [2001] for an application with stock returns.

combination scheme by computing Δ with respect to three static benchmark strategies: holding only stock (r_s), holding only oil (r_o), and holding only free risk bond (r_f).

Finally we compute the certainty equivalent return (CER) for each strategy, in formula:

$$CER_{t+h} = u^{-1}(\mathbb{E}_t(u(R_{t+h}))) \quad (12)$$

where u^{-1} is the inverse of the power utility function defined in (7). Strategy with maximum CER is preferred. We do not report the final value as for other alternative measures, but plot how cumulative difference between any strategy A and investing 100% of the portfolio in the stock market:

$$CERD_{k,t+1} = \sum_{s=t}^t (CER_{A,t+h} - CER_{s,t+h}), \quad (13)$$

where $k = RW, AR, \dots, FRR$. If $CERD_{k,t+h}$ increases at observation $t+h$, this indicates that the strategy k gives higher CER than the benchmark strategy.

Results in Table 2 strength the evidence reported in the section on point and density forecasting. The VAR and DCC models give higher SR than the RW and have positive fees relative to passive strategies for all the horizons. Gains are robust to reasonable transaction costs. The DCC provides the highest gains for horizons longer than two weeks and, in particular, they are substantial compared to alternative models for our middle horizons.

Our analysis shows that a joint oil and stock returns model with time-varying volatility can produce statistical and economic gains. To shed light how such gains are made, Figure 6 plots CERD relative to a passive strategy of investing 100% of the portfolio in stock prices over the sample period. For all the six investment horizons we consider, the gains are mainly in the second part of 2008 during the most turbulent time of the recent financial crisis. The joint strategies give lower CER before the crisis and at beginning of 2009, and

similar CERs to the ones of the passive strategies from the second semester of 2009. Plots also confirm that the DCC has the highest CER for horizons from 4 to 12 weeks ahead.

4 Concluding remarks

We have investigated the use of stock prices to predict oil prices. The analysis shows that the correlation between the two series has substantially increases after the beginning of recent financial crisis in the second semester of 2008. A bivariate Bayesian Dynamic Conditional Correlation model which can model such time variation produces statistically more accurate density forecasts and gives large economic gains in an asset allocation exercise relative to the benchmark random walk model. The value of an active strategy based on the DCC forecasts is substantially large compared to passive strategies during turbulent times.

The research question can be extended in several directions. The variable set in the DCC model could be extended by using interest rates and global demand index such as in Kilian [2009]. The accuracy of the model could be tested using other loss functions, such as risk measures for policy analysis.

References

- G. Amisano and R. Giacomini. Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics*, 25(2):177–190, 2007.
- D. Avramov. Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64:423–458, 2002.
- N. Barberis. Investing for the long run when returns are predictable. *Journal of Finance*, 55:255–264, 2000.

- O. E. Barndorff-Nielsen and N. Shephard. Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica*, 72(3):pp. 885–925, 2004.
- C. Baumeister and L. Kilian. Real-time forecasts of the real price of oil. *Journal of Business and Economic and Statistics*, 30(2):326–336, 2012.
- D. Bicchetti and N. Maystre. The synchronized and long-lasting structural change on commodity markets: evidence from high frequency data. Technical Report 37486, MPRA Paper, 2012.
- M. Billio, R. Casarin, F. Ravazzolo, and van Dijk H.K. Time-varying combinations of predictive densities using nonlinear filtering. Technical report, Norges Bank, 2012.
- J. Cassassus and F. Higuera. Stock return predictability and oil prices. Technical Report 406, Instituto de Economia, Pontificia Universidad Catolica de Chile, 2011.
- C.-L. Chang, M. McAleer, and R. Tansuchat. Crude oil hedging strategies using dynamic multivariate garch. *Energy Economics*, 33(5):912–923, 2011.
- T. Clark and M. W. McCracken. Advances in forecast evaluation. In A. Timmermann and G. Elliott, editors, *Handbook of Economic Forecasting*. Elsevier, Amsterdam, 2012.
- T. Clark and F. Ravazzolo. Comparing models of time-varying macroeconomic volatility. Technical Report 12-18, FRB of Cleveland Working Paper, 2012.
- K.J.M. Cremers. Stock return predictability: A bayesian model selection perspective. *Review of Financial Studies*, 15:1223–1249, 2002.
- P. Della Corte, L. Sarno, and I. Tsiakas. Correlation timin in asset allocation with bayesian learning. Technical report, Warwick Business School, 2010.

- P. Della Corte, L. Sarno, and I. Tsiakas. Spot and forward volatility in foreign exchange. *Journal of Financial Economics*, 100:496–513, 2011.
- F.X. Diebold and R.S. Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13:253–263, 1995.
- R. Engle. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350, 2002.
- B. Fattouh, L. Kilian, and L. Mahadeva. The role of speculation in oil markets: What have we learned so far? CEPR Discussion Papers 8916, C.E.P.R. Discussion Papers, 2012.
- D. Ferraro, K. Rogoff, and B. Rossi. Can oil prices forecast exchange rates? NBER Working Paper 17998, NBER, 2012.
- J. Fleming, C. Kirby, and B. Ostdiek. The Economic Value of Volatility Timing. *Journal of Finance*, 56:329–352, 2001.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.
- T. Gneiting and R. Ranjan. Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business and Economic Statistics*, 29:411–422, 2011.
- J.J.J Groen, R. Paap, and F. Ravazzolo. Real-time inflation forecasting in a changing world. *Journal of Business and Economic Statistics*, forthcoming, 2012.
- M. Guidolin and C. F. Na. The economic and statistical value of forecast combination under regime switching: An application to u.s. predictable returns. In D. Rapach and

- M. Wohar, editors, *Forecasting in the Presence of Structural Breaks and Uncertainty*, volume 3. Emerald Publishing Group, 2007.
- M. Guidolin and A. Timmermann. Asset allocation under multivariate regime switching. Technical report, 2007.
- S. G. Hall and J. Mitchell. Combining density forecasts. *International Journal of Forecasting*, 23:1–13, 2007.
- D. Harvey, S. Leybourne, and P. Newbold. Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13:281–291, 1997.
- E. Jacquier and N. Polson. Asset allocation in finance: A bayesian perspective. In D. Dellaportas, N. Polson, and G. Stephen, editors, *Hierarchical models and MCMC: a Tribute to Adrian Smith*. Oxford University Press, 2012.
- R. Kan and G. Zhou. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 51:385–424, 2007.
- S. Kandel and R. Stambaugh. On the predictability of stock returns: An asset allocation perspective. *Journal of Finance*, 51:385–424, 1996.
- C. Kascha and F. Ravazzolo. Combining Inflation Density Forecasts. *Journal of Forecasting*, 29(1-2):231–250, 2010.
- L. Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, 2009.
- L. Kilian and C. Park. The impact of oil price shocks on the u.s. stock market. *International Economic Review*, 50(4):1267–1287, 2009.

- Lutz Kilian and Clara Vega. Do energy prices respond to u.s. macroeconomic news? a test of the hypothesis of predetermined energy prices. *The Review of Economics and Statistics*, 93(2):660–671, 2011.
- Y. Kitamura. Econometric Comparisons of Conditional Models. Discussion paper,, University of Pennsylvania, 2002.
- J. Mitchell and S. G. Hall. Evaluating, comparing and combining density forecasts using the KLIC with an application to the Bank of England and NIESER “fan” charts of inflation. *Oxford Bulletin of Economics and Statistics*, 67:995–1033, 2005.
- D. Pettenuzzo and A. Timmermann. Predictability of stock returns and asset allocation under structural breaks. *Journal of Econometrics*, forthcoming, 2011.
- F. Ravazzolo and S.V. Vahey. Forecast densities for economic aggregates from disaggregate ensembles. Technical Report 2010/02, Norges Bank, 2012.
- F. Ravazzolo, R. Paap, D. van Dijk, and P.H. Franses. Bayesian model averaging in the presence of structural breaks. In D. Rapach and M. Wohar, editors, *Forecasting in the Presence of Structural Breaks and Uncertainty*, volume 3. Emerald Publishing Group, 2007.
- T. Tokuda, B. Goodrich, I. Van Mechelen, A. Gelman, and F. Tuerlinckx. Visualizing distributions of covariance matrices. Technical report, Columbia, 2012.
- I. Welch and A. Goyal. A Comprehensive Look at the Empirical Performance of Equity Premium prediction. *Review of Financial Studies*, 21(4):253–303, 2008.
- K. D. West, H. J. Edison, and D. Cho. A utility-based comparison of some models of exchange rate volatility. *Journal of International Economics*, 35(1-2):23–45, 1993.

K.D. West. Asymptotic inference about predictive ability. *Econometrica*, 64:1067–1084, 1996.

Appendix

DCC: a Bayesian algorithm

The DCC model for the $(N \times 1)$ vector y_t is formulated as:

$$\begin{aligned}
 y_t &= \Phi(L)y_{t-1} + v_t \\
 v_t &= H_t^{0.5}\varepsilon_t, \quad \varepsilon_t \sim N(0, I_N), \quad H_t = D_t R_t D_t \\
 D_t^2 &= \text{diag}\{\omega_i\} + \text{diag}\{\kappa_i\}v_{t-1}v'_{t-1} + \text{diag}\{\lambda_i\}D_{t-1}^2, \quad i = 1, \dots, N \\
 Q_t &= S(u' - A - B) + A\varepsilon_{t-1}\varepsilon'_{t-1} + BQ_{t-1} \\
 R_t &= \text{diag}\{Q_t\}^{-1}Q_t\text{diag}\{Q_t\}^{-1}
 \end{aligned} \tag{14}$$

where S is the unconditional correlation matrix of ε_t , A , B and $S(u' - A - B)$ are positive semidefinite matrices. In our exercise $N = 2$, therefore the parameters A and B reduce to scalar and above conditions to $A > 0$, $B > 0$, $A + B < 1$. Following Engle [2002], the log likelihood can be expressed as:

$$\ln L = -\frac{1}{2} \sum_{t=1}^T (N \ln(2\pi) + 2 \ln |D_t| + v'_t D_t^{-1} D_t^{-1} v_t - \varepsilon_t \varepsilon'_t + \ln |R_t| + \varepsilon'_{t-1} R_t^{-1} \varepsilon_{t-1}) \tag{15}$$

We estimate the DCC model using a Metropolis-Hastings algorithm. Define the vector $\alpha_i = (\Phi_0, \dots, \Phi_L, \omega_i, \kappa_i, \lambda_i, A, B)'$, with $\Phi(L) = (\Phi_0, \dots, \Phi_L)$, $i = 1, \dots, N$, and α_j the j -th element of it. The sampling scheme consists of the following iterative steps.

Step 1: At iteration s , generate a point α_j^* from the random walk kernel

$$\alpha_j^* = \alpha_j^{s-1} + \epsilon_j, \quad \epsilon \sim N(0, Q), \tag{16}$$

where Q is a diagonal matrix and σ_j^2 is its j -th diagonal element, and α_j^{s-1} is the $(s - 1)$ th iterate of α_j . Therefore, we draw row elements of Φ_0, \dots, Φ_L and $\omega_i, \kappa_i, \lambda_i, A, B$

independently. Then accept α_j^* as α_j^s with probability $p = \min [1, f(\alpha_j^*)/f(\alpha_j^{s-1})]$, where $f()$ is the likelihood of model (14) times priors. Otherwise, set $\alpha_j^* = \alpha_j^{s-1}$. The elements of Q are tuned by monitoring the acceptance rate to lie between 25% and 50%.

Step 2: After M iterations, we apply the following independent kernel MH algorithm.

Generate α_j^* from

$$\alpha_j^* = \mu_{\alpha_j}^{i-1} + \epsilon_j, \quad \epsilon_j \sim N(0, Q_{\alpha_j}), \quad (17)$$

where μ_{α_j} and Q_{α_j} are, respectively, the sample mean and the sample covariance of the first M iterates for α_j . Then accept α_j^* as α_j^i with probability

$$p = \min \left[1, \frac{f(\alpha_j^*)g(\alpha_j^{s-1})}{f(\alpha_j^{s-1})g(\alpha_j^*)} \right], \quad (18)$$

where $g()$ is a Gaussian proposal density (17).

Priors

We set normal priors for $\Phi(L)$ with mean and variance equal to OLS estimates. The priors for $\omega_i, \kappa_i, \lambda_i, A, B$ are uniform distributed and satisfy the restrictions $\omega_i > 0, \kappa_i > 0, \lambda_i > 0, \kappa_i + \lambda_i < 1, A > 0, B > 0, A + B < 1$. We note that different priors for the coefficients A and B of the correlation matrix should be considered if the dimension of the model is larger than two, see discussion in, e.g., Tokuda et al. [2012].

Table 1: Oil and stock return forecast accuracy

Oil return												
	h=1				h=2				h=4			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
MSPE	15.628	0.996	0.954	0.956	15.525	1.040	0.986	0.987	15.447	1.044	1.002	0.993
LS	-2.794	0.999	0.987	0.978*	-2.792	1.007	0.995	0.968**	-2.790	1.007	0.997	0.985**
CRPS	2.087	1.001	0.986	0.978*	2.085	1.015	0.998	0.985	2.075	1.020	1.003	0.991
	h=8				h=12				h=24			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
MSPE	15.503	1.026	0.986	0.981	15.560	0.993	0.998	0.992	15.564	1.006	1.004	1.004
LS	-2.794	1.004	0.996	0.980*	-2.793	1.001	1.001	0.986*	-2.795	1.004	1.003	1.007
CRPS	2.083	1.011	0.992	0.982*	2.079	0.997	0.998	0.989**	2.072	1.004	1.001	0.999
Stock return												
	h=1				h=2				h=4			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
MSPE	8.054	1.025	1.038	1.043	8.095	1.009	1.032	1.033	8.084	1.024	1.042	1.045
LS	-2.505	1.005	1.012	1.028	-2.496	1.010	1.015	1.009	-2.495	1.015	1.021	1.077
CRPS	1.451	1.020	1.025	1.000	1.454	1.016	1.022	1.006	1.456	1.018	1.027	1.018
	h=8				h=12				h=24			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
MSPE	8.218	1.007	1.016	1.011	8.289	0.997	1.002	1.000	8.497	1.001	1.000	0.999
LS	-2.511	1.010	1.016	1.164	-2.518	1.009	1.012	1.104	-2.534	1.020	1.022	1.203
CRPS	1.466	1.005	1.007	1.008	1.472	0.996	0.999	1.001	1.491	1.003	1.001	1.017

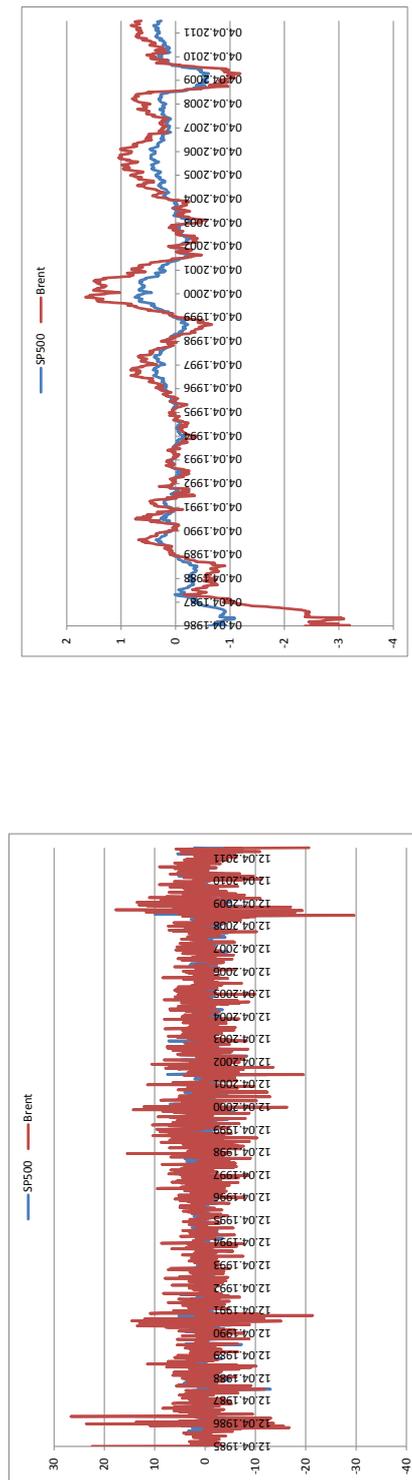
Notes: *RW*, *AR*, *VAR*, *DCC*: individual models defined in Section 3. MSPE: Mean Square Prediction Error. LS: average Logarithmic Score. CRPS: cumulative rank probability score. Absolute statistics for the columns RW; relative statistics to the RW for the other models. One * and two ** represent rejections of the null hypothesis of equal predictability following the Harvey et al. [1997] type of test at 10% and 5% respectively.

Table 2: Economic value

	h=1				h=2				h=4			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
Mean Ret	5.09	13.12	30.23	20.32	2.46	7.45	19.04	16.45	0.91	2.43	8.17	9.89
St dev	18.03	20.96	23.00	20.65	18.31	18.97	20.52	18.65	18.65	22.58	20.04	20.57
SR	0.16	0.54	1.26	0.91	0.01	0.28	0.85	0.79	-0.08	0.01	0.30	0.38
Δ_s	-2.87	101.61	274.18	174.78	-46.94	-24.71	161.82	164.76	-95.37	-136.03	24.57	134.36
Δ_o	140.35	244.84	417.40	318.01	166.28	188.52	375.05	377.99	189.27	148.62	309.22	419.00
Δ_f	66.85	171.34	343.90	244.51	64.02	86.26	272.79	275.73	54.13	13.47	174.07	283.85
	h=8				h=12				h=24			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
Mean Ret	1.06	0.34	5.35	7.32	1.49	0.27	4.12	5.04	2.59	3.22	3.69	3.33
St dev	19.37	24.58	20.36	20.47	20.68	27.00	22.48	20.24	23.32	26.30	26.14	22.38
SR	-0.07	-0.08	0.15	0.25	-0.04	-0.08	0.08	0.14	0.01	0.03	0.05	0.05
Δ_s	-89.91	-276.56	30.88	191.74	4.70	-213.90	98.18	237.01	207.29	220.95	263.16	271.83
Δ_o	270.13	83.47	390.91	551.78	331.10	112.50	424.59	563.42	449.07	462.72	504.93	513.60
Δ_f	73.50	-113.16	194.28	355.15	105.77	-112.83	199.25	338.08	159.18	172.83	215.04	223.71
Transaction costs												
	h=1				h=2				h=4			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
Mean Ret	2.54	8.11	25.76	15.86	1.45	5.48	17.03	14.64	0.53	1.61	7.44	9.15
St dev	18.01	21.02	23.04	20.65	18.31	18.94	20.54	18.62	18.65	22.59	20.05	20.55
SR	0.01	0.29	1.06	0.68	-0.05	0.17	0.74	0.69	-0.10	-0.03	0.27	0.35
Δ_s	-39.52	29.55	210.03	110.67	-67.23	-64.60	121.02	128.09	-106.38	-159.72	3.67	113.36
Δ_o	103.70	172.77	353.26	253.90	146.00	148.63	334.24	341.32	178.26	124.92	288.31	398.00
Δ_f	30.20	99.27	279.76	180.40	43.74	46.37	231.99	239.06	43.12	-10.23	153.16	262.85
	h=8				h=12				h=24			
	RW	AR	VAR	DCC	RW	AR	VAR	DCC	RW	AR	VAR	DCC
Mean Ret	0.92	0.09	5.12	7.07	1.41	0.16	4.01	4.90	2.56	3.19	3.65	3.28
St dev	19.37	24.58	20.37	20.47	20.68	26.99	22.48	20.23	23.32	26.30	26.14	22.38
SR	-0.08	-0.09	0.14	0.24	-0.05	-0.08	0.08	0.13	0.01	0.03	0.05	0.07
Delta_s	-95.68	-286.88	21.54	181.51	0.80	-219.40	92.57	230.23	205.31	218.50	260.85	268.72
Delta_b	264.36	73.15	381.57	541.55	327.20	107.00	418.97	556.63	447.08	460.27	502.63	510.49
Delta_f	67.73	-123.47	184.94	344.92	101.87	-118.33	193.64	331.30	157.19	170.38	212.73	220.60

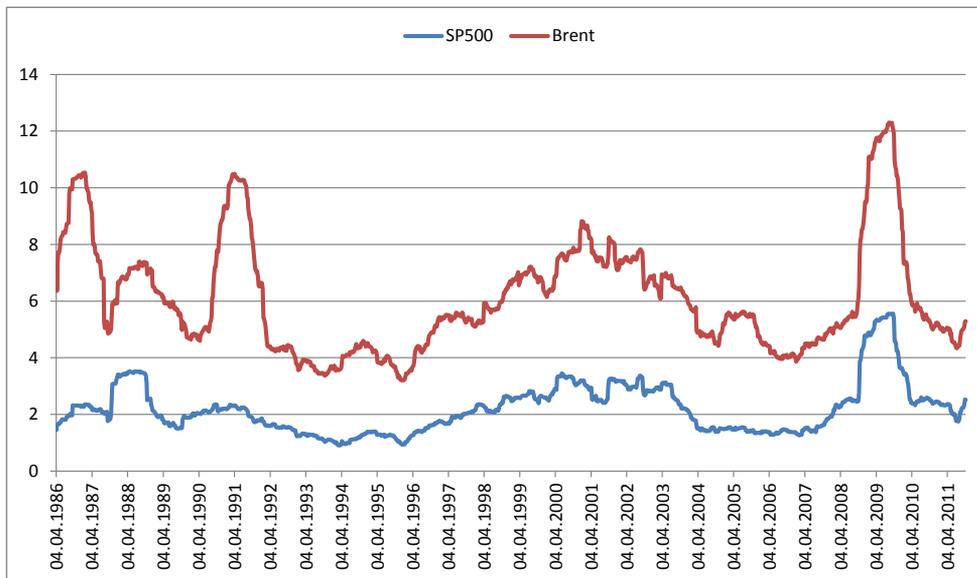
Note: *RW*, *AR*, *VAR*, *DCC*: individual models defined in Section 3. Mean Ret: annualized mean portfolio return. St dev: annualized standard deviation. SR: Sharpe ratio. Δ_s , Δ_o , Δ_f : performance fee from switching from an active strategy to passive strategies holding only stock (r_s), holding only oil (r_o), and only free risk bond (r_b).

Figure 1: Oil and stock returns



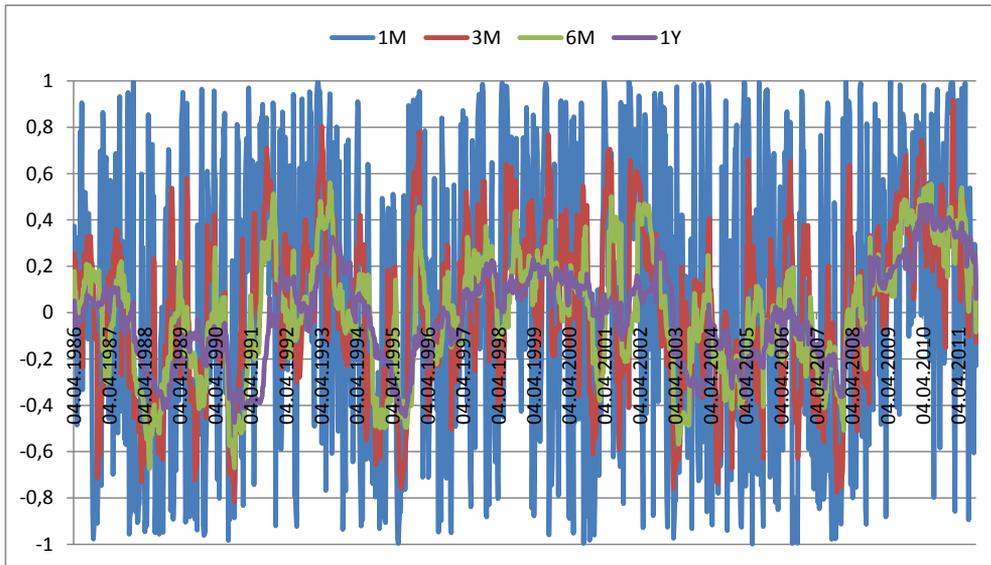
Note: This Figure shows weekly returns of oil and stock prices (left panel) and 1-year moving window average returns (right panel).

Figure 2: Oil and stock volatilities



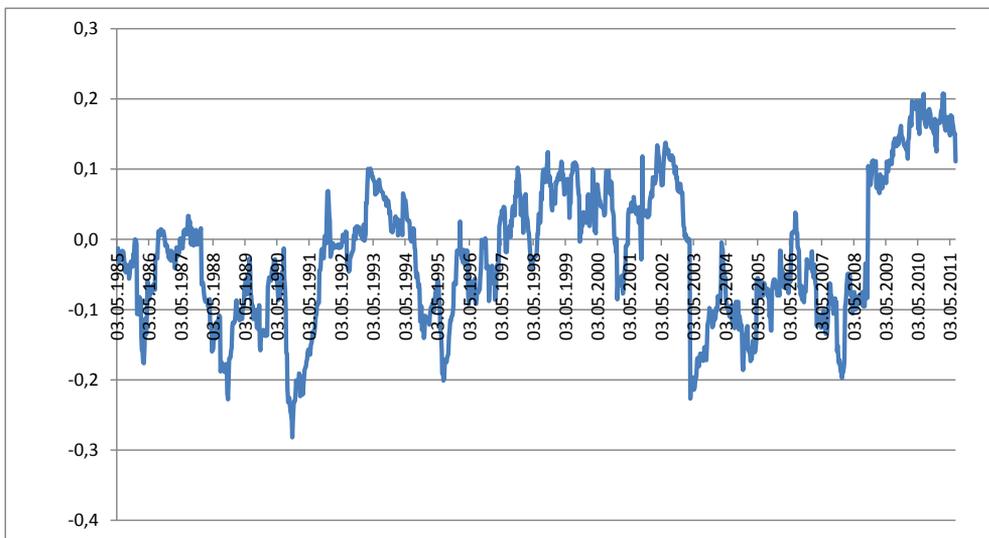
Note: This Figure shows 1-year moving window standard deviations for oil and stock returns.

Figure 3: Oil and stock correlation



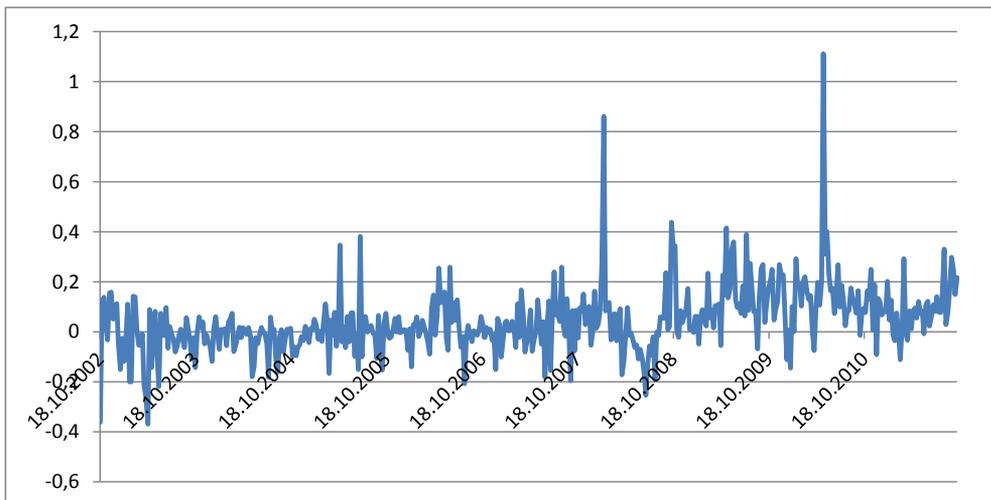
Note: This Figure shows sample correlations computed on moving windows of different length: 1 month (1M); 3-months (3M); 6 months (6M); and 1-year (1Y).

Figure 4: DCC correlation




Note: This Figure shows correlation estimated using a Dynamic Conditional Correlation model.

Figure 5: Weekly realized correlation




Note: This Figure shows weakly realized correlation using hourly data.


Figure 6: CER Differentials



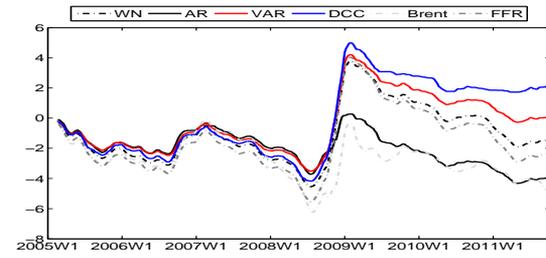
$h = 1$



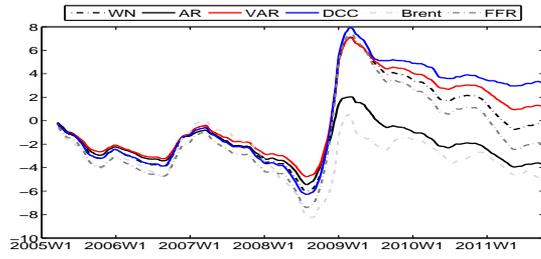
$h = 2$



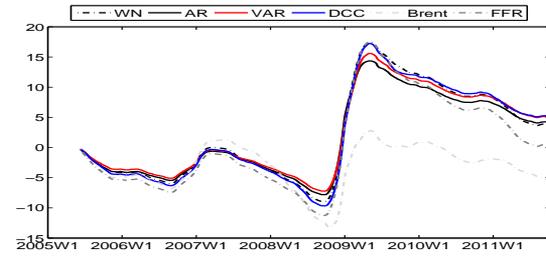
$h = 4$



$h = 8$



$h = 12$



$h = 24$

Note: Certainty Equivalent Return (CER) Differentials versus the CER of the passive strategy holding only stock with transaction costs $c = 0.05\%$.