The New Benchmark for Forecasts of the Real Price of Crude Oil

Amor Aniss Benmoussa
Economist, Bank of Canada

Reinhard Ellwanger, Ph.D.
Senior Economist, Bank of Canada

Stephen Snudden, Ph.D.
Assistant Professor, Wilfrid Laurier University, Canada

Full paper available at: https://www.bankofcanada.ca/2020/09/staff-working-paper-2020-39/

How can we assess the quality of a forecast? We propose a new benchmark to evaluate forecasts of averaged series and show that the real price of oil is more difficult to predict than we previously thought.

Is the Real Price of Crude Oil Predictable?

The payoff to investments in new oil production, oil-intensive goods purchases, and oil-related research all hinge critically on the quality of oil-price forecasts. Forecasts can be derived from a variety of approaches, including expert knowledge, economic or statistical models, or the prices of financial assets. But how should a forecaster assess the quality of a specific forecast?

A common way to address this question is to compare the accuracy of a forecast to that of a benchmark forecast. For forecasts of the real price of crude oil, this benchmark has typically been the no-change benchmark – a naïve forecast that simply takes the last observed value of the series of interest to predict future values. Indeed, an increasing number of studies document that model-based forecasts of the real price of crude oil outperform the simple no-change benchmark (Baumeister and Kilian, 2012; Alquist *et al*., 2013; Baumeister *et al*., 2014; Baumeister and Kilian, 2014, 2015; Snudden, 2018; Funk, 2018; Garratt *et al*., 2019). This evidence has been used to conclude that the real price of oil is predictable in general, and that econometrics models are more useful to forecast prices than a naïve approach.

Our paper challenges this conclusion by observing that the real price of crude oil is typically constructed as an average monthly price of daily data. We show that this seemingly innocent transformation invalidates the conventional interpretation of forecast comparisons with the no-change benchmark. Instead, we propose an alternative no-change benchmark that is based on monthly closing prices. The new benchmark re-establishes meaningful forecast comparisons and has large effects on the evaluation of different oil-price forecasts.

Why We Need a Different Benchmark to Evaluate Forecasts of the Real Price of Crude Oil

The appeal of the no-change benchmark originates from its relationship with the random walk model. The random walk model has been used to approximate the behavior of many economic and financial series. It assumes that any future observation of a series is given by its last observed value plus a random

innovation. Under the random walk hypothesis, any future changes in the series are unpredictable, and the no-change forecast is the optimal forecast for all future observations. Consequently, forecast-improvements over the no-change benchmark allow forecasters to reject the random walk hypothesis and to claim that the series of interest is predictable.

Dr. Reinhard Ellwanger, Ph.D., Senior Economist, Bank of Canada, presenting at a J.P. Morgan Center for Commodities (JPMCC) international commodities symposium at the University of Colorado Denver Business School.

However, this general logic fails when the series of interest is constructed by averaging higher-frequency data. This is the case for many macroeconomic variables such as real interest rates and real commodity prices, which are often expressed as deflated monthly or quarterly averages of daily observations. It can be shown that under the random walk hypothesis, the averaged data do not follow a random walk, but rather a cumulative sum of a moving-average process (Working, 1960). This implies that averaged series have a predictable pattern by design, even when all future changes in the daily data are completely random. For such series, improvements over the conventional no-change forecast are not informative about the predictability of the underlying series or the practical usefulness of a specific forecasting approach.

Our paper shows that the original interpretation of forecast comparisons can be restored by introducing an alternative no-change benchmark. This benchmark is not the last value of the averaged monthly or quarterly series that we want to predict, but rather the last value of the underlying high-frequency

observation. Under the random walk null hypothesis, the new no-change benchmark is the optimal forecast for all future observations, including the averaged series.

The difference in the forecast accuracy between these two no-change benchmarks can be sizeable. When the underlying series follows a random walk, the theoretical improvements in the one-step-ahead mean squared prediction error (MSPE) are larger than 45 percent when using the last observed daily value instead of the last observed monthly or quarterly average value. A simple change of the benchmark can thus have large effects on assessments of different oil-price forecasting models.

Closing Prices Drastically Improve Model-Based Forecasts of the Real Price of Crude Oil, But the New Benchmark is Difficult to Beat

We study the importance of these effects in the case of the real price of crude oil. The focus of the empirical application is real-time forecasts of monthly averages of oil prices, which is the standard approach in the literature. For this purpose, we update a real-time dataset for oil-market and other economic variables created by Baumeister and Kilian (2012).

As in the typical setup of existing studies, it is assumed that the forecaster uses the available information at the end of each month to form their prediction for the following months. Based on our theoretical insights, we construct a new benchmark from real monthly closing prices and revisit the claim that these models can predict the real price of crude oil. We also investigate the extent to which the use of closing prices can improve model-based forecasts more generally.

The main empirical result from these exercises is that replacing average prices with closing prices considerably improves traditional forecasting approaches for the real price of oil. A simple no-change forecast based on the last closing price reduces the MSPE of the conventional no-change forecast computed from average monthly prices by 40 percent for one-month-ahead forecasts. The directional accuracy for the one-month-ahead forecasts is higher than 70 percent. The gains decrease with the forecast horizon but are still apparent up to the 12-months-ahead forecast.

The magnitude and the pattern in the forecast-improvements are roughly consistent with the theoretical predictions of a random walk model for daily oil prices. For example, the theoretical improvements from using the new benchmark instead of the conventional, average-price benchmark for the 1-, 3- and 12-month horizons are 46, 12 and 3 percent, respectively; see Figure 1 on the next page. The empirical counterparts for the real price of WTI (West Texas Intermediate) crude oil are 39, 11 and 4 percent, respectively. We show that this pattern arises because all forecasting gains from using closing prices are realized at the one-step-ahead prediction and become relatively less important for longer-horizon forecasts.

Figure 1
Accuracy-Improvements for Forecasts of the Real Price of WTI Crude Oil

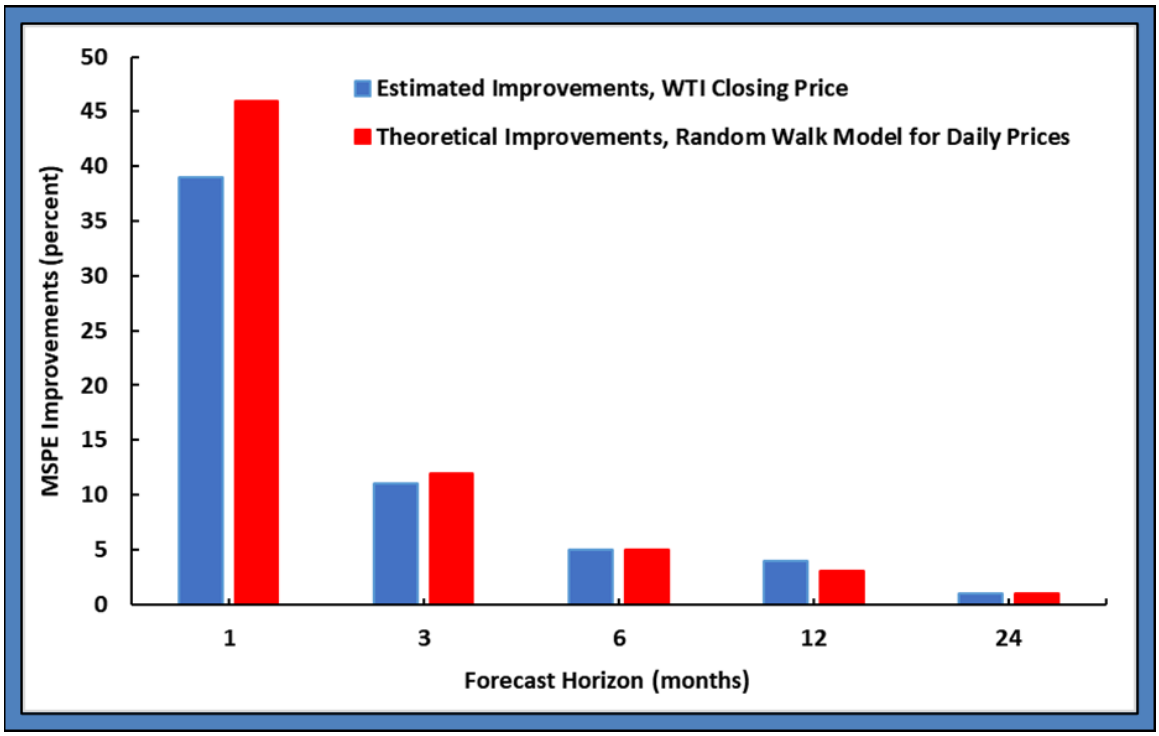

Another major result is that forecasting models of the price of oil should be estimated with the monthly closing price, even if the goal is to predict average prices. Traditionally, models of the real price of crude oil have been estimated with the same series the forecaster wants to predict: the monthly average price of oil. We show that the forecasts derived from several popular forecasting models – including univariate time-series models, vector-autoregressive models, a futures-based forecast and a simple forecast combination – improve considerably when these models are estimated with closing prices instead.

At the one-month horizon, forecasts from models that are estimated with closing prices produce large improvements of about 40 percent over the average no-change forecast. As for the new no-change forecast, accuracy-gains are especially significant for shorter forecast horizons and become less pronounced for longer-horizons forecasts. We document that these gains are remarkably robust to the choice of the crude oil benchmark and the sample period. By contrast, most of the forecasts that are derived from the same models are unable to beat the conventional no-change forecast when models are estimated with closing prices. This suggests that closing prices should be used to estimate these models even if the forecaster's goal is to predict average prices.

How do models that are estimated with closing prices fare against the new benchmark? Although the model-based forecasts in some cases show lower MSPE ratios and better directional accuracy than the new no-change benchmark, these improvements are rarely statistically significant and consistent across both criteria. Only the futures-based forecasts for 1 and 2 years are both economically and statistically more accurate than the new benchmark. This shows that the choice of the benchmark matters and that

improvements over the conventional benchmark are not necessarily indicative that oil prices are predictable.

Implications for Forecasts of the Real Price of Crude Oil and of Other Commodity Prices

Our findings have two broader implications. First, the introduction of a new benchmark can raise the bar for model-based forecasts to claim improvements over the no-change forecast. We show that this is indeed the case for the real price of crude oil. Forecasts that are generated from several popular models often outperform the conventional no-change benchmark, especially when these models are estimated using closing prices instead of average prices. However, they generally do not improve upon the new benchmark. Only the futures-based forecast provides better forecasts than the monthly closing-price benchmark and only for horizons of one year and beyond.

These results suggest that real oil prices are more difficult to predict and, in this sense, closer to asset prices than implied by the previous literature. They also suggest that closing prices provide better measures of oil price expectations than many models that rely on average prices. As such, the use of closing prices could shed new light on the transmission of oil price shocks and on the predictive power of oil prices for other macroeconomic variables.

The second implication concerns policymakers and applied forecasters who forecast averaged data. Our results highlight that incorporating information from high-frequency observations can yield large gains even in the context of the simple models many practitioners prefer. Such gains are likely to occur in any setting where forecasters work with averaged data and the underlying series are very persistent. This includes the prices of other primary commodities. In this environment, one would expect forecasts from econometric models to beat the conventional no-change forecasts that are based on averaged data, even if the underlying data used to obtain the averaged series is entirely or approximately unpredictable. For policymakers and applied forecasters, the easiest way to improve traditional forecasts for such series – particularly for short- and medium-term horizon – is to rely on the last higher-frequency observation rather than on lower-frequency averages.

References

Alquist, R., Kilian, L. and R. Vigfusson, 2013, "Forecasting the Price of Oil," in G. Elliott and A. Timmermann (eds) Handbook of Economic Forecasting, Vol. 2A, Amsterdam: North-Holland, pp. 427-507. [Drs. Kilian and Vigfusson are members of the JPMCC's Research Council.]

Baumeister, C. and L. Kilian, 2012, "Real-Time Forecasts of the Real Price of Oil," *Journal of Business & Economic Statistics*, Vol. 30, No. 2, pp. 326-336.

Baumeister, C. and L. Kilian, 2014, "What Central Bankers Need to Know about Forecasting Oil Prices," *International Economic Review*, Vol. 55, No. 3, August, pp. 869-889.

Baumeister, C. and L. Kilian, 2015, "Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach," *Journal of Business & Economic Statistics*, Vol. 33, No. 3, pp. 338-351.

Baumeister, C., Kilian, L. and T.K. Lee, 2014, "Are There Gains from Pooling Real-Time Oil Price Forecasts?", *Energy Economics*, Vol. 46, Supp. 1, December, pp. S33-S43. [Dr. Lee is a member of the *GCARD*'s Editorial Advisory Board.]

Funk, C., 2018, "Forecasting the Real Price of Oil - Time-Variation and Forecast Combination," *Energy Economics*, Vol. 76(C), pp. 288-302.

Garratt, A., Vahey, S. and Y. Zhang, 2019, "Real-Time Forecast Combinations for the Oil Price," *Journal of Applied Econometrics*, Vol. 34, No. 3, April-May, pp. 456-462.

Snudden, S., 2018, "Targeted Growth Rates for Long-Horizon Crude Oil Price Forecasts," *International Journal of Forecasting*, Vol. 34, No. 1, January-March, pp. 1-16.

Working, H., 1960, "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica*, Vol. 28, No. 4, October, pp. 916-918.

Author Biographies

AMOR ANISS BENMOUSSA
Economist, Bank of Canada

Mr. Amor Aniss Benmoussa is an Economist in the International Projection Division at the Bank of Canada. He holds an M.Sc. in Economics from the Université du Québec à Montréal (UQAM).

REINHARD ELLWANGER, Ph.D.
Senior Economist, Bank of Canada

Dr. Reinhard Ellwanger is a Senior Economist in the International Economic Analysis department at the Bank of Canada. He received a Ph.D. in Economics from the European University Institute in Florence (Italy) in 2015. Dr. Ellwanger had last co-authored an article for the *GCARD* on "The Effect of Oil-Price Shocks on Asset Markets: Evidence from Oil Inventory News."

STEPHEN SNUDDEN, Ph.D.
Assistant Professor, Wilfrid Laurier University, Canada

Dr. Stephen Snudden is an Assistant Professor at the Department of Economics at Wilfrid Laurier University. He received his Ph.D. in Economics from Queen's University (Canada) in 2019 and has six years of combined work experience at the Bank of Canada and the International Monetary Fund.