# The Art and Science of Forecasting the Real Price of Oil

*Christiane Baumeister, International Economic Analysis*


- Forecasts of the price of crude oil play a significant role in the conduct of monetary policy, especially for commodity-producing countries such as Canada.
- This article explores a range of recently developed forecasting models that can generate, on average, accurate forecasts of the price of oil. Instead of relying on a single model, this article shows that forecast combinations outperform the oil futures curve.
- In addition to accurate forecasts of the price of oil, policy-makers are interested in evaluating the risks associated with the baseline forecast to gauge the implications of alternative oil price paths for the economic outlook. A structural model of the global oil market can be used to develop risk scenarios for oil price forecasts, based on hypothetical assumptions about future demand and supply conditions in the crude oil market.
- Based on this structural model, it can also be shown that changes in demand associated with the global business cycle are the primary determinant of changes in oil prices.


Given the importance of oil prices for the Canadian economy, understanding what drives fluctuations in oil prices and how best to forecast them is critical for monetary policy. Specifically, oil price forecasts play an important role in assessing the future developments of inflation and economic activity in Canada and its trading partners, with implications for Canada's terms of trade.

Until recently, central banks and international organizations tended to rely exclusively on the oil futures curve to forecast the price of oil. Recent research, however, demonstrates that models that include the economic determinants of the price of oil, such as changes in oil inventories, oil production and global real economic activity, may provide more accurate out-of-sample forecasts than oil futures prices (Alquist, Kilian and Vigfusson 2013; Baumeister and Kilian 2014b; Baumeister, Kilian and Zhou 2013). This finding holds even in a real-time forecasting environment, where predictors of the price of oil become available only with a delay and are subsequently revised repeatedly (Baumeister and Kilian 2012).

An important limitation of all these forecasting approaches is that they provide limited insight into what is affecting the oil price forecast. It is, however, possible to derive a structural model of the global oil market from one of the forecasting models that helps policy-makers to interpret oil price forecasts. Such an economic model can also be used to evaluate the risks associated with the baseline forecast—that is, how the future path of the price of oil changes with alternative hypothetical scenarios for the economic environment.

This article begins by describing recent advances in forecasting the real price of oil. It stresses the benefits of combining the forecasts from alternative models that have different strengths and weaknesses, rather than relying on only one forecasting method. It then outlines a framework for constructing forecast scenarios that enhances policy-makers' understanding of the factors affecting oil prices and allows them to formally assess the risks associated with oil price forecasts.

## Forecasting Models

The volatility of the real price of oil since 2003 has renewed interest in how best to forecast oil prices (**Chart 1**). This section presents the traditional approach that uses oil futures prices as predictors of the real price of oil, as well as three forecasting methods developed recently at the Bank of Canada. It then compares the relative accuracy of a combination of these forecasting methods with the no-change benchmark forecast.

**Chart 1: The real price of oil, 1974 to 2013, in November 2013 U.S. dollars per barrel**

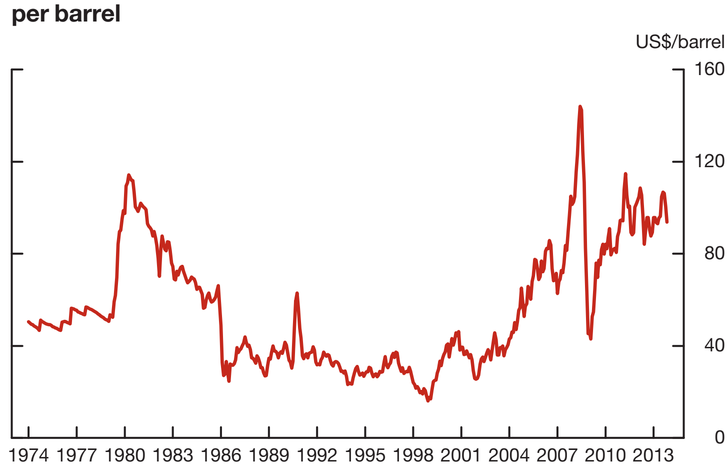


Note: The real price of oil is the nominal refiner acquisition cost of crude oil imports deflated by U.S. CPI.
Sources: U.S. Energy Information Administration, *Monthly Energy Review*; and the Federal Reserve Bank of St. Louis
Last observation: November 2013

### Oil futures curve

The traditional approach to constructing out-of-sample forecasts of the real price of oil is to rely on the oil futures curve. Since the oil futures market plays an important role in information aggregation and price discovery, the prices of crude oil futures contracts traded on exchanges such as the New York Mercantile Exchange or the Intercontinental Exchange are commonly perceived to reflect the expectations of market participants about the future course of oil prices (Alquist and Arbatli 2010). When communicating policy decisions, many central banks have highlighted the importance of oil futures prices for the future evolution of inflation.

When the forecasting performance of oil futures prices is evaluated over a period of 20 years against a simple model that postulates that prices will remain unchanged over the forecast horizon (the no-change forecast), at shorter horizons, there is no significant evidence that the oil futures curve achieves gains in forecast accuracy. Moreover, Alquist, Kilian and Vigfusson (2013) show that, at longer-term horizons that matter for policy-makers, the forecasting performance of the oil futures curve is inferior when compared with the no-change forecast. A possible explanation for this finding is that oil futures prices contain a time-varying risk premium. In fact, Hamilton and Wu (2014) find evidence of considerable changes in risk premiums in oil futures prices after 2005.

*When the forecasting performance of oil futures prices is evaluated against a simple model postulating that prices will remain unchanged over the forecast horizon, at shorter horizons, there is no significant evidence that the oil futures curve achieves gains in forecast accuracy*

## Model of the global oil market

The first of the recently developed alternative approaches uses a model of the global market for crude oil that includes the key determinants of oil prices based on economic theory. Specifically, the current real price of oil is modelled as a function of its own past values and the past values of world oil production, an index of real economic activity that captures fluctuations in the global business cycle and changes in global above-ground inventories of crude oil.

Out-of-sample forecasts generated by this model tend to be more accurate than the no-change forecast at short horizons, even when real-time data constraints are taken into account (Baumeister and Kilian 2012; 2014b).

## Spot price of raw industrial materials

The second alternative method is based on the observation that prices of non-oil industrial commodities such as copper and zinc are indicators of shifts in the demand for all industrial commodities, including oil. To the extent that persistent fluctuations in the global business cycle move together with industrial commodity prices, recent cumulative changes in the price indexes of non-oil industrial commodities are expected to have predictive power for the real price of oil.

Based on this insight, Baumeister and Kilian (2012) show that forecasts that extrapolate cumulative changes in the spot price for raw industrial materials adjusted for expected inflation perform well at short horizons relative to the no-change forecast, but become increasingly less accurate at horizons beyond three months. The ability of these forecasts to accurately predict whether the price of oil is increasing or decreasing is consistently high for horizons of up to 12 months.

## Refined product spreads

The third promising forecasting approach is based on the idea that the demand for crude oil is driven by the demand for refined petroleum products, such as gasoline, heating oil and diesel. This relationship suggests that spot market prices for petroleum products will ultimately determine the price for crude oil. In fact, many oil industry analysts believe that a widening of the spread between product prices and the price of crude oil signals upward pressures on future oil prices. This insight may be exploited by analyzing whether changes in these price spreads, defined as the extent to which today's price of gasoline or heating oil deviates from today's price of crude oil, have predictive power for future changes in the price of oil.

There are, however, many reasons to expect this predictive relationship to be unstable over time. Given that refined products are produced in approximately fixed proportions, the price of crude oil is likely to be determined by the refined product in highest demand, and this product changes gradually. Another concern is that oil supply shocks, local capacity constraints in refining, changes in environmental regulations or other market turmoil may temporarily undermine the predictive content of these price spreads. To take these possibilities into account, Baumeister, Kilian and Zhou (2013) allow the weights assigned to gasoline price spreads and heating oil price spreads in the forecasting model to evolve smoothly. They find that this model delivers more accurate forecasts relative to the no-change forecast, especially at horizons between one and two years.

## Combining forecasts from different models

Given the range of models available for forecasting the real price of oil, the question that arises is which model policy-makers should rely on to obtain the most accurate point forecasts and to correctly predict whether the oil price will go up or down over the projection horizon.

Rather than selecting a single model, it might be beneficial to pool the information contained in the four individual models (Baumeister and Kilian 2013). Combining forecasts from different models is promising for several reasons. First, even the most accurate forecasting models do not work equally well at all times. The global oil market forecasting model, for example, works well during periods when economic fundamentals show persistent variation, as was the case between 2002 and 2011, but not as well at other times. Similarly, there is considerable variation over time in the ability of oil futures prices to forecast the price of oil.

Second, the forecasting performance of individual models differs depending on the forecast horizon. For example, models based on economic fundamentals tend to be more accurate at short horizons, while models based on the spread between the prices of refined products and the price of crude oil tend to be more accurate at longer horizons. Since the policy horizon usually spans a period of two years, no single model provides the most accurate forecasts for the entire projection horizon.

Third, even a forecasting model with a better forecasting track record could be improved upon by incorporating additional information from other models that perform less well, on average.

*Pooling information from different models produces more robust forecasts than relying on any one individual model*

These arguments suggest that combining forecasts from different models should be more reliable than individual models.[1] To evaluate the forecasting performance of equal-weighted forecast combinations, two criteria are considered. The first is the relative mean-squared prediction error (MSPE), which measures the average squared deviation between the pooled forecasts and the actual realization, relative to the no-change benchmark. An MSPE ratio below 1 indicates that the pooled forecasts are more accurate than the benchmark forecast. Second, the directional accuracy of the pooled forecasts is assessed by the success ratio, which represents the number of times that a method correctly predicts whether the real price of oil is increasing or decreasing. If there were no directional accuracy, the model should be no more successful at predicting the direction of price changes

1 Baumeister and Kilian (2013) conclude that the best way to pool individual forecasts of oil prices is by assigning equal weight to them. This approach is more accurate than basing weights on the recent forecasting performance of each model.

**Table 1: Real-time forecast accuracy of pooled forecasts with equal weights**

| Horizon (in months) | Real U.S. refiner acquisition cost (RAC) for crude oil imports | Real West Texas Intermediate (WTI) price |
|---|---|---|
| | Relative mean-squared prediction error (MSPE) ratios | |
| 1 | **0.897** | **0.880** |
| 3 | **0.874** | **0.873** |
| 6 | **0.949** | **0.956** |
| 9 | **0.939** | **0.943** |
| 12 | **0.892** | **0.902** |
| 15 | **0.893** | **0.906** |
| 18 | **0.957** | **0.959** |
| 21 | 1.065 | 1.064 |
| 24 | 1.029 | 1.017 |
| | Success ratios | |
| 1 | **0.554*** | **0.517** |
| 3 | **0.609*** | **0.592*** |
| 6 | **0.556** | **0.543** |
| 9 | **0.580**** | **0.562** |
| 12 | **0.609*** | **0.605*** |
| 15 | **0.650*** | **0.617*** |
| 18 | **0.601*** | **0.577**** |
| 21 | **0.550** | **0.550** |
| 24 | **0.561** | **0.551** |

Notes: Numbers in **bold** indicate improvements relative to the no-change forecast.
* denotes significance at the 5 per cent level and ** at the 10 per cent 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.
Source: Baumeister and Kilian (2013)

than a coin toss, with a success probability of 0.5 (or 50 per cent). Thus, success ratios above 0.5 indicate an improvement relative to the no-change forecast.

The forecasts are generated for two alternative measures of the real price of crude oil. The first is the U.S. refiner acquisition cost (RAC) of imported crude oil, which is considered a good proxy for a truly global oil price. The second is the West Texas Intermediate (WTI) spot price, which is the U.S. benchmark. Equal-weighted forecast combinations for the real RAC and WTI yield considerable reductions in MSPE ratios for horizons up to 18 months, ranging from 4 per cent to 13 per cent (Table 1).[2] These forecast combinations are also successful at predicting the direction of change for these horizons. For the RAC, improvements in directional accuracy are statistically significant at all but one horizon up to 18 months, and range from 55 per cent to 65 per cent. For the WTI, the highest success ratio is 62 per cent, but the directional accuracy is statistically significant at only four of these horizons.

2 Chart A-1 in the appendix shows the evolution of the recursive root-mean-squared prediction errors for the equal-weighted combination forecast in comparison with the no-change forecast for selected forecast horizons.

## A Structural Model for Assessing Risks to Oil Prices

An important limitation of these forecasting approaches is that they do not help policy-makers explore how the forecast would change relative to the baseline forecast under hypothetical assumptions about future economic conditions. For example, it is important to know how much the real price of oil would be affected by civil unrest in the Middle East or by a period of unexpectedly low global demand for crude oil caused by a worldwide recession. Similarly, policy-makers want to understand what drove changes in the real price of oil in the recent past, such as the persistent increase in oil prices between 2003 and 2008, or the rapid decline between 2008 and 2012. To address questions such as these, a structural model of the global oil market is required.

### A structural model of the global oil market

Kilian and Murphy (2014) propose a dynamic structural model that includes the key determinants of the real price of oil: changes in global oil production, real economic activity worldwide and above-ground crude oil inventories. This econometric model is motivated by a standard stock-flow model with an explicit role for expectations and can be directly derived from the forecasting model of the global oil market by imposing additional economic structure. Within this framework, it is possible to decompose past fluctuations in oil prices into structural driving forces stemming from supply and demand. In particular, the authors distinguish between four types of shocks:

*Using a dynamic structural model of the global oil market, it is possible to decompose past fluctuations in oil prices into structural driving forces stemming from supply and demand*

(i) **Flow supply shock**—a classic oil supply shock that captures disruptions to the flow of oil production resulting, for example, from exogenous political events in oil-producing countries, such as war or civil unrest.

(ii) **Flow demand shock**—a shock to the demand for crude oil that is associated with unexpected fluctuations in the global business cycle. A prominent example is the surprisingly persistent demand from emerging-market economies, particularly China.

(iii) **Speculative demand shock**—a shock to the demand for oil inventories arising from shifts in expectations about future demand for and supply of oil that is not otherwise captured in the model. Such shifts could arise, for example, from the anticipation of several factors, including political unrest in oil-exporting countries in the Middle East, peak oil effects or the depletion of oil reserves. This shock thus captures forward-looking behaviour and speculation.

(iv) **Other demand shock**—a residual shock that has no economic interpretation but is designed to capture idiosyncratic oil demand shocks not otherwise accounted for. Examples of such shocks include weather shocks, changes in inventory technology or preferences, or politically motivated releases of the U.S. Strategic Petroleum Reserve.

This model provides a coherent framework both to understand past oil price fluctuations and to assess risks associated with oil price forecasts.

### The contribution of each shock to cumulative price changes

Kilian and Lee (2014) use this model to quantify the contribution of each type of shock to cumulative changes in the real U.S.-dollar price of oil during specific historical episodes.

**Chart 2: Contribution of each structural shock to the cumulative change in the real price of oil**

a. January 2003 to June 2008, in May 2012 U.S. dollars per barrel

b. June 2008 to May 2012, in May 2012 U.S. dollars per barrel

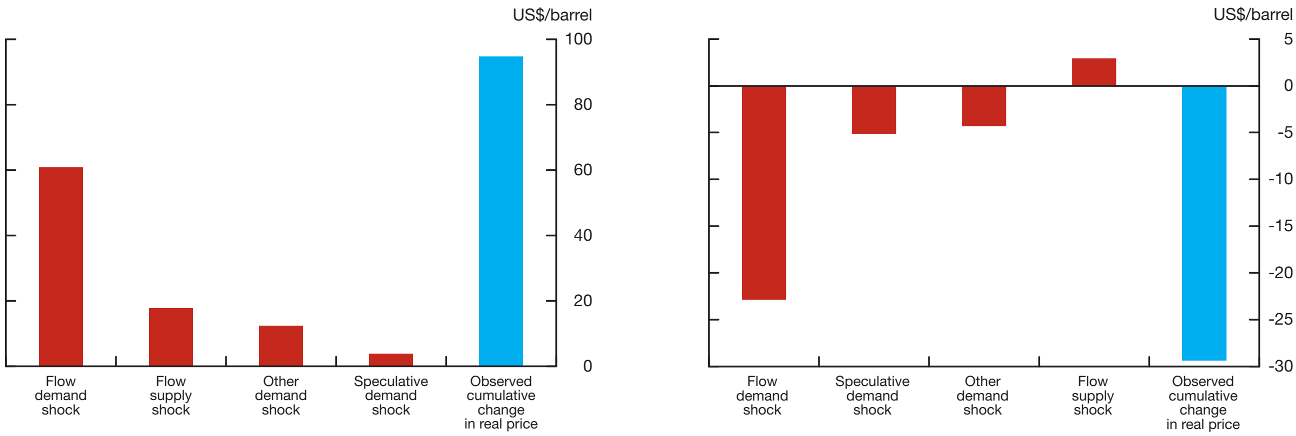


Source: Kilian and Lee (2014)

From 2003 to mid-2008, oil prices experienced an unprecedented surge. This development triggered a debate about whether the run-up in oil prices could be explained by elevated demand from China, or whether it was the result of the financialization of physical oil markets.[3]

**Chart 2a** provides compelling evidence that an unexpectedly strong world economy was the main cause for the rise in global oil prices. In fact, flow demand shocks associated with shifts in the global demand for oil from emerging Asia and from member countries of the Organisation for Economic Co-operation and Development accounted for US$60 of the observed US$95 increase in the real per-barrel price of oil during that period. While supply-side factors contributed somewhat to the upswing in oil prices, they accounted for less than US$20 of the increase. Speculation by oil consumers, in contrast, was negligible.

A similar picture emerges for the decline in the price of oil between mid-2008 and 2012. **Chart 2b** shows that most of the US$29 decrease in the real per-barrel price of oil can be attributed to a series of unexpected negative flow demand shocks associated with the weak global economy in the wake of the financial crisis. As before, other factors played a limited role. These findings suggest that changes in the demand for crude oil associated with the global business cycle are the primary determinant of changes in oil prices.

## From explaining the past to assessing future risks

Over the projection horizon, unpredictable variations in the demand for and supply of crude oil can lead to deviations of the future oil price from its forecasted path. It is therefore useful to assess the sensitivity of the baseline forecast to potential events involving future demand and supply conditions in the crude oil market.

3 Financialization refers to the large increase in investors' participation in commodities as an asset class, as reflected, for instance, in the inflow of investment funds to oil futures markets in the past decade. This trend has led to a debate about the possible influence of financialization on oil price dynamics.

**Chart 3: Scenarios for the projected path of the real price of oil, in December 2010 U.S. dollars per barrel**

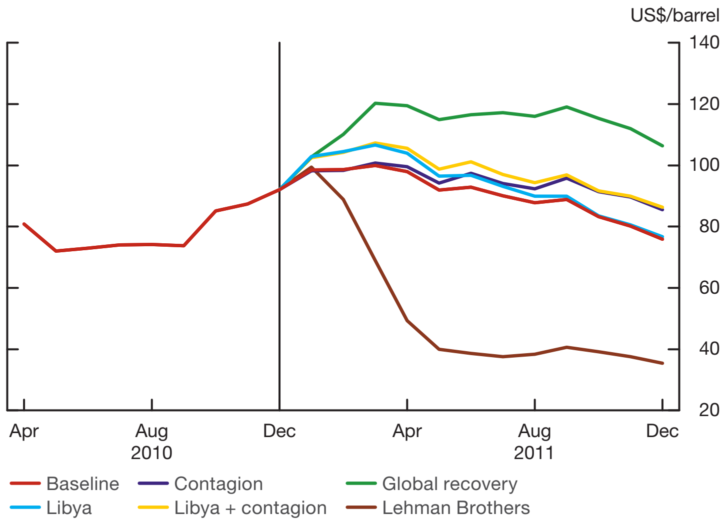


Note: The red line is the real-time out-of-sample forecast for the real U.S. refiner acquisition cost of crude oil in December 2010 U.S. dollars. The vertical line indicates the point in time when the forecast is made.

Source: Baumeister and Kilian (2014a)

To model such departures from the baseline forecast, Baumeister and Kilian (2014a) present alternative forecast scenarios based on the structural model of the oil market (Kilian and Murphy 2014). These scenarios examine the percentage deviation from the baseline forecast if a certain sequence of oil demand or supply shocks were to occur over the projection horizon. They are intended to help policy-makers gauge the possible consequences of unlikely events.

*Alternative forecast scenarios based on hypothetical sequences of oil demand and supply shocks help policy-makers gauge the possible consequences of unlikely events*

The baseline oil price forecast is generated as of December 2010 (**Chart 3**). To this baseline forecast we add five scenarios taken from Baumeister and Kilian (2014a). The first scenario relates to the supply side of the oil market and is motivated by the political unrest in Libya in early 2011. The authors ask what would have happened to the real price of oil if Libyan production, which accounts for 2.2 per cent of global oil production, were unexpectedly taken off the market. The results from the model show that such a shortfall in Libyan production would raise the price of oil by only 7 per cent after three months. This example illustrates that the observed increase in the price of oil of 21 per cent over that same period (**Chart 1**) cannot be attributed to supply disruptions alone.

Events such as the Arab Spring or the ongoing civil war in Syria can affect the oil price by triggering speculative demand, driven by fears of contagion of political unrest in the Middle East. Such an expectations-driven contagion scenario would increase the real price of oil by 20 per cent after about a year and a half, if the shift in speculative demand were comparable with the sustained speculative frenzy that began in mid-1979 following the Iranian revolution. The third scenario is a combination of the previous two scenarios.

The fourth and fifth scenarios relate to the role of the global business cycle. The global recovery scenario illustrates that an unexpected full recovery of the world economy would raise the real price of oil by an additional 40 per cent after about one year. The prospect of a global collapse shows that the recurrence of an event such as the financial crisis following the bankruptcy of Lehman Brothers in 2008 would be expected to lower the real price of oil by close to 60 per cent, as global demand drops dramatically.

For expository purposes, it is assumed that all of the scenarios begin in January 2011. Each scenario results in a different projected path for the real per-barrel price of oil, providing the full range of alternative outcomes. The real price of oil may fall as low as US$69 or rise as high as US$120 after one quarter, depending on the scenario. After one year, the range is between US$35 and US$106. Consistent with earlier results, the more extreme movements correspond to scenarios with large shifts in flow demand.

Obviously, policy-makers will not consider all scenarios equally likely; some scenarios will be mutually exclusive, while others might occur in conjunction. Assessing by how much such an alternative path deviates from the baseline, and how sensitive this deviation is to alternative assumptions about the relative likelihood of the underlying scenarios, allows policy-makers to get a better sense of the nature of the upside and downside risks involved. This information can also be used as input into more comprehensive risk scenarios that policy-makers might use to assess potential macroeconomic outcomes.

## Conclusion

Combinations of forecasts generated by different models are a useful tool for obtaining more accurate and robust out-of-sample forecasts of the real price of oil. These baseline forecasts can be supplemented by forecast scenarios from a structural model of the global oil market to evaluate upside and downside risks at various horizons. Such an approach is important because central bankers care not only about forecast accuracy but also about the economic interpretation underlying the past, present and future evolution of the real price of oil.

Appendix 1

## Evolution of Root-Mean-Squared Prediction Errors

**Chart A-1: Recursive root-mean-squared prediction errors for the combination forecast with equal weights and for the no-change forecast for horizons of 1, 12 and 24 months**

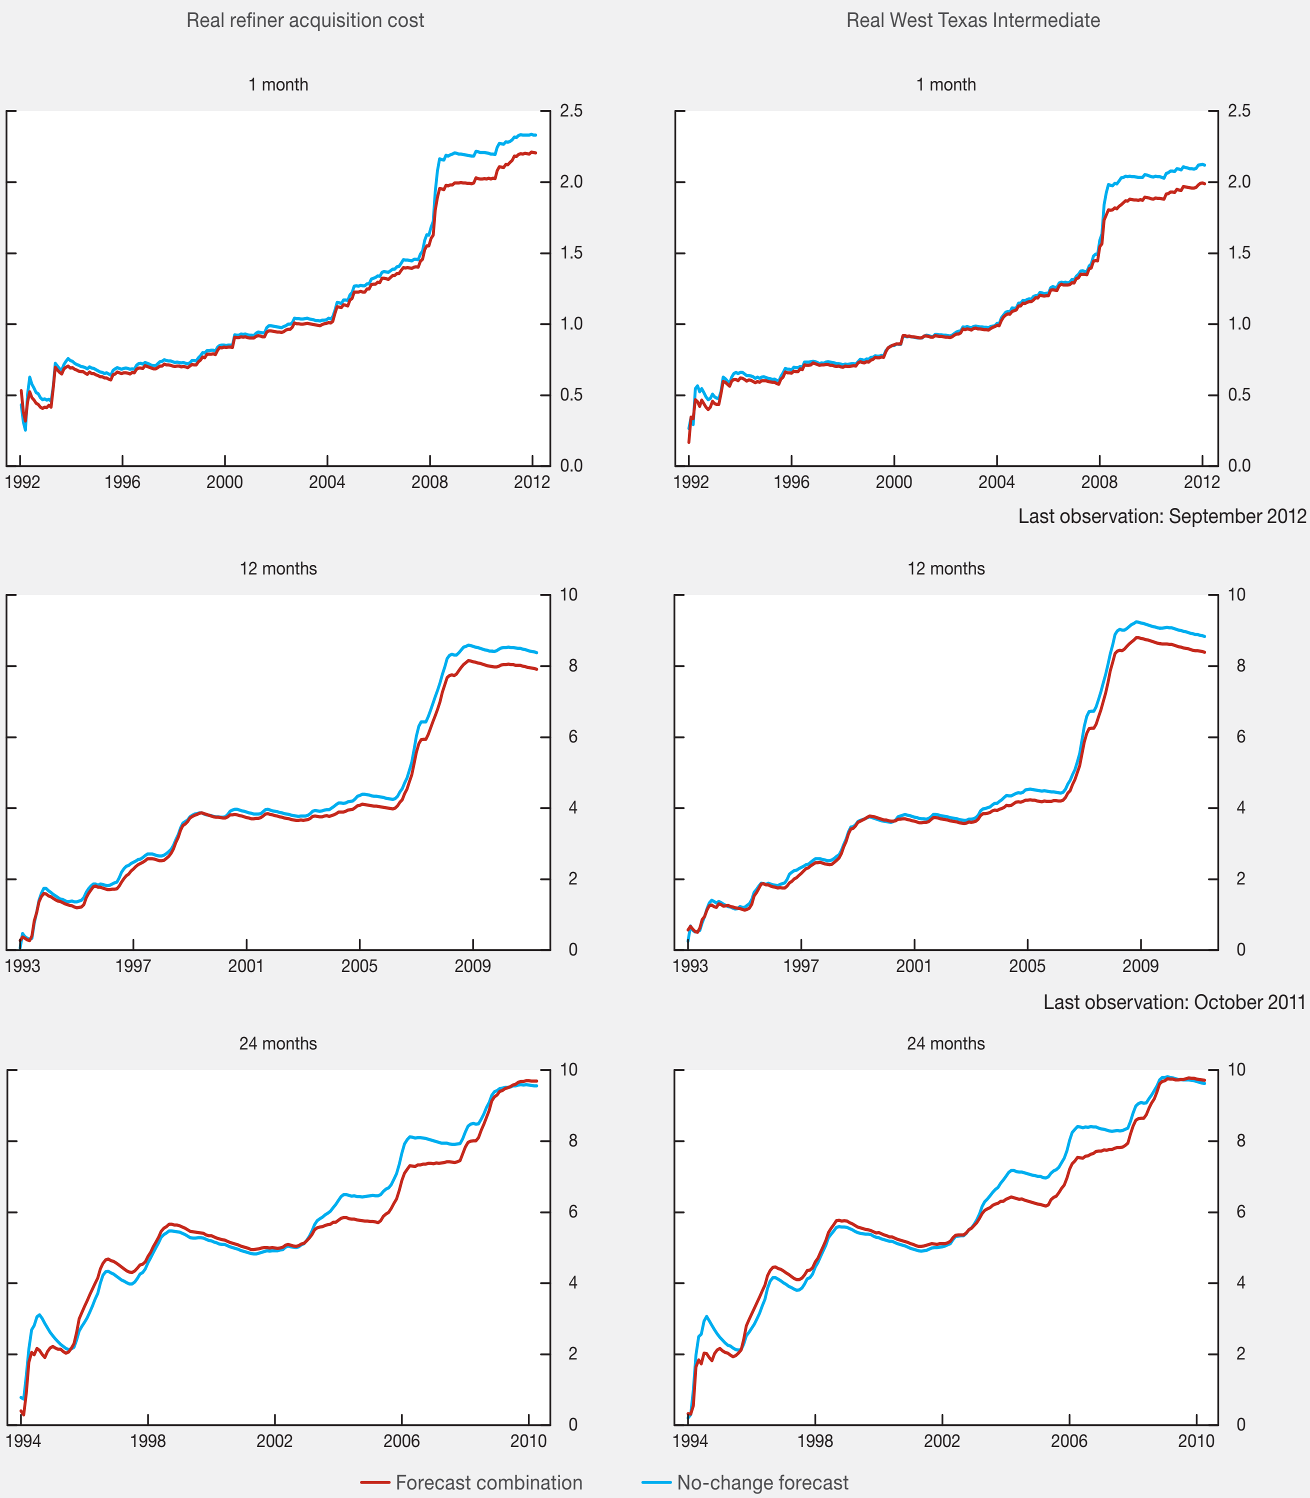


Source: Bank of Canada calculations

Last observation: October 2010

## Literature Cited

Alquist, R. and E. Arbatli. 2010. "Crude Oil Futures: A Crystal Ball?" *Bank of Canada Review* (Spring): 3–11.

Alquist, R., L. Kilian and R. J. Vigfusson. 2013. "Forecasting the Price of Oil." In *Handbook of Economic Forecasting*, 2A, edited by G. Elliott and A. Timmermann, 427–507. Amsterdam: North-Holland.

Baumeister, C. and L. Kilian. 2012. "Real-Time Forecasts of the Real Price of Oil." *Journal of Business and Economic Statistics* 30 (2): 326–36.

———. 2013. "Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach." Bank of Canada Working Paper No. 2013-28.

———. 2014a. "Real-Time Analysis of Oil Price Risks Using Forecast Scenarios." *IMF Economic Review* (forthcoming).

———. 2014b. "What Central Bankers Need to Know About Forecasting Oil Prices." *International Economic Review* 55 (3): 869–89.

Baumeister, C., L. Kilian and X. Zhou. 2013. "Are Product Spreads Useful for Forecasting? An Empirical Evaluation of the Verleger Hypothesis." Bank of Canada Working Paper No. 2013-25.

Hamilton, J. D. and J. C. Wu. 2014. "Risk Premia in Crude Oil Futures Prices." *Journal of International Money and Finance* 42: 9–37.

Kilian, L. and T. K. Lee. 2014. "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories." *Journal of International Money and Finance* 42: 71–87.

Kilian, L. and D. P. Murphy. 2014. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics* 29 (3): 454–78.

Pesaran, M. H. and A. Timmermann. 2009. "Testing Dependence Among Serially Correlated Multicategory Variables." *Journal of the American Statistical Association* 104 (485): 325–37.