

Are there Gains from Pooling Real-Time Oil Price Forecasts?

Christiane Baumeister, Bank of Canada

Lutz Kilian, University of Michigan

Thomas K. Lee, U.S. Energy Information Administration

February 12, 2014

This paper is released to encourage discussion and critical comment. The analysis and conclusions expressed here are those of the authors and not necessarily those of the U.S. Energy Information Administration.

Abstract

The answer depends on the objective. The approach of combining five of the leading forecasting models with equal weights dominates the strategy of selecting one model and using it for all horizons up to two years. Even more accurate forecasts, however, are obtained when allowing the forecast combinations to vary across forecast horizons. While the latter approach is not always more accurate than selecting the single most accurate forecasting model by horizon, its accuracy can be shown to be much more stable over time. The MSPE of real-time pooled forecasts is between 3% and 29% lower than that of the no-change forecast and its directional accuracy as high as 73%. Our results are robust to alternative oil price measures and apply to monthly as well as quarterly forecasts. We illustrate how forecast pooling may be used to produce real-time forecasts of the real and the nominal price of oil in a format consistent with that employed by the U.S. Energy Information Administration in releasing its short-term oil price forecasts and we compare these forecasts during key historical episodes.

JEL Code: Q43, C53

Key Words: Forecast combination; real-time data; WTI; refiners' acquisition cost; model selection.

Acknowledgements: The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Bank of Canada or the views of the U.S. Energy Information Administration. We thank Dmitri Tchebotarev for excellent research assistance.

Corresponding author: Lutz Kilian, Department of Economics, 611 Tappan Street, Ann Arbor, MI 48109-1220, USA. Email: lkilian@umich.edu.

Christiane Baumeister, International Economic Analysis Department, Bank of Canada, 234 Laurier Avenue West, Ottawa, ON K1A 0G9, Canada. Email: CBaumeister@bankofcanada.ca.

Thomas K. Lee, U.S. Energy Information Administration, 1000 Independence Avenue, Washington, DC 20585, USA. Email: Thomas.Lee@eia.gov.

1. Introduction

Accurate forecasts of the price of oil are important to firms and consumers as well as state and national governments. There are many alternative approaches to forecasting oil prices ranging from the use of oil futures prices and survey forecasts to atheoretical time series models and econometric models.¹ Our approach in this paper is to focus on short-term oil price forecasting models that can be motivated based on economic grounds. Among these models we restrict attention to models that have been shown in the literature to produce more accurate real-time forecasts than the random walk benchmark model at least for some forecast horizons. Our objective is to examine the forecast accuracy of weighted averages of these models, as measured by the mean-squared prediction error (MSPE) at monthly and quarterly horizons up to two years. We also report results for the directional accuracy of these combined forecasts.

Forecast combinations (also known as pooled forecasts) have a long tradition in macroeconomic forecasting (see, e.g., Timmermann 2006). With regard to short-term oil price forecasts, Baumeister and Kilian (2013a) established that an equal-weighted combination of four recently proposed oil price forecasting models is systematically more accurate than the no-change forecast as well as forecast combinations based on recursive or rolling inverse MSPE weights. The forecasting models considered in that study are included a vector autoregressive (VAR) forecast, forecasts based on the spread between oil futures prices and the spot price of oil, forecasts based on non-oil industrial commodity prices, and forecasts based on a time-varying parameter (TVP) model of the spreads between the U.S. spot prices of gasoline and heating oil and the spot price of crude oil. More recent work by Baumeister, Guérin and Kilian (2014), which explored the predictive content of high-frequency data from financial and energy markets, uncovered evidence that an important additional source of real-time information about future oil prices is the cumulative change in U.S. crude oil inventories. In the current paper, we extend the set of models to be combined to include the latter forecast, which performs particularly well at horizons between one and two years.

Baumeister and Kilian (2013a) compared equal-weighted forecast combinations to individual forecasting models and showed that only pooled forecasts are systematically more accurate than the no-change forecast at all horizons up to 18 months or 6 quarters. This approach satisfies the requirement that there should be no discontinuities across horizons in the forecast path of the type that would arise if we switched forecasting models or forecast combinations from one horizon to the next. Using the same approach as in Baumeister and Kilian (2013a), in this paper we show that including in addition forecasts based on U.S. crude oil inventories in the forecast combination substantially improves the accuracy of the pooled forecast at horizons between one and two years.

¹For a comprehensive review the reader is referred to Alquist, Kilian and Vigfusson (2013). Subsequent contributions include Baumeister and Kilian (2013a,b, 2014), Baumeister, Kilian and Zhou (2013), Chen (2013), Baumeister, Guérin and Kilian (2014), and Bernard, Khalaf, Kichian and Yelou (2014), among others.

Insisting on a continuous forecast path comes at the price of higher MSPEs, however. If a low MSPE is all we care about in forecasting, one clearly can improve on equal-weighted combinations of all oil price forecasting models. In this paper, we extend the analysis of pooled forecasts by allowing for different forecast combinations to be chosen at each horizon. This strategy takes advantage of the fact that some oil price forecasting models perform well at short horizons, but were never intended for longer horizons, whereas other models perform best at longer horizons. We show that relaxing the continuity constraint substantially reduces the MSPE of the pooled forecast at all horizons, but especially at horizons beyond one year.

This fact raises the question of how forecast pooling by horizon compares with simply selecting for each horizon the individual forecasting model with the lowest MSPE. The latter comparison is the relevant benchmark when evaluating the benefits of pooling in the absence of the continuity constraint. We find that pooled forecasts often, but not always have lower MSPE than the best individual forecast. The superior accuracy of the forecast combination at some horizons is not surprising in that pooled forecasts provide insurance against failures of individual models. Our results show that this insurance has a price in the form of lower forecast accuracy in some dimensions, however. For example, at horizons beyond 18 months, the individual forecasts are clearly more accurate. This drawback is offset by the fact that the accuracy of the pooled forecasts is more stable over time, as revealed by plots of the recursive MSPE ratios over time.

The MSPE of the real-time pooled forecasts is up to 29% lower than that of the no-change forecast even at horizons as high as two years. The pooled forecasts also predict the direction of change in the real price of oil correctly with probabilities as high as 73%. Our qualitative results are robust to alternative oil price measures and apply to monthly as well as quarterly forecasts. In addition to presenting these summary statistics, we use graphical methods to examine how the pooled real-time forecasts performed in recent years when the real price of oil fluctuated substantially. We compare these model-based pooled forecasts to the U.S. Energy Information Administration's (EIA) short-term oil price forecasts, as released in the Short-Term Energy Outlook. Finally, we discuss how real-time pooled forecasts of the nominal oil price may be derived from the forecasts of the real price, and we illustrate that both real and nominal oil price forecasts may be presented in a format already used by the EIA.

The remainder of the paper is organized as follows. In section 2 we review the forecasting models considered. Section 3 evaluates our monthly forecasts of the real U.S. refiners' acquisition cost for oil imports and of the West Texas Intermediate (WTI) price of crude oil. In section 4 we extend the analysis to quarterly forecasts. Section 5 examines how stable the accuracy of these oil price forecasts is over time. In section 6 we visually compare the accuracy of our pooled oil price forecasts to that of the EIA oil price forecasts during key episodes. In section 7 we illustrate how these forecasting tools may be used to produce real-time forecasts of the real and the nominal price of oil in a format consistent with that employed by the EIA in releasing its short-term oil price forecasts. The concluding remarks are in section 8.

2. The forecasting environment

All forecasting models are estimated at monthly frequency. We consider monthly forecast horizons up to two years. Forecasts at the corresponding quarterly horizons are obtained by averaging the monthly forecasts at quarterly frequency, as recommended in Baumeister and Kilian (2013b). The forecasting models are estimated recursively and subject to real-time data constraints. All data are obtained from the real-time database developed in Baumeister and Kilian (2012, 2013a) and extended in Baumeister, Kilian and Zhou (2013). The reader is referred to the latter references for a detailed description of the data sources and definitions. The evaluation period is January 1992 through September 2012 (or equivalently the first quarter of 1992 through the third quarter of 2012). Our objective is to forecast the ex-post revised real price of oil, as measured by the observations in the March 2013 vintage of the real-time database.

The real-time forecasts are evaluated based on their recursive MSPE and their directional accuracy, as measured by the success ratio. The success ratio is the fraction of times that a method correctly predicts the direction of change in the real price of oil. Success ratios above 0.5 indicate an improvement relative to the no-change forecast. The MSPE results are normalized relative to the no-change forecast, with a ratio below 1 indicating a gain in accuracy. There is no valid test for judging the statistical significance of the MSPE reductions in our context, but we examine the stability of our results across horizons, across specifications and over time. The statistical significance of the success ratios is assessed based on the test proposed in Pesaran and Timmermann (2009).

Building on the comprehensive analysis of forecast combination methods in Baumeister and Kilian (2013a), we consider five forecasting models with proven credentials.

2.1. Forecasts based on a VAR model of the global oil market

The first model is a reduced-form VAR model of the form:

$$B(L)y_t = \nu + u_t$$

where $y_t = [\Delta prod_t, rea_t, r_t^{oil}, \Delta inv_t]'$ refers to a vector including the percent change in global crude oil production, a measure of global real economic activity, the log of the U.S. refiners' acquisition cost for crude oil imports deflated by the log of the U.S. CPI, and the change in global crude oil inventories, ν denotes the intercept, $B(L) = I_4 - B_1L - \dots - B_pL^p$ denotes the autoregressive lag order polynomial, p is the autoregressive lag order, L is the lag operator, and u_t is a white noise innovation.² This VAR model may be viewed as the reduced-form representation of the structural global oil market model developed in Kilian and Murphy (2013).

We estimate the unrestricted VAR model with 12 autoregressive lags by the method of least squares. Forecasts $\hat{r}_{t+h|t}^{oil}$ of the log of the real price of oil are constructed iteratively from the estimated VAR model conditional on the most recent data and converted to levels, resulting in the forecast:

² The inventory data are constructed by multiplying U.S. crude oil inventories by the ratio of OECD petroleum inventories to U.S. petroleum inventories. Petroleum inventories are defined to include both stocks of crude oil and stocks of refined products. The global real activity index is constructed from data on global dry cargo ocean shipping freight rates as described in Kilian (2009).

$$\hat{R}_{t+h|t}^{oil} = \exp(\hat{r}_{t+h|t}^{oil,VAR}) \quad (1)$$

Forecasts for the real WTI price are constructed from the same VAR model by assuming that the most recent spread between the log WTI price and the log of the U.S. refiner's acquisition cost remains unchanged in the future. By rescaling the forecasts of the U.S. refiners' acquisition cost in this manner, we allow the relationship between the two oil price measures to evolve as a random walk (see Baumeister and Kilian 2013b).

2.2. Forecasts based on the price of non-oil industrial raw materials

An alternative forecasting method – based on the intuition that there are broad-based predictable shifts in the demand for globally traded commodities – exploits real-time information from recent cumulative changes in non-oil industrial commodity price indices. As discussed in Baumeister and Kilian (2012), such a forecast of the real price of oil may be constructed as follows:

$$\hat{R}_{t+h|t}^{oil} = R_t^{oil} \left(1 + \pi_t^{h, industrial\ raw\ materials} - E_t(\pi_{t+h}^h) \right) \quad (2)$$

where R_t^{oil} denotes the current level of the real price of oil and $\pi_t^{h, industrial\ raw\ materials}$ stands for the percent change of the Commodity Research Bureau (CRB) index of the spot price of industrial raw materials (other than oil) over the preceding h months. The term $E_t(\pi_{t+h}^h)$ is the expected U.S. inflation rate over the next h periods. In practice, this expectation is proxied by recursively constructed averages of past U.S. CPI inflation data, starting in July 1986.³

2.3. Forecasts based on oil futures prices

Yet another approach is to exploit information from oil futures markets. Many practitioners rely on the price of oil futures contracts in generating forecasts of the nominal price of oil or, equivalently, on the futures spread for forecasting the change in the nominal price of oil. This forecast can then be converted to a forecast for the real price of oil by subtracting expected inflation. This approach is embodied in the forecasting model

$$\hat{R}_{t+h|t}^{oil} = R_t^{oil} \left(1 + f_t^h - s_t - E_t(\pi_{t+h}^h) \right), \quad (3)$$

where f_t^h is the log of the current WTI oil futures price for maturity h , s_t is the log of the corresponding WTI spot price, and $E_t(\pi_{t+h}^h)$ is again the expected inflation rate over the next h periods. Monthly WTI oil futures price data for our evaluation period are available only up to a horizon of 18 months. This means that for horizons beyond 18 months the futures-based forecast receives zero weight in the forecast combinations.

³ Undoubtedly, the inflation forecast could be refined further, but there is little loss in generality in our approach, as shown in Baumeister and Kilian (2012).

2.4. Time-varying parameter model of the gasoline and heating oil spreads

Many market practitioners believe that rising spreads between the price of refined products (such as gasoline or heating oil) and the price of crude oil signal upward pressures on the price of crude oil. For example, Goldman Sachs in April 2013 cut its oil price forecast, citing significant pressure on product spreads, which it interpreted as an indication of reduced demand for products (see Strumpf 2013). There are many reasons to expect forecasts based on product spreads to be unstable over time. One concern is that the price of crude oil is likely to be determined by the refined product in highest demand and that product has changed over time. Another concern is that crude oil supply shocks, local capacity constraints in refining, changes in environmental regulations, or other market turmoil may all temporarily undermine the predictive power of product spreads. We therefore follow Baumeister et al. (2013) in first recursively estimating the time-varying regression model

$$\Delta s_{t+h|t} = \beta_{1t} [s_t^{gas} - s_t] + \beta_{2t} [s_t^{heat} - s_t] + \varepsilon_{t+h}$$

where the s_t^{gas} is the log of the nominal U.S. spot price of gasoline and s_t^{heat} is the log of the nominal U.S. spot price of heating oil.⁴ Given the TVP estimates, we then construct the TVP model forecast:

$$\hat{R}_{t+h|t}^{oil} = R_t^{oil} \exp \left\{ \hat{\beta}_{1t} [s_t^{gas} - s_t] + \hat{\beta}_{2t} [s_t^{heat} - s_t] - E_t(\pi_{t+h}^h) \right\} \quad (4)$$

by Monte Carlo integration as the mean of the forecasts simulated based on 1,000 Gibbs iterations conditional on the most recent data.⁵

2.5. Forecasts based on U.S. crude oil inventories

The final forecasting model can be motivated by the theoretical analysis in Alquist and Kilian (2010), which shows that changes in crude oil inventories capture shifts in expectations about the future real price of oil. As shown in Baumeister, Guérin and Kilian (2014), forecasts of the real price of oil may be constructed from U.S. crude oil inventory data as follows:

$$\hat{R}_{t+h|t}^{oil} = R_t^{oil} \left(1 + \hat{\beta} \Delta inv_t^h \right), \quad (5)$$

⁴In estimating the model, we postulate that $\varepsilon_{t+h} \sim NID(0, \sigma^2)$, while the time-varying coefficients $\theta_t = [\beta_{1t} \ \beta_{2t}]'$ evolve according to a random walk as $\theta_t = \theta_{t-1} + \xi_t$, and ξ_t is independent Gaussian white noise with variance Q. The intercept has been restricted to zero, following Baumeister et al. (2013) who show that this restriction greatly improves the out-of-sample accuracy. This state-space model is estimated using a Gibbs sampling algorithm. The conditional posterior of θ_t is normal, and its mean and variance can be derived via standard Kalman filter recursions (see Kim and Nelson 1999). Conditional on an estimate of θ_t , the conditional posterior distribution of σ^2 is inverse Gamma and that of Q is inverse Wishart.

⁵Our forecasts take into account that the model parameters continue to drift over the forecast horizon according to their law of motion. The first 30 observations of the initial estimation period are used as a training sample to calibrate the priors and to initialize the Kalman filter.

where Δinv_t^h denotes the percent change in U.S. crude oil inventories over the preceding h months and $\hat{\beta}$ is obtained by regressing cumulative percent changes in the real price of oil on the lagged cumulative percent change in U.S. inventories without intercept. The latter restriction improves the accuracy.

3. Monthly forecasts of the real price of oil

It is useful to focus on the monthly forecasts of the real price of oil first, before considering the quarterly forecasts.

3.1. U.S. refiners' acquisition cost for crude oil imports

The U.S. refiners' acquisition cost for crude oil imports is a commonly used proxy for the global price of crude oil and as such is of central interest for oil price forecasters. Table 1 shows that an equal-weighted combination of all five oil price forecasting models unambiguously reduces the MSPE of the forecast at all horizons relative to the no-change forecast much like the simpler forecast combination considered in Baumeister and Kilian (2013a). The difference is that the larger forecast combination including the additional forecasting model based on U.S. crude oil inventories generates further improvements at horizons beyond 18 months. Not only are there MSPE reductions at these long horizons, but there also is statistically significant directional accuracy.

The MSPE reductions in Table 1 for this forecast combination occur at all horizons and are reasonably large by the standards of the literature on oil price forecasting. The largest MSPE reduction is 11%. The largest success ratio is 68%. Table 2 reveals which models contribute to the improved accuracy at which horizon. It shows how the MSPE ratio of the combination forecast changes as we drop one model at a time from the pooled forecast. Increases in the ratio mean that the model left out would have improved forecast accuracy if included, whereas decreases mean that it would have worsened forecast accuracy. For example, the VAR improves forecast accuracy at horizons 1 through 7 and industrial raw materials prices at horizons 1 through 5. Oil futures prices worsen accuracy at short horizons, but improve accuracy at horizons 5 through 18, while the TVP product spread model improves forecast accuracy at horizons 5 through 24. Finally, the inventory model improves accuracy only at horizons 14 through 24.

Table 2 shows that clearly not every forecasting model is suitable for every horizon. A natural proposal is to select an improved forecast combination for each horizon by dropping all models that lower forecast accuracy in Table 2 and by retaining the others. The second set of results in Table 1 confirms that choosing equal-weighted subsets by horizon further improves forecast accuracy. In that case, the largest MSPE reduction increases from 11% to 29% and the largest success ratio from 68% to 73%. These statistics are very large by the standards of the oil price forecasting literature. The most stunning improvement in accuracy occurs beyond the 18-month horizon. It is related to the fact that we dropped the VAR forecast and the forecast based on industrial commodity prices which were never intended to be accurate at such long horizons (see Baumeister and Kilian 2012).

The difference between the first and the second set of results in Table 1 is that we have relaxed the implicit constraint that forecasts across horizons should be generated by the same forecasting model or the same forecast combination. Relaxing this constraint by construction must reduce the MSPE of the pooled forecast. This raises the question of what the benefits are of forecast pooling in the absence of that continuity restriction. If we are not concerned with continuity the relevant benchmark is the

forecast accuracy obtained by generating the most accurate individual oil price forecasting model for each horizon. The third set of results in Table 1 shows that by selecting the best individual forecasting model, we can reduce the MSPE even further. In that case, the MSPE reductions in Table 1 reach 36% in some cases and the success ratios an astonishing 75%. This means that we predict the direction of change in the real price of oil three times out of four. These impressive gains in accuracy are not uniform across forecast horizons, however. In fact, for many horizons the forecast combination chosen by horizon has slightly lower MSPE than the best individual model. Only at horizons beyond 18 months, the individual forecasting models are clearly more accurate. Thus, these results seem ambiguous and either approach seems reasonable in practice. In section 5 we will present additional evidence, however, that in fact the accuracy of the pooled forecasts is much more stable over time, making them the preferred approach in practice. Before addressing this point, it is useful to examine how robust these findings are to the use of other oil price measures.

3.2. WTI spot price of crude oil

Table 3 shows the corresponding results for the real spot price of WTI crude oil, which is a commonly used reference price in oil markets. Unlike the U.S. refiners' acquisition cost, the nominal WTI price is available without delays and not subject to revisions. It has been subject to regulation until the 1980s, however, and has recently become unrepresentative of the global price of crude oil.

Table 3 documents that the equal-weighted combination of all five models is about as accurate for the real WTI price as for the real refiners' acquisition cost. The largest MSPE reduction is 12% and the largest success ratio 67%. Many of the gains in directional accuracy are statistically significant, even at horizons as high as 24 months. Some are not, especially between horizons of 5 and 10 months. The pooled forecast generates MSPE reductions at every horizon between 1 month and 24 months.

To conserve space, we do not present a table showing the contributions of each forecasting model to the pooled forecast analogous to Table 2, but note that the pattern is very similar except that the VAR model contributes to the accuracy of the forecast combination at all horizons but horizons 16, 17, and 18. Allowing the forecast combination to vary by horizon increases the MSPE reductions at all horizons, as in Table 1. At some horizons the MSPE reductions reach 19% relative to the no-change forecast. The largest success ratio is 70%. Compared with the results in Table 1, these gains in accuracy are still substantial, but more modest.

Finally, when choosing the best individual forecasting model by horizon, the MSPE reductions reach 30% in some cases and the success ratios 82%. These large gains are all obtained at horizons beyond 18 months. At shorter horizons, the relative performance of this last method compared with the forecast combination chosen by horizon is ambiguous and the differences are modest. We conclude that our findings are robust with respect to the choice of oil price measure.

4. Quarterly forecasts of the real price of oil

For many purposes, users of oil price forecasts are interested in quarterly rather than monthly forecasts. It is straightforward to construct forecasts at quarterly horizons by averaging the monthly forecasts. This approach is more accurate than constructing quarterly forecasts from models estimated at quarterly frequency (see Baumeister and Kilian 2013b). It is important to note that the accuracy of quarterly forecasts may not be inferred from that of monthly forecasts because of the unknown covariance of the monthly forecasts across horizons and because averaging lowers the variance of the forecasts.

The upper panel of Table 4, nevertheless, suggests that quarterly forecasts are about as accurate as monthly forecasts. The success ratios are systematically above 0.5 and all statistically significant except for the longest horizons. All MSPE ratios are below 1. The maximum MSPE reductions of the pooled forecast range from 12% for the equal-weighted combination of all models to 26% for the combination chosen by horizon. The corresponding success ratios may be as high as 70% and 71%. Similar results hold for the real WTI price in the lower panel. For the best individual model forecasts chosen by horizon, in the upper panel of Table 3, the MSPE reductions may be as large as 35% and the success ratio as high as 74% at the quarterly frequency. The corresponding results for the real WTI price are 29% and 71%. We conclude that our results are not affected by the choice between monthly and quarterly forecast horizons.

5. How stable is the accuracy of pooled forecasts compared to other forecasts?

One of the perceived benefits of pooling oil price forecasts is that the forecast accuracy tends to become more stable over time than forecasts from individual models. Tables 1 and 3 illustrate that substantial further reductions in the MSPE ratio are feasible when selecting the preferred forecast combination or the preferred forecasting model by horizon. These accuracy gains, however, may come at the expense of instability over time. Figure 1 examines this concern by plotting the recursive real-time MSPE ratios over time. We focus on the real U.S. refiners' acquisition cost for oil imports, but note that similar results hold for the real WTI price. In recognition of the fact that MSPE ratios are uninformative when based on too short of an evaluation period, we discard the first five years of the evaluation period. The last observation shown in Figure 1 corresponds to the MSPE ratios shown in Table 1. To conserve space, we focus on horizons of 1, 3, 6, 9, 12, 15, 18, 21 and 24 months.

Figure 1 illustrates that the equal-weighted combination of all five forecasting models virtually always has been more accurate than the no-change forecast at every horizon shown, especially since 2000. To a lesser extent this has also been true for the subset of models chosen for each horizon. Figure 1 shows that the reductions in the overall MSPE associated with that second method come at the expense of a somewhat more erratic forecast performance. Overall, however, the accuracy of this method still remains quite robust. Only occasionally the recursive MSPE ratios become positive. Such episodes are short-lived and the losses in accuracy are comparatively small.

In sharp contrast, the accuracy of the third method we considered, which involves selecting the single most accurate forecasting model by horizon, is highly erratic. Especially at horizons 6, 9, 12, and 15, this method may generate recursive MSPE ratios as high as 1.7 in some cases, making this approach unfit for applied work, despite its strong performance on average over the evaluation sample. We conclude that there are clear gains from pooling forecasts not captured by the MSPE ratios in Tables 1 and 3. Very similar results also hold for quarterly forecasts, as shown in Figure 2.

6. Pooled forecasts and EIA forecasts in historical perspective

A different perspective on the accuracy of pooled real-time forecasts may be obtained by comparing them to the short-term oil price forecasts published by the EIA. The data source for the latter forecasts is the EIA's *Short-Term Energy Outlook*. Most importantly for our purposes, this publication provides quarterly forecasts of the U.S. refiners' acquisition cost for imports for horizons up to 6 quarters. For the purpose of comparing the EIA forecasts with the pooled forecasts, the nominal EIA oil price forecasts have been deflated by the same measure of expected inflation used in constructing the pooled forecasts of the real price of oil. Figure 3 compares the forecasts produced by the EIA against the realizations of the real price of oil. For each quarter, for which a forecast is available, we show the entire forecast path. The EIA nowcast for the current quarter is marked as a circle. In Figure 4 we conduct a similar comparison for the pooled real-time forecast selected by horizon.

Figures 3 and 4 illustrate the differences between the pooled forecast and the EIA's forecast. For example, the plots show that the quarterly nowcasts produced by the EIA tend to differ from the realizations much more than the nowcasts constructed in the Baumeister and Kilian (2012) real-time database. There is little systematic difference between the EIA forecasts and the pooled forecasts during 1992-2002, except that the EIA forecasts often are based on misleading nowcasts, causing them to overstate or understate the level of the real price of oil. With the surge in the real price of oil starting in 2003, however, strong qualitative differences emerge. Whereas the EIA's forecast paths for the next five years always point downward, the pooled forecasts almost always predict increases at least in the short run. This is particularly evident in the first quarter of 2008, when the pooled forecast anticipates a sharp increase in the real price of oil, while the EIA predicted a strong decline. Similarly, the pooled forecast is better at predicting the downturn of the real price of oil in late 2008 and the recovery in 2009. While neither the pooled forecast nor the EIA forecast is good at predicting turning points, the pooled forecast quickly adapts, once a turning point has occurred. A case in point is the third quarter of 2008, one quarter after the peak of mid-2008. Whereas the EIA predicted continued oil price increases following a small decline, the pooled forecast correctly anticipated a sharp decline in the real price of oil in the fourth quarter, although not as steep as the actual drop.

7. Real versus nominal oil price forecasts

It can be shown that much of the variation in the nominal price of crude oil reflects variation in the real price of oil. Only at longer horizons, the inflation component matters (see Alquist et al. 2013). Whereas we have focused on the real price of oil, which is the price that ultimately matters for users of oil price forecasts, the EIA publishes nominal oil price forecasts. In this section we illustrate how both the real and the nominal pooled oil price forecasts may be presented in the same format in which the EIA presents its own oil price forecasts. We again focus on the U.S. refiners' acquisition cost for oil imports.

Table 5 shows three examples. It displays forecast paths generated as of the last quarter of 2000, the third quarter of 2008 and the second quarter of 2009. When the forecasts are generated, many prior values of the real price of oil are still unknown and have to be inferred by the forecaster. These nowcasts differ from the actual values of the price of oil revealed later. The forecast path and the corresponding realizations of the ex-post revised price of oil are shown for five quarters, reflecting the limited availability of the EIA forecasts. All forecasts are normalized such that the nominal and the real

price of oil coincide in the last nowcast period. The difference in the forecast paths of the real and the nominal price is expected U.S. inflation. Inflation adjustments for the nowcasts are based on the real-time U.S. inflation data in Baumeister and Kilian (2012). We convert pooled forecasts of the real price of oil to nominal price forecasts by adjusting the predicted changes in the real price of oil by the same expected rate of inflation used in constructing Figure 3.

The first example in Table 5 shows that in the last quarter of 2000 both the EIA forecast and the pooled forecast predicted a decline in the price of oil. One difference is that the EIA relied on a nowcast that was much too high, whereas our nowcast was close to the realized value of the price of oil. The other difference is that the pooled forecast predicted a steeper decline, more in line with the actual evolution of the price of oil.

The other two examples in Table 5 involve situations in which both the EIA nowcasts and our nowcasts are reasonably accurate, allowing us to abstract from the role of the nowcast. The second example shows that in the third quarter of 2008, amidst the financial crisis, the EIA predicted that the price of oil would quickly recover and rise to 122 dollars by mid-2009. The pooled forecast, in contrast, predicted a sharp decline, although not nearly as steep as the actual decline. The third example is the second quarter of 2009, after the recovery of the price of oil had started. The EIA forecast path shows a modest increase, whereas the path predicted by the pooled forecast shows a much steeper increase, closer to the actual evolution of the price of oil.

These three examples make the point that the choice of forecasting methods matter. Clearly, there are other examples in which the EIA forecast and the pooled forecasts are more similar, and there even are rare examples in which the EIA forecast is more accurate than the pooled forecast. Nevertheless, overall, the benefits of relying on model-based real-time forecasts compared with the EIA forecasts are readily apparent.⁶ We conclude that it is straightforward to construct real-time pooled forecasts of both the nominal and the real price of oil in a format the EIA relies on. Table 5 also illustrates that the distinction between real and nominal oil price forecasts matters. While that difference is negligible at short horizons, at the 5-quarter horizon, for example, it accounts for somewhere between four and five dollars in our examples. This difference is small compared with the overall level of the price of oil, but large enough to matter for economic decisions.

8. Concluding remarks

We compared three approaches to generating short-term real-time oil price forecasts. One is a pooled forecast obtained by assigning equal weight to all forecasting models under consideration. Another allows the subset of models selected for the forecast combination to vary by horizons according to its ability to reduce the MSPE. A third approach involves selecting for each horizon the model with lowest recursive MSPE. Of these approaches only the first two can be recommended. While tailoring the forecast combination to each horizon improves the forecast accuracy especially at longer horizons

⁶ We do not formally compare the MSPEs and directional accuracy of the EIA forecast and the pooled forecast. The reader is referred to the related analysis in Baumeister and Kilian (2013a) which shows that the EIA forecasts have much higher MSPE than the no-change forecast and lack directional accuracy. This result overturns the substantive conclusion of Sanders, Manfredo and Boris (2009) based on a much shorter evaluation period and a different econometric approach. Moreover, including the EIA forecast in the forecast combination would systematically lower the accuracy of the pooled forecast.

without seriously impairing the reliability of the pooled forecast, the accuracy of forecasts based on selecting the best individual model at each horizon tends to be unstable over time.

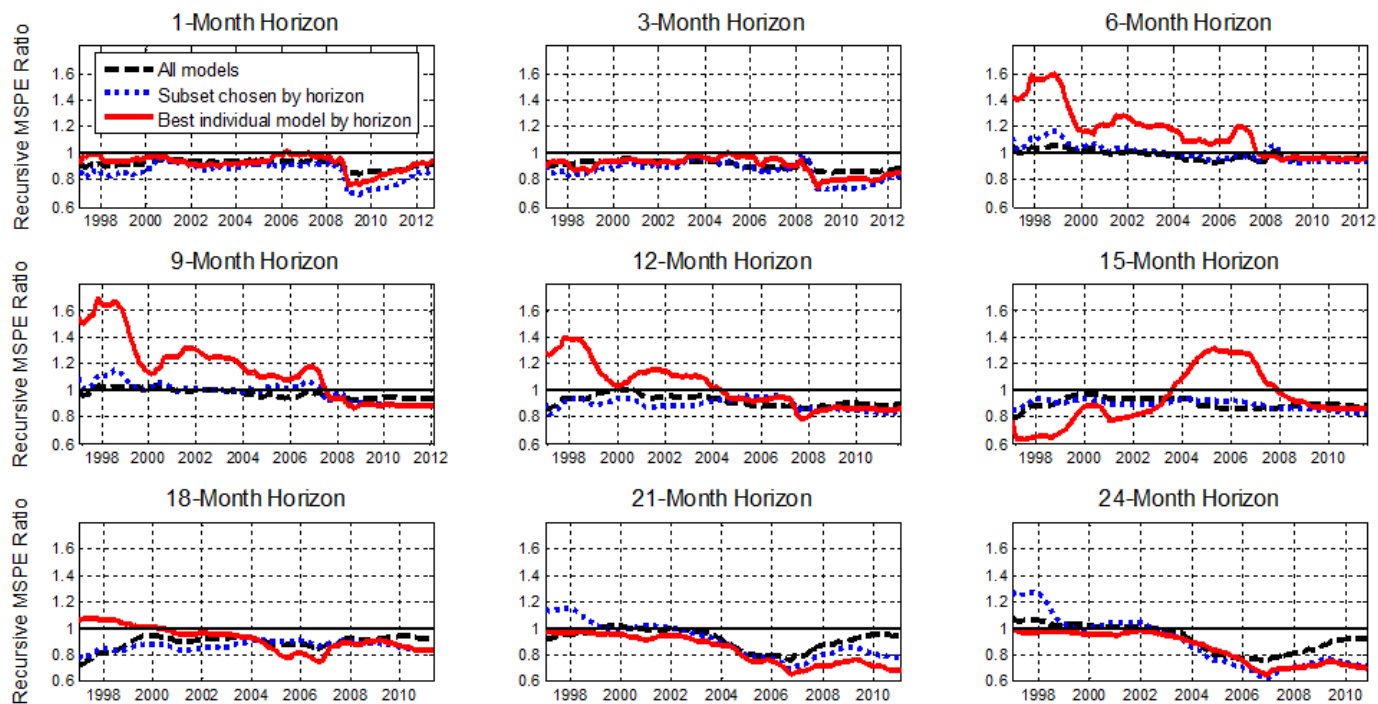
Pooled forecasts were shown to be at least as accurate as the no-change forecast of the real price of oil at all horizons up to two years and often substantially more accurate. We illustrated that pooled oil price forecasts not only are more accurate than the no-change forecast, but also perform better than the EIA's own short-term forecasts during key historical episodes. Our analysis suggests that recently developed model-based forecasts have become a promising alternative to forecasts of the type traditionally employed by the EIA. We illustrated how these forecasting tools may be used to produce real-time forecasts of the real and the nominal price of oil in a format consistent with that employed by the EIA in releasing its short-term oil price forecasts.

References

- Alquist, R., and Kilian, L. (2010), "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, 25, 539-573.
- Alquist, R., Kilian, L., and R.J. Vigfusson (2013), "Forecasting the Price of Oil," in: Elliott, G., and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland, 427-507.
- Baumeister, C., Guérin, P., and L. Kilian (2013), "Do High-Frequency Financial Data Help Forecast Oil Prices? The MIDAS Touch at Work," mimeo, University of Michigan.
- Baumeister, C., and L. Kilian (2012), "Real-Time Forecasts of the Real Price of Oil," *Journal of Business and Economic Statistics*, 30, 326-336.
- Baumeister, C., and L. Kilian (2013a), "Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach," mimeo, University of Michigan.
- Baumeister, C., and L. Kilian (2013b), "What Central Bankers Need to Know about Forecasting Oil Prices," forthcoming: *International Economic Review*.
- Baumeister, C., and L. Kilian (2014), "Real-Time Analysis of Oil Price Risks using Forecast Scenarios," mimeo, University of Michigan.
- Baumeister, C., Kilian, L., and X. Zhou (2013), "Are Product Spreads Useful for Forecasting? An Empirical Evaluation of the Verleger Hypothesis," mimeo, University of Michigan.
- Bernard, J.-T., Khalaf, L., Kichian, M., and C. Yelou (2014), "On the Long-Term Dynamics of Oil Prices: Learning from Combination Forecasts," mimeo, Carleton University.
- Chen, S.S. (2013), "Forecasting Crude Oil Price Movements with Oil-Sensitive Stocks," forthcoming: *Economic Inquiry*.
- Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053-1069.
- Kilian, L., and Murphy, D.P. (2013), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," forthcoming: *Journal of Applied Econometrics*.
- Kim, C.J., and C.R. Nelson (1999), *State Space Models with Regime Switching: Classical and Gibbs Sampling Approaches with Applications*. Cambridge, MA: MIT Press.
- Pesaran, M.H., and A. Timmermann (2009), "Testing Dependence Among Serially Correlated Multicategory Variables," *Journal of the American Statistical Association*, 104, 325-337.
- Sanders, D.R., Manfredo, M.R., and K. Boris (2009), "Evaluating Information in Multiple Horizon Forecasts: The DOE's Energy Price Forecasts," *Energy Economics*, 31, 189-196.
- Strumpf, D. (2013), "Goldman Cuts the Near-Term Brent Crude Forecast to \$100 a Barrel," *Wall Street Journal*, April 23.

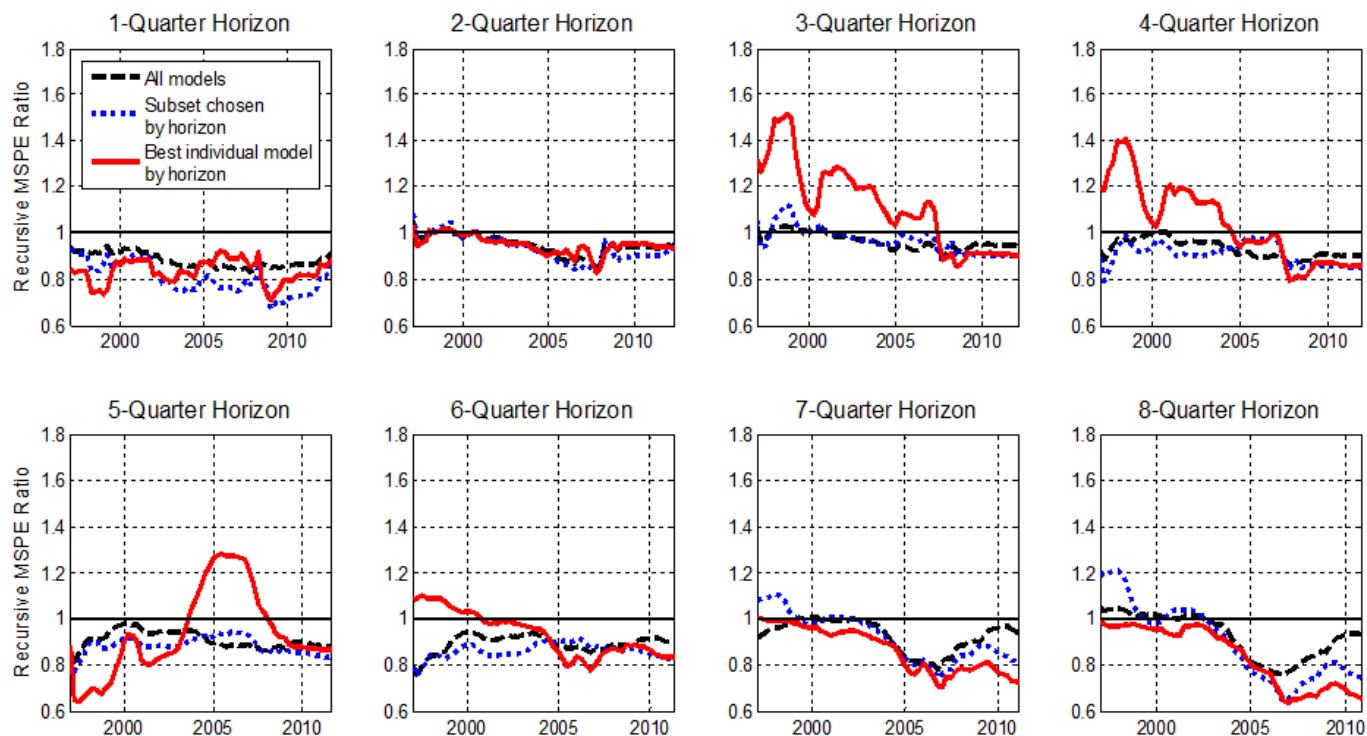
Timmermann, A. (2006), "Forecast Combinations," in: Elliott, G., Granger, C.W.J., and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 1, Amsterdam: North-Holland, 135-196.

Figure 1. Real-time recursive MSPE ratio relative to no-change forecast real U.S. refiner's acquisition cost for oil imports



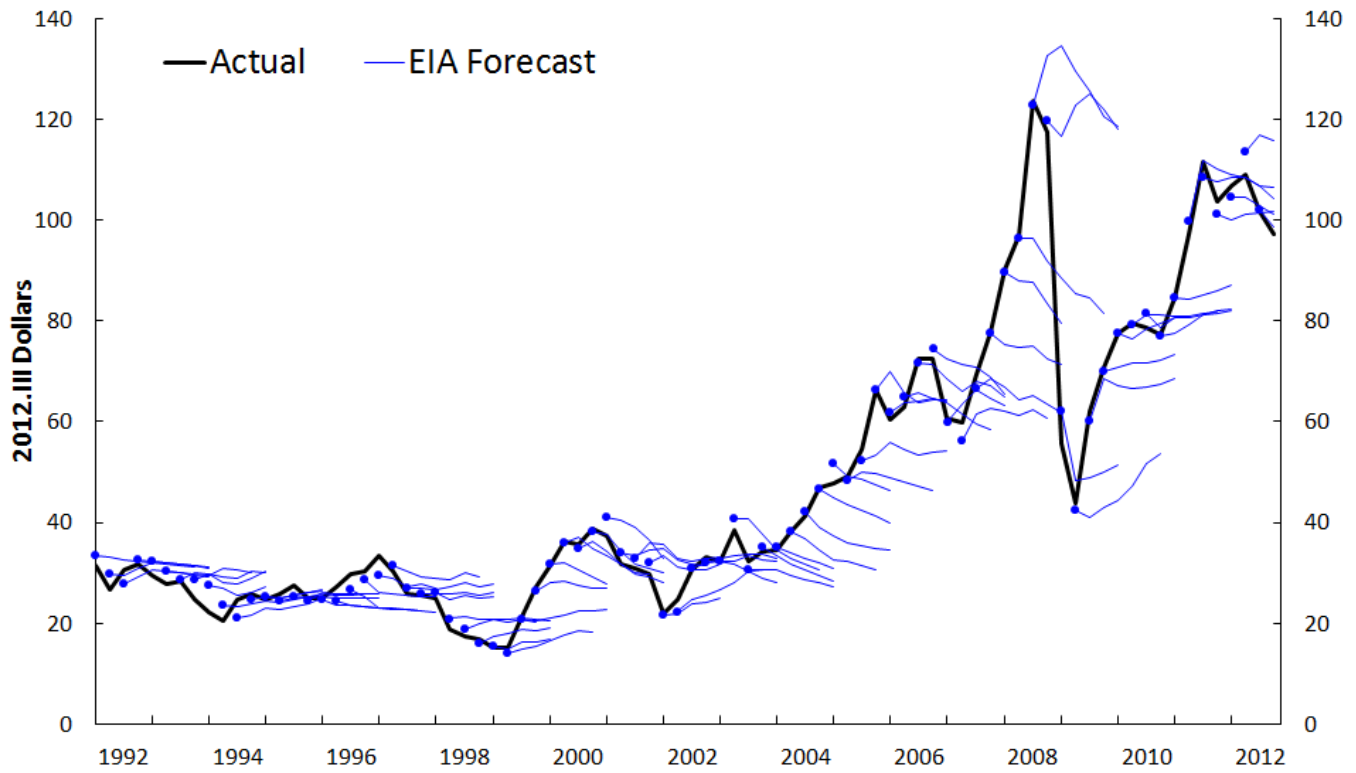
NOTES: All results are based on the methods examined in Table 1. A ratio below 1 indicates an improvement relative to the no-change forecast. The plot shows the evolution of the recursive MSPE ratio over time for the forecast evaluation period since 1997. This increases the reliability of the MSPE estimates and allows the MSPE ratio to stabilize.

Figure 2. Real-time recursive MSPE ratio relative to no-change forecast real U.S. refiner's acquisition cost for oil imports



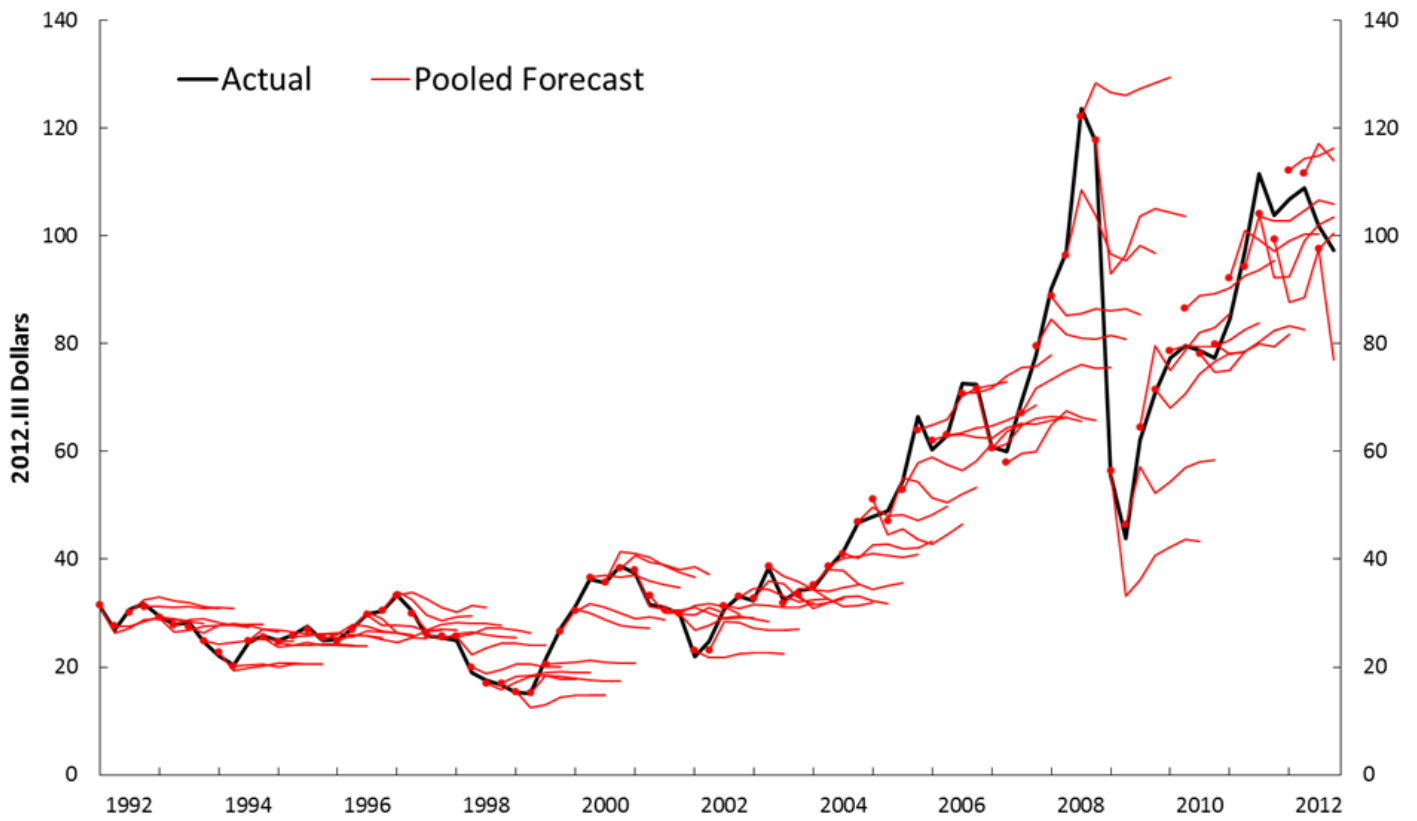
NOTES: All results based on the methods examined in Table 4. A ratio below 1 indicates an improvement relative to the no-change forecast. The plot shows the evolution of the recursive MSPE ratio over time for the forecast evaluation period since 1997. This increases the reliability of the MSPE estimates and allows the MSPE ratio to stabilize.

Figure 3. Real-time recursive forecasts and realizations for the real U.S. refiner's acquisition cost for oil imports quarterly EIA forecasts



NOTES: The nominal EIA forecasts from the *Short-Term Energy Outlook* have been adjusted for expected inflation. The EIA's nowcasts are marked by a circle.

Figure 4. Real-time recursive forecasts and realizations for the real U.S. refiner's acquisition cost for oil imports quarterly pooled real-time forecasts chosen by horizon



NOTES: The nowcasts constructed as in Baumeister and Kilian (2012) are marked by a circle.

Table 1. Real-time forecast accuracy of alternative forecast combinations evaluation period: 1992.1-2012.9

Monthly horizon	Real U.S. Refiners' Acquisition Cost for Crude Oil Imports					
	Equal-weighted combination of all 5 forecasting models		Equal-weighted subset of models chosen by horizon		Best individual forecasting model by horizon	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio (model)	Success ratio
1	0.912	0.578*	0.872	0.550*	0.934 (2)	0.546**
2	0.887	0.573*	0.824	0.577*	0.848 (2)	0.552**
3	0.894	0.559*	0.849	0.599*	0.864 (2)	0.628*
4	0.914	0.561	0.881	0.585*	0.938 (2)	0.598*
5	0.945	0.527	0.926	0.514	0.979 (3)	0.498
6	0.960	0.557	0.944	0.545	0.971 (4)	0.541
7	0.962	0.535	0.946	0.531	0.935 (4)	0.527
8	0.956	0.533	0.905	0.574	0.914 (4)	0.579
9	0.944	0.573	0.884	0.598**	0.892 (4)	0.560
10	0.927	0.584**	0.869	0.613*	0.879 (4)	0.588
11	0.906	0.603*	0.845	0.636*	0.856 (4)	0.623
12	0.894	0.626*	0.837	0.639*	0.865 (4)	0.613
13	0.890	0.637*	0.834	0.654*	0.869 (3)	0.629*
14	0.885	0.636*	0.846	0.661*	0.861 (3)	0.640*
15	0.885	0.677*	0.836	0.694*	0.860 (3)	0.634*
16	0.893	0.680*	0.832	0.727*	0.862 (5)	0.632*
17	0.907	0.682*	0.836	0.717*	0.849 (5)	0.644*
18	0.917	0.655*	0.835	0.720*	0.835 (5)	0.629**
19	0.961	0.593	0.843	0.680*	0.802 (5)	0.623
20	0.955	0.635*	0.812	0.696*	0.742 (5)	0.709*
21	0.939	0.625**	0.772	0.664	0.680 (5)	0.716*
22	0.919	0.645*	0.731	0.671*	0.642 (5)	0.746*
23	0.910	0.665*	0.710	0.648	0.654 (5)	0.740*
24	0.919	0.611	0.716	0.637	0.695 (5)	0.708*

NOTES: We consider five forecasting models based on monthly data: (1) a VAR forecast, (2) a forecast based on non-oil industrial commodity prices, (3) a forecast based on oil futures prices, (4) a forecast based on the spread of product prices relative to the price of crude oil, and (5) a forecast based on U.S. crude oil inventories. All forecasts are generated recursively from data subject to real-time data constraints. Boldface indicates improvements relative to the no-change forecast. * denotes significance at the 5% level and ** at the 10% level based on the Pesaran and Timmermann (2009) test for the null hypothesis of no directional accuracy. The statistical significance of the MSPE reductions cannot be assessed because none of the currently available tests of equal predictive accuracy applies in this setting.

Table 2. Changes in real-time recursive MSPE ratios of “leave-one-out” forecast combinations with equal weights evaluation period: 1992.1-2012.9

Monthly Horizon	Real U.S. Refiners' Acquisition Cost for Crude Oil Imports				
	VAR	Oil Futures Spread	Industrial Commodity Prices	TVP Product Price Spread	U.S. Crude Oil Inventories
1	+0.046	-0.015	+0.013	-0.008	-0.016
2	+0.039	-0.013	+0.032	-0.016	-0.023
3	+0.032	-0.009	+0.030	-0.011	-0.020
4	+0.033	-0.005	+0.023	-0.005	-0.020
5	+0.028	+0.003	+0.009	+0.005	-0.019
6	+0.015	+0.009	-0.001	+0.015	-0.011
7	+0.001	+0.012	-0.005	+0.023	-0.008
8	-0.005	+0.016	-0.007	+0.027	-0.006
9	-0.009	+0.018	-0.007	+0.030	-0.004
10	-0.010	+0.019	-0.008	+0.028	0
11	-0.011	+0.020	-0.006	+0.029	-0.002
12	-0.011	+0.024	-0.003	+0.023	-0.001
13	-0.011	+0.029	-0.002	+0.018	-0.001
14	-0.015	+0.031	-0.004	+0.019	+0.001
15	-0.020	+0.033	-0.007	+0.018	+0.008
16	-0.023	+0.035	-0.012	+0.016	+0.021
17	-0.027	+0.036	-0.016	+0.014	+0.029
18	-0.033	+0.036	-0.019	+0.016	+0.038
19	-0.028	NA	-0.032	+0.030	+0.082
20	-0.024	NA	-0.055	+0.030	+0.104
21	-0.022	NA	-0.078	+0.034	+0.124
22	-0.027	NA	-0.091	+0.045	+0.133
23	-0.031	NA	-0.097	+0.064	+0.123
24	-0.033	NA	-0.098	+0.081	+0.108

NOTES: The models are described in the text. Boldface indicates increases relative to the MSPE ratio in column (1) of Table 1. Increases mean that the model left out would have improved forecast accuracy if included, whereas decreases mean that it would have worsened forecast accuracy. The statistical significance of the MSPE changes cannot be assessed because none of the currently available tests of equal predictive accuracy applies in this setting.

Table 3. Real-time forecast accuracy of alternative forecast combinations evaluation period: 1992.1-2012.9

Real WTI Price of Crude Oil						
Monthly horizon	Equal-weighted combination of all 5 forecasting models		Equal-weighted subset of models chosen by horizon		Best individual forecasting model by horizon	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio (model)	Success ratio
1	0.898	0.518	0.846	0.494	0.866 (2)	0.530
2	0.884	0.565*	0.814	0.565*	0.813 (2)	0.565*
3	0.894	0.587*	0.841	0.591*	0.862 (2)	0.595*
4	0.919	0.578**	0.877	0.589*	0.959 (2)	0.602*
5	0.949	0.543	0.928	0.535	0.986 (4)	0.584
6	0.968	0.533	0.939	0.541	0.973 (4)	0.553
7	0.969	0.547	0.940	0.539	0.943 (4)	0.547
8	0.959	0.537	0.927	0.550	0.907 (4)	0.558
9	0.949	0.552	0.919	0.564	0.896 (4)	0.564
10	0.930	0.538	0.903	0.554	0.876 (4)	0.571
11	0.911	0.603*	0.884	0.582*	0.859 (4)	0.611
12	0.905	0.605*	0.883	0.576**	0.892 (4)	0.597
13	0.899	0.616*	0.876	0.599*	0.893 (3)	0.603*
14	0.894	0.610*	0.868	0.614*	0.885 (3)	0.606*
15	0.898	0.634*	0.873	0.647*	0.884 (3)	0.626*
16	0.906	0.667*	0.851	0.697*	0.884 (5)	0.662*
17	0.916	0.674*	0.856	0.678*	0.880 (5)	0.635*
18	0.926	0.638*	0.859	0.668*	0.873 (5)	0.651*
19	0.963	0.593**	0.916	0.641*	0.845 (5)	0.745*
20	0.965	0.639*	0.899	0.683*	0.795 (5)	0.739*
21	0.953	0.625**	0.867	0.703*	0.743 (5)	0.773*
22	0.924	0.645*	0.824	0.680*	0.703 (5)	0.816*
23	0.916	0.648*	0.812	0.678*	0.704 (5)	0.819*
24	0.916	0.628**	0.809	0.646**	0.731 (5)	0.792*

NOTES: We consider five forecasting models based on monthly data: (1) a VAR forecast, (2) a forecast based on non-oil industrial commodity prices, (3) a forecast based on oil futures prices, (4) a forecast based on the spread of product prices relative to the price of crude oil, and (5) a forecast based on U.S. crude oil inventories. All forecasts are generated recursively from data subject to real-time data constraints. Boldface indicates improvements relative to the no-change forecast. * denotes significance at the 5% level and ** at the 10% level based on the Pesaran and Timmermann (2009) test for the null hypothesis of no directional accuracy. The statistical significance of the MSPE reductions cannot be assessed because none of the currently available tests of equal predictive accuracy applies in this setting.

Table 4. Real-time forecast accuracy of alternative forecast combinations evaluation period: 1992.1-2012.9

Quarterly horizon	Equal-weighted combination of all 5 forecasting models		Equal-weighted subset of models chosen by horizon		Best individual forecasting model by horizon	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
Real U.S. Refiners' Acquisition Cost for Crude Oil Imports						
1	0.921	0.687*	0.894	0.711*	0.903	0.735*
2	0.945	0.646*	0.920	0.610*	0.945	0.646*
3	0.938	0.654*	0.901	0.593**	0.907	0.543
4	0.898	0.700*	0.848	0.663*	0.858	0.625*
5	0.880	0.620*	0.837	0.633*	0.866	0.620*
6	0.897	0.628*	0.831	0.641*	0.835	0.603
7	0.947	0.558	0.814	0.649*	0.728	0.623*
8	0.927	0.540	0.741	0.618	0.649	0.592
Real WTI Price of Crude Oil						
1	0.901	0.699*	0.845	0.687*	0.856	0.711*
2	0.949	0.646*	0.914	0.622*	0.971	0.659*
3	0.945	0.654*	0.910	0.593*	0.909	0.593**
4	0.902	0.663*	0.873	0.650*	0.856	0.638**
5	0.893	0.608*	0.869	0.570	0.890	0.608*
6	0.908	0.628*	0.847	0.628*	0.864	0.603**
7	0.961	0.584	0.895	0.546	0.780	0.623*
8	0.933	0.566	0.833	0.618	0.706	0.632*

NOTES: We consider five forecasting models based on monthly data: (1) a VAR forecast, (2) a forecast based on non-oil industrial commodity prices, (3) a forecast based on oil futures prices, (4) a forecast based on the spread of product prices relative to the price of crude oil, and (5) a forecast based on U.S. crude oil inventories. All forecasts are generated recursively from data subject to real-time data constraints. Boldface indicates improvements relative to the no-change forecast. * denotes significance at the 5% level and ** at the 10% level based on the Pesaran and Timmermann (2009) test for the null hypothesis of no directional accuracy. The statistical significance of the MSPE reductions cannot be assessed because none of the currently available tests of equal predictive accuracy applies in this setting.

Table 5: Selected real-time forecast paths for the U.S. refiners' acquisition cost for crude oil imports

Example 1: 2000.IV		1999		2000				2001				2002
		3rd	4th	1st	2nd	3rd	4th	1st	2nd	3rd	4th	1st
Real	<i>Pooled</i>	20.58	23.84	27.44	26.91	29.33	28.82	23.69	23.21	22.61	22.47	23.46
	<i>EIA</i>	20.54	23.85	27.41	26.94	29.33	31.06	30.69	29.56	27.61	24.94	N.A.
	<i>Actual</i>	20.57	23.83	27.44	26.94	29.32	28.30	23.93	23.46	22.57	16.61	18.77
Nominal	<i>Pooled</i>	19.75	23.04	26.79	26.52	29.12	28.82	23.90	23.61	23.18	23.23	24.45
	<i>EIA</i>	19.70	23.01	26.84	26.55	29.11	31.06	30.94	30.04	28.29	25.76	N.A.
	<i>Actual</i>	19.75	23.04	26.79	26.52	29.12	28.30	24.16	23.86	23.01	16.92	19.19
Example 2: 2008.III				2007				2008				2009
		2nd	3rd	4th	1st	2nd	3rd	4th	1st	2nd	3rd	4th
Real	<i>Pooled</i>	66.24	74.26	85.84	92.46	117.92	112.24	88.63	91.96	98.87	100.14	99.50
	<i>EIA</i>	66.05	74.16	85.57	92.43	116.93	114.01	111.25	117.16	119.16	116.31	112.52
	<i>Actual</i>	66.10	74.16	85.73	92.30	117.58	111.78	52.90	41.67	59.01	67.48	73.67
Nominal	<i>Pooled</i>	62.41	70.45	82.45	89.75	115.90	112.24	89.26	93.37	101.16	103.28	103.42
	<i>EIA</i>	62.30	70.38	82.44	89.73	115.70	114.01	112.13	119.02	122.00	120.02	117.01
	<i>Actual</i>	62.41	70.45	82.45	89.73	115.83	111.78	51.84	40.45	57.57	66.40	73.04
Example 3: 2009.II				2008				2009				2010
		1st	2nd	3rd	4th	1st	2nd	3rd	4th	1st	2nd	3rd
Real	<i>Pooled</i>	89.79	114.66	108.98	51.51	40.55	59.76	73.79	69.60	72.88	76.03	76.92
	<i>EIA</i>	89.74	114.13	110.26	52.82	40.64	55.77	63.57	62.09	61.64	61.85	62.36
	<i>Actual</i>	90.06	114.73	109.08	51.63	40.67	57.57	65.85	71.89	73.87	73.10	71.83
Nominal	<i>Pooled</i>	89.75	115.90	111.79	51.84	40.45	59.76	74.34	70.63	74.50	78.29	79.79
	<i>EIA</i>	89.74	115.93	112.85	52.31	40.46	55.77	64.03	63.00	63.00	63.67	64.66
	<i>Actual</i>	89.73	115.83	111.78	51.84	40.45	57.57	66.40	73.04	75.19	74.36	73.31

NOTES: Nowcasts are shown in bold; forecasts in italics. By construction the nominal and the real price coincide in the last nowcast period.