Futures prices are useful predictors of the spot price of crude oil*

Reinhard Ellwanger† Stephen Snudden‡§

September 21, 2021

Abstract: How well do futures prices forecast the spot price of crude oil? Contrary to the established view, futures prices significantly improve upon the accuracy of monthly no-change forecasts. This results from two innovations. First, we document that independent of the construction of futures-based forecasts, longer-horizon futures prices have become better predictors of crude oil spot prices since the mid-2000s. Second, we show that futures curves constructed using end-of-month prices instead of average prices have consistently been able to generate large accuracy-improvements for short-horizon forecasts of average prices. These findings are remarkably robust and apply to both WTI and Brent crude oil prices.

JEL classification: C1, C53, G13, G15, Q47
Keywords: Forecasting and Prediction Methods, Oil Price Forecast, Futures Market

---

*The views expressed in this paper are those of the authors; no responsibility for them should be attributed to the Bank of Canada.

†Bank of Canada, International Economic Analysis Department

‡Corresponding author; Wilfrid Laurier University, Department of Economics; ORCID iD 0000-0002-4033-3235; Email: ssnudden@wlu.ca

§We thank Lutz Kilian, Ron Alquist, Christiane Baumeister, Edouard Djeutem, Konrad Zmitrowicz, Pierre Siklos, Wing Chan and Willy Chetwin for detailed discussions, comments and suggestions. This paper supersedes results contained in the unpublished Bank of Canada Working Paper "The New Benchmark for Forecasts of the Real Price of Crude Oil".

1 Introduction

Given the importance of crude oil as a macroeconomic determinant used in models of central banks, in investment decisions, and in oil-intensive goods purchases, there is wide interest in accurately predicting the price of crude oil.[1] A frequently cited result in this literature is that futures prices are not particularly useful to forecast the spot price of crude oil (Alquist and Kilian, 2010; Reeve and Vigfusson, 2011; Baumeister and Kilian, 2012; Alquist et al., 2013; Baumeister and Kilian, 2014). In this paper, we document that future-based forecasts have always been useful for short-horizon forecasts of the average spot price of crude oil and are now very accurate at longer horizons. This occurs for two reasons.

First, we show that that futures curves constructed using end-of-month prices contain substantive predictive power of future average prices at short-run horizons. Incorporating end-of-period information improves the mean-squared prediction error (MSPE) and the directional accuracy of no-change forecast for average spot prices by 40 percent at the one-month horizon and remains statistically significant for forecasts up to the 12-month-ahead horizon. These improvements are remarkably robust and independent of the sample period.

Second, the predictive content of crude oil futures prices at longer forecast horizons has improved since the mid-2000s. We show that whenever the end of the forecast evaluation period is extended beyond 2014, futures-based forecasts are found to be statistically significantly more accurate predictors of the spot price than the no-change forecast at horizons of one-year to five-years ahead. This result holds for forecasts of both nominal prices (as in Alquist and Kilian, 2010) and real prices (as in Baumeister and Kilian, 2012). It is particularly strong for multi-year-ahead forecasts, which were previously difficult to evaluate due the lack of a proper evaluation period for longer-dated futures contracts. This and corroborates and generalizes some existing evidence on the usefulness of future-based forecasts of real prices at longer horizons.[2]

A key insight from our forecast exercises is that futures-based forecasts should always be constructed using end-of-the-period prices rather than the average price that is standard in forecasts of the real price of oil (Baumeister and Kilian, 2012; Alquist et al., 2013). This seemingly innocuous difference yields substantial improvements for short-horizon forecasts. The use of closing prices

[1] The literature on forecasting oil prices is vast and includes Ye et al. (2005); Baumeister and Kilian (2012); Alquist et al. (2013); Baumeister et al. (2014); Baumeister and Kilian (2014, 2015); Wang et al. (2015); Yin and Yang (2016); Snudden (2018); Zhang et al. (2018); Funk (2018); Garratt et al. (2019), among many others.

[2] For example, Baumeister and Kilian (2016) document statistically significant forecast-improvements for the real price of crude oil at the one-year horizon. Similar improvements for predicting real prices between the 9-month and two-year horizon have been document for quarterly by Manescu et al. (2016) and for monthly forecasts by Funk (2018).

is preferable because averaging changes the underlying data process (Rossana and Seater, 1995), which leads to a mechanical loss of information when forecasting persistent processes (Lütkepohl, 1984; Ellwanger and Snudden, 2021). Futures curves constructed using end-of-month prices has precedent in Alquist and Kilian (2010) in the context of forecasting of the end-of-month nominal crude oil price. However, the model was never used to forecast average spot prices, which are standard for real prices.[3] While the distinction between end-of-period and average futures prices is crucial, the role of the deflator appears to be negligible for these results: using information from end-of-period futures prices work similarly well for forecasting average nominal and real prices.

The results contribute to a recurring debate on the role of futures prices in forecasting oil prices. Because of their simplicity and ease of implementation, futures-based forecast are popular among policy makers, investors and market participants. Early results from the academic literature documented some predictability from oil futures prices, at least at specific horizons (Ma, 1989; Kumar, 1992; Chinn et al., 2005; Coppola, 2008). However, influential reviews that extended the sample period beyond the early 2000s found little evidence that futures-based forecasts are helpful to forecast oil prices and recommended against their use (Alquist and Kilian, 2010; Baumeister and Kilian, 2012). Still, futures prices have remained a steady ingredient in the construction of forecast combinations (Baumeister and Kilian, 2015; Funk, 2018; Garratt et al., 2019); while other approaches focused on improving futures-based forecast by separating their expectations component from the risk premium (Baumeister and Kilian, 2016). Our results show that a decade's worth of additional data and as well as a simple modification to the originally proposed implimentation change the assessment of futures-based forecasts of the price of oil.

2 Futures-Based Forecasts

The goal of this paper is to study the predictive content of oil futures for the spot price of crude oil. The spot price of oil is determined in market transactions for immediate delivery and used as a key macroeconomic indicator by both academics and policy makers (see, e.g. Alquist and Kilian, 2010). By contrast, futures contracts are financial instruments that allow traders to lock in today a price at which to buy or sell a fixed quantity of oil on a predetermined date in the future. The inherent forward-looking nature of these contracts contributes to the wide-spread interest of futures as predictors of spot prices.

[3] The only exception in the literature seems to be Funk (2018), who documents gains in the accuracy of one-month-ahead forecasts from using closing prices. We show that when closing prices are applied to the futures spread, these gains are measurable and statistically significant for up to 1 year ahead.

Following existing practice, we construct futures-based forecasts for monthly data using the percentage spread of the futures price with maturity h, $F^h_{t,x}$, over the spot price, $S_{t,x}$.[4] For nominal prices, the h-step-ahead forecast is

$$\hat{S}_{t+h,x|t} = S_{t,x}\left(1 + \ln F^h_{t,x} - \ln S_{t,x}\right), \tag{1}$$

where t denotes the current month, h denotes the forecast horizon, and x is an indicator for the price series that distinguishes between end-of-period observations ($x = n$) and monthly average observations ($x = a$). The monthly average price is the average of the n daily closing prices within the month, $S_{t,a} \equiv \frac{1}{n}\sum_{i=1}^{n} S_{t,i}$. In a similar fashion, futures-based forecasts for real prices are constructed via:

$$\hat{R}_{t+h,x|t} = R_{t,x}\left(1 + \ln F^h_{t,x} - \ln S_{t,x} - E_t\left(\pi^h_t\right)\right), \tag{2}$$

where $E_t\left(\pi^h_t\right)$ is the expected U.S. inflation rate over the next h periods, and $R_{t,x}$, $x \in \{a, n\}$, is the monthly measure of the real price of crude oil.

The distinction between average and end-of-period futures and spot prices clarifies the use of alternative futures-spreads and spot prices in the literature. For example, Alquist and Kilian (2010) forecast the end-of-month nominal spot price, $\hat{S}_{t+h,n|t}$ using end-of-month futures prices $F^h_{t,x} = F^h_{t,n}$.[5] By contrast, average prices are typically considered to be more economically relevant for the construction of macroeconomic variables and total pay-offs to oil investments. As such, the standard series for the real price of oil used in structural work and forecasting applications has always been the average monthly price, $R_{t,a}$ (see, e.g., Kilian, 2009; Alquist et al., 2013; Baumeister and Kilian, 2014, 2015).

In most applications, the futures-based forecast for real prices, $R_{t,a}$, has been constructed with average futures prices, $F^h_{t,a}$.[6] However, averaging oil prices can lead to a loss of information about price levels relative to end-of-period observations (Benmoussa et al., 2020). A key contribution of this study is to systematically compare the relative forecast performance of end-of-period futures prices, $F^h_{t,n}$, and average futures prices, $F^h_{t,a}$, for the real price of oil.

[4] The use of monthly data and log-percentage spreads is standard in the literature, see, e.g., Alquist and Kilian (2010); Baumeister and Kilian (2012); Alquist et al. (2013).

[5] Alquist and Kilian (2010) average over the last 3-5 trading days of the month. See also Pak (2018). However, we find that any averaging reduces forecast accuracy of future average prices. Thus, our analysis focuses on the closing price of the last trading day of the month.

[6] See, e.g., Alquist et al. (2013); Baumeister and Kilian (2014, 2015). The only exception to this rule seems to be Funk (2018), who used monthly average real spot prices, $R_{t,x} = R_{t,a}$ and $S_{t,x} = S_{t,a}$ but end-of-the month futures prices $F^h_{t,x} = F^h_{t,n}$ when documenting significant one-month-ahead forecasts of the real price of crude oil.

3 Application to Real-Time Forecasts

We construct monthly, recursive, real-time, and out-of-sample forecasts following Baumeister and Kilian (2012). Our baseline estimates are for West Texas Intermediate (WTI) crude oil. Forecasts for spot prices of another popular global crude oil benchmark, Brent, are considered in the robustness analysis. Monthly spot prices for both benchmarks were obtained from the Energy Information Administration (EIA). Daily futures prices were collected from Haver. Contracts for different delivery months are traded with maturies up to 7 years. Contracts are used as continuous series, $F_{t,i}^{h}$, with $h = 1$ referring to the front contract, $h = 2$ referring to the second back contract, etc. Because trading for the front contract stops several days prior to the last business day of the month preceding delivery, the front contract rolls over to the next contract before the end of the month, typically around the 21st of the month.[7] This means that the standard average futures price is based on prices for two adjacent contracts.[8]

For forecast horizons of up to one year, our sample starts in 1992M1. For longer-term contracts, the sample starts when the contract began to be traded at least once per day: 1995M5 for the 2-year-contract, 2006M3 for the 3-year contract, and 2007M8 for the 5-year contract.[9] For all forecasts, the sample period ends in 2021M1.

To construct real prices, monthly prices are deflated using real-time vintages of seasonally adjusted U.S. consumer price index obtained from the FRASER database of the Federal Reserve Bank of St. Louis and the Philadelphia Federal Reserve. Expected inflation is derived from the CPI price index which is projected using the historical average for CPI inflation from 1986M7.

Forecast evaluation is conducted using the mean-squared prediction error (MSPE) ratio and the success ratio for directional accuracy. Consistent with the literature, both measures are compared against the no-change constructed from the last available observation of the forecasted price series. Diebold-Mariano (Diebold and Mariano, 1995) tests are used for the MSPE ratios to test the null hypothesis of equal predictability, whereas Pesaran and Timmermann (2009) tests are used to test the null hypothesis of random directional accuracy.

[7] Trading terminates 3 business day prior to the 25th calendar day of the month prior to the contract month. If the 25th calendar day is not a business day, trading terminates 4 business days prior to the 25th calendar day of the month prior to the contract month. See https://www.cmegroup.com/markets/energy/crude-oil/light-sweet-crude.contractSpecs.html.

[8] Forecasts based on average prices from a single contract are discussed in the robustness section.

[9] An alternative approach to dealing with sparsely traded futures is to impute missing observations. However, while this approach extends the sample size, we found that it does not materially impact our results.

4 Forecast Results

Our main results focus on three different forecasting exercises, displayed in columns one to three in Table 1. The first column considers forecasts of the nominal end-of-month price of oil using end-of-the-month futures prices, as in Alquist and Kilian (2010). The futures-based forecasts exhibit lower MSPEs than the no-change forecasts at all forecast horizons. While the improvements are modest in magnitude and not statistically significant at short forecast horizons, they are large and significant for longer horizon. At the 1-year forecast horizon, the futures-based forecast reduces the MSPE of the no-change forecast by 13%, which is significant at the 5% significance level. Forecasts beyond one year are even better, with improvements close to 20% for the 2-year and 6-year horizon and more than 55% at the 3-year horizon. These improvements are statistically significant at the 1% significance level. A similar pattern emerges from the success ratios which compare the relative directional accuracy of the futures-based forecast with the no-change forecast. The futures-based forecast is slightly worse for short-term forecasts, but both economically and statistically significantly more accurate than the no-change forecasts for forecasts of 1-year ahead and longer. For these forecasts, the directional accuracy is over 60%, peaking at 87% at the 3-year horizon.

These results differ from Alquist and Kilian (2010), who focused on forecasts of up to 1-year, but found that futures prices typically performed worse than no-change forecasts in terms of the MSPE and only marginally better in terms of directional accuracy. We show that for a longer sample period that extends to 2021, futures prices are significantly (both economically and statistically) better predictors than the no-change forecast beginning at the 1-year horizon.

The results thus far apply to the full sample evaluations. To see how the forecast performance evolved, Figure 1 displays the evolution of the MSPE ratios and directional accuracy from recursively updated sample periods. Interestingly, the relative forecast performance at the one-year horizon started to improve considerably after 2007, just when the sample of Alquist and Kilian (2010) ended. By 2010, the gains were close to 10% and remained relatively stable. The p-values with the test for equal predictability associated with the recursively updated 1-year-ahead forecasts indicate that the null hypothesis of equal predictability can be rejected at the 5% confidence level for any estimation period that extends beyond December 2013. One possible explanation for these results is that volume of traded future contracts with longer maturities have gradually increased over the last decades, in particular during the 2000s (Alquist and Kilian, 2010). Much of this increase was associated with a rising participation of investors in futures markets, which caused

Table 1. Futures-based forecasts of monthly prices of WTI crude oil

Forecasted Price	Nominal, EoM	Real, Average		Sample Start
Futures Model	EoM	Average	EoM	
Horizon		MSPE Ratio		
1	**0.99 (0.207)**	1.00 (0.425)	**0.59 (0.000)**	1992m1
3	**0.96 (0.060)**	**0.96 (0.078)**	**0.87 (0.016)**	1992m1
6	**0.95 (0.083)**	**0.95 (0.130)**	**0.91 (0.053)**	1992m1
12	**0.87 (0.010)**	**0.85 (0.025)**	**0.84 (0.015)**	1992m1
24	**0.79 (0.012)**	**0.81 (0.068)**	**0.80 (0.054)**	1995m5
36	**0.58 (0.000)**	**0.45 (0.000)**	**0.44 (0.000)**	2006m3
48	**0.67 (0.000)**	**0.48 (0.000)**	**0.47 (0.000)**	2007m3
60	**0.80 (0.053)**	**0.48 (0.000)**	**0.48 (0.000)**	2007m8
		Success Ratio		
1	0.47 (0.274)	0.47 (0.458)	**0.72 (0.000)**	1992m1
3	0.49 (0.703)	0.50 (0.436)	**0.60 (0.001)**	1992m1
6	**0.54 (0.184)**	**0.56 (0.074)**	**0.56 (0.050)**	1992m1
12	**0.62 (0.003)**	**0.64 (0.000)**	**0.66 (0.000)**	1992m1
24	**0.67 (0.000)**	**0.62 (0.001)**	**0.63 (0.000)**	1995m5
36	**0.87 (0.000)**	**0.86 (0.000)**	**0.85 (0.000)**	2006m3
48	**0.70 (0.000)**	**0.80 (0.000)**	**0.78 (0.000)**	2007m3
60	**0.62 (0.014)**	**0.76 (0.009)**	**0.77 (0.008)**	2007m8

Notes: Futures-based forecasts of the monthly price of WTI crude oil. *EoM* stands for end-of-the-month prices, *Average* for average prices. The start of sample is 1992m1 or the month when at least one contract at the respective maturity is traded per day; the end of the sample is 2021M1. Bold values indicate improvements over no-change forecast. The brackets show the p-values for serial-dependence-robust tests of Diebold and Mariano (1995) for the null hypothesis of equal MSPEs and of Pesaran and Timmermann (2009) for the null of no directional accuracy, both relative to the no-change forecast.

broad-based concerns about potential market distortions. Contrary to these concerns, our results suggest that this influx could have provided additional liquidity and potentially even improved the informational efficiency of longer-term futures.

Qualitatively similar results are obtained for forecasts of the real price of oil when the forecasts are constructed with the average-futures price as in Baumeister and Kilian (2012). The results are displayed in column 2 of Table 1. Except for very short forecast horizons up to 6 months, the futures based-forecasts outperform the no-change forecast, both in terms of the MSPE ratio and the success ratio. These improvements are statistically significant at conventional levels for all horizons of one year and beyond. For the 3-year and 5-year horizon, the futures-based forecast reduce the MSPE by over 50% and have a directional accuracy of more then 80%.

Figure 1 shows that similar to the forecasts for nominal end-of-the-month prices, the standard futures-based forecasts for the real price of oil improved considerably after the mid-2000s. Using a sample period until 2010, Baumeister and Kilian (2012) found that at forecast horizons up to

Figure 1. Recursive MSPEs and success ratios for futures-based forecasts of WTI crude oil prices

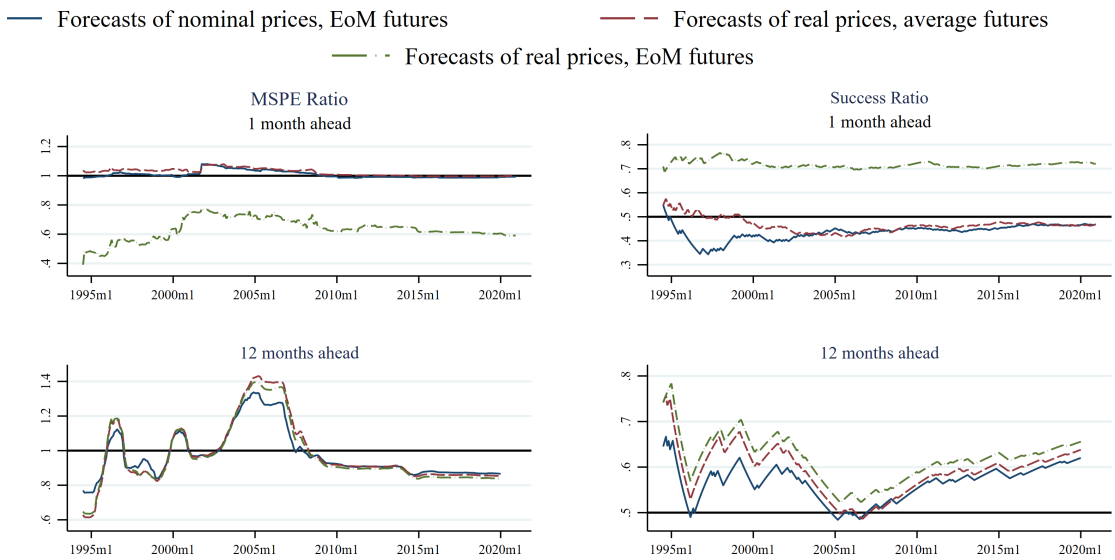

Notes: Evolution of recursively estimated forecast criteria for 1-month-ahead and 1-year-ahead forecasts of WTI oil prices. *EoM futures* stands for end-of-the-month futures prices, *average futures* for the monthly average futures price. The MSPEs are expressed as a ratio relative to the no-change forecast. The success ratios represent the fraction of times the forecast correctly predicts the direction of the change in the price of oil. The estimation period is 1992M1–2021M1, the first 30 months are not displayed to reduce starting-point effects.

one year, futures-based forecast were only marginally, and generally insignificantly, better than the no-change forecast. For larger sample periods, however, these forecasts are economically and statistically significant even at the one-year horizon. Our results show that the forecast performance of the standard futures-based forecast of the real price of oil steadily improved over the 2010 and became statistically significant for samples ending in June 2014, which explains more favorable results reported in more recent studies.[10] At all horizons, the improvements in the MSPE and in the directional accuracy remained robust throughout the COVID-19 episode.

Finally, the 3rd column of Table 1 considers the forecasts of the real price of oil that are constructed using end-of-the-month futures. Contrary to the standard futures-based forecast, which are constructed with the average price, forecasts based on end-of-the-month futures significantly outperform the no-change forecast at all forecasting horizons. For the one-month-ahead forecast, the MSPE reduction is over 40%, and the directional accuracy is 72%. These gains in forecast accuracy exceed gains typically found in studies advocating for introducing new models, predictor variables, or forecast combination techniques (see, e.g. Baumeister and Kilian, 2015; Funk, 2018; Garratt et al., 2019).

The short-run forecast accuracy of end-of-the-month futures-based forecasts is exceptionally

[10] e.g., Baumeister and Kilian (2016); Manescu et al. (2016); Funk (2018); Garratt et al. (2019).

robust. Figure 1 shows that at the one-month ahead horizon, the forecasting gains according to both criteria are large throughout the sample period. The MSPE gains are statistically significant at the 1 and 5 percent level for 97.2 and 100 percent of the sample, respectively. This degree of robustness is impressive in a literature that is plagued with unstable forecast performances across different samples.[11]

At all horizons, the forecast based on end-of-the-month futures is at least as good as the average futures-based forecast. For the one-step ahead prediction, the improvements from introducing information from end-of-the-month prices are statistically significant at the 1% significance level for 93.6 percent of the sample, including the end of the sample. Likewise, for all forecast-horizons up to one year, the null hypothesis of equal forecast performance in terms of the MSPE and the success ratios can be rejected at the 10% significance level. Even though the improvements are somewhat weaker for longer-horizon forecasts, our results show that futures-based forecasts should always be constructed with end-of-month futures prices rather than average-price futures prices. Our results are remarkably robust to alternative oil price series and modelling assumptions. End-of-the-month futures-based forecasts work similarly well when real prices of Brent or the average nominal WTI or are forecasted (rows "Nominal WTI, EoM" and "Real Brent, EoM" in Table 2, respectively). It can also be shown that our results are robust to alternative nowcasts of the real CPI index or the use of ex-post revised data instead of real-time data.[12] More generally, the role of the price deflator is minimal, which is unsurprising given the relatively low variability of US inflation compared that of nominal oil prices. Finally, the improvements in the one-step-ahead prediction are even larger for quarterly and annual data, which are of primary interest to policymakers (Baumeister and Kilian, 2014, 2015).

The forecast performance is also robust to alternative futures-based model assumptions. Using end-of-month futures prices directly (rows "EoM, Direct" in Table 2) instead of the futures-spread model has a negligible quantitative effect. Moreover, model extensions to include recursive estimates of the intercept and/or slope parameter fail to consistently improve forecasts for either nominal or real prices, which is consistent with the exercises presented in Alquist and Kilian (2010).

Averaging futures prices could be helpful if prices are distorted by random noise arising from measurement error or market-microstructure dynamics. However, introducing any averaging over daily futures prices decreases the performance of the futures-based forecasts. For example, the

[11] The fact that the performance of oil-price forecasts is sensitive to the sample period is well-documented, e.g., in Baumeister et al. (2014); Baumeister and Kilian (2015); Snudden (2018); Funk (2018); Garratt et al. (2019).

[12] These results are not displayed for the sake brevity.

Table 2. Robustness to alternative oil-price series and modelling assumptions

Horizon	1	3	6	12	36
Model			MSPE Ratio		
EoM	**0.59 (0.000)**	**0.87 (0.016)**	**0.91 (0.053)**	**0.84 (0.015)**	**0.44 (0.000)**
EoM, Direct	**0.59 (0.000)**	**0.87 (0.015)**	**0.91 (0.046)**	**0.83 (0.011)**	**0.46 (0.000)**
EoM 3 Day Ave.	**0.62 (0.000)**	**0.87 (0.011)**	**0.92 (0.055)**	**0.85 (0.014)**	**0.47 (0.000)**
End of Contract	1.10 (0.834)	**0.97 (0.198)**	**0.97 (0.235)**	**0.87 (0.029)**	**0.48 (0.000)**
Contract Average	1.43 (0.999)	1.08 (0.921)	**0.99 (0.372)**	**0.86 (0.038)**	**0.48 (0.000)**
Real Brent, EoM	**0.62 (0.000)**	**0.93 (0.048)**	**0.97 (0.233)**	**0.79 (0.013)**	**0.59 (0.000)**
Nominal WTI, EoM	**0.57 (0.000)**	**0.86 (0.007)**	**0.91 (0.032)**	**0.84 (0.005)**	**0.58 (0.000)**
			Success Ratio		
EoM	**0.72 (0.000)**	**0.60 (0.001)**	**0.56 (0.050)**	**0.66 (0.000)**	**0.85 (0.000)**
EoM, Direct	**0.72 (0.000)**	**0.60 (0.001)**	**0.56 (0.052)**	**0.65 (0.000)**	**0.85 (0.000)**
EoM 3 Day Ave.	**0.72 (0.000)**	**0.60 (0.000)**	**0.56 (0.057)**	**0.66 (0.000)**	**0.85 (0.000)**
End of Contract	**0.53 (0.102)**	**0.55 (0.052)**	**0.53 (0.135)**	**0.62 (0.000)**	**0.85 (0.000)**
Contract Average	0.39 (1.000)	0.47 (0.803)	**0.51 (0.350)**	**0.61 (0.001)**	**0.86 (0.000)**
Real Brent, EoM	**0.70 (0.000)**	**0.55 (0.039)**	**0.57 (0.024)**	**0.65 (0.000)**	**0.81 (0.000)**
Nominal WTI, EoM	**0.72 (0.000)**	**0.60 (0.003)**	**0.59 (0.024)**	**0.60 (0.009)**	**0.87 (0.000)**

Notes: Futures-based forecasts for alternative models and measures of the monthly price of crude oil. *EoM* stands forecasts based on end-of-the-month closing prices; "EoM, Direct" stands for direct forecasts based on end-of-month closing prices; "EoM 3 Day Ave." stands for forecasts based on the average futures price over the three last business days of each month; "End of Contract" stands for forecasts based on the last available price of the front contract before the rollover date; "Contract Average" stands for forecasts based on the average monthly price of the front contract before the rollover date, "Real Brent" stands for forecasts of real Brent prices; "Nominal WTI" stands for forecasts of nominal WTI prices. The start of sample is 1992m1, apart from the 36-months-ahead forecasts for WTI and Brent which begin in 2006m3 and 2005m2, respectively. Bold values indicate improvements over no-change forecast. The brackets show the p-values for serial-dependence-robust tests of Diebold and Mariano (1995) for the null hypothesis of equal MSPEs and of Pesaran and Timmermann (2009) for the null of no directional accuracy, both relative to the no-change forecast.

futures-based forecast based on an average price over the last trading days yields slightly worse MSPE ratios and does not improve on the Success Ratio (rows "EoM 3 day average" in Table 2). The forecast improvements we obtain from the use of end-of-period observations instead of averaged observations indicate that random noise in daily futures prices is largely irrelevant in practice. The result is not unexpected given that oil futures prices are determined in liquid markets. It suggests that similar to other assets, market-microstructure noise in oil futures prices can safely be ignored at the daily frequency (see, e.g. Hansen and Lunde, 2006).

Because the front contract changes over the course of the month (and hence the $h-th$ contract more generally) , we also consider forecasts based on the available information from a single contract. Using the daily or the average monthly price for a contact based on the information on the date when the front contact rolls over (rows "End of Contract" and "Contract Average", respectively, in Table 2) significantly deteriorates the short run forecasts. This is consistent with the idea that the closing price on the last trading day of the month contains important additional information

that is not available on the roll-over date.

5 Conclusion

Contrary to established views, futures prices are useful predictors of the spot price of crude oil. Futures curves constructed with end-of-month prices have always been able to generate robust one-step-ahead forecast accuracy improvements of over 40 percent for the average spot price. Moreover, the predictive content of longer horizon forecasts has improved considerably since the mid-2000s. Forecasts several years ahead perform particularly well, with reductions in the MSPE relative to the no-change forecast of over 55 percent at the 3-year horizons. One potential explanation for this result is that oil futures have become more effective predictors with the influx of investors over the 2000s and the associated decrease in the risk premium component of futures prices (Hamilton and Wu, 2014).

In the view of these results, previous recommendations against futures-based forecasts seem outdated. Futures-based forecast are easy to implement in real-time, offering policymakers, investors, and market participants straightforward forecasts of average crude oil prices that are significantly better than no-change forecasts. As such, they also provide a natural point of comparison to evaluate the usefulness of model-based forecasts of the price of crude oil, which have become increasingly popular over the last decade.

Our results also raise the question if similar improvements could be found for futures-based forecasts of other commodity prices. While historically, futures-based forecasts have proven to perform quite differently across different commodities (Chinn and Coibion, 2014), there is reason to believe that the insights of our paper apply more broadly. First, as suggested by Ellwanger and Snudden (2021), the gains from using end-of-period prices instead of average should apply to any persistent series, including the prices of other commodities. Second, the futures markets of non-oil commodities have experienced a similar process of financialization as the oil futures markets, which could, in principle, have contributed to similar changes in the predictive content of futures in these markets. Revisiting the predictive content of other commodity prices in the light of these findings opens a promising avenue for future research.

References

Alquist, R. and Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4):539–573.

Alquist, R., Kilian, L., and Vigfusson, R. (2013). Forecasting the price of oil. In Elliott, G., Granger, C., and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 2, chapter 8, pages 427–507. Elsevier, 1 edition.

Baumeister, C. and Kilian, L. (2012). Real-time forecasts of the real price of oil. *Journal of Business & Economic Statistics*, 30(2):326–336.

Baumeister, C. and Kilian, L. (2014). What central bankers need to know about forecasting oil prices. *International Economic Review*, 55(3):869–889.

Baumeister, C. and Kilian, L. (2015). Forecasting the real price of oil in a changing world: A forecast combination approach. *Journal of Business & Economic Statistics*, 33(3):338–351.

Baumeister, C. and Kilian, L. (2016). A general approach to recovering market expectations from futures prices with an application to crude oil. *Bank of Canada Staff Working Paper, 2016-18*.

Baumeister, C., Kilian, L., and Lee, T. K. (2014). Are there gains from pooling real-time oil price forecasts? *Energy Economics*, 46:S33–S43.

Benmoussa, A. A., Ellwanger, R., and Snudden, S. (2020). The new benchmark for forecasts of the real price of crude oil. *Bank of Canada Staff Working Paper, 2020-39*.

Chinn, M. D. and Coibion, O. (2014). The predictive content of commodity futures. *Journal of Futures Markets*, 34(7):607–636.

Chinn, M. D., LeBlanc, M., and Coibion, O. (2005). The predictive content of energy futures: an update on petroleum, natural gas, heating oil and gasoline. *NBER Working Paper No. 11033*.

Coppola, A. (2008). Forecasting oil price movements: Exploiting the information in the futures market. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, 28(1):34–56.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.

Ellwanger, R. and Snudden, S. (2021). Predictability of aggregated time series. *LCERPA Working Paper No. 2021-5*.

Funk, C. (2018). Forecasting the real price of oil-time-variation and forecast combination. *Energy Economics*, 76:288–302.

Garratt, A., Vahey, S. P., and Zhang, Y. (2019). Real-time forecast combinations for the oil price. *Journal of Applied Econometrics*, 34(3):456–462.

Hamilton, J. D. and Wu, J. C. (2014). Risk premia in crude oil futures prices. *Journal of International Money and Finance*, 42:9–37.

Hansen, P. R. and Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business & Economic Statistics*, 24(2):127–161.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–69.

Kumar, M. S. (1992). The forecasting accuracy of crude oil futures prices. *IMF Staff Papers*, 39(2):432–461.

Lütkepohl, H. (1984). Forecasting contemporaneously aggregated vector ARMA processes. *Journal of Business & Economic Statistics*, 2(3):201–214.

Ma, C. W. (1989). Forecasting efficiency of energy futures prices. *The Journal of Futures Markets (1986-1998)*, 9(5):393.

Manescu, C. B., Van Robays, I., et al. (2016). Forecasting the brent oil price: Addressing time-variation in forecast performance. *CESifo Working Paper No. 6242*.

Pak, A. (2018). Predicting crude oil prices: Replication of the empirical results in "What do we learn from the price of crude oil?". *Journal of Applied Econometrics*, 33(1):160–163.

Pesaran, M. H. and Timmermann, A. (2009). Testing dependence among serially correlated multicategory variables. *Journal of the American Statistical Association*, 104(485):325–337.

Reeve, T. A. and Vigfusson, R. J. (2011). Evaluating the forecasting performance of commodity futures prices. *FRB International Finance Discussion Paper*, (1025).

Rossana, R. J. and Seater, J. J. (1995). Temporal aggregation and economic time series. *Journal of Business & Economic Statistics*, 13(4):441–451.

Snudden, S. (2018). Targeted growth rates for long-horizon crude oil price forecasts. *International Journal of Forecasting*, 34(1):1–16.

Wang, Y., Liu, L., Diao, X., and Wu, C. (2015). Forecasting the real prices of crude oil under economic and statistical constraints. *Energy Economics*, 51:599–608.

Ye, M., Zyren, J., and Shore, J. (2005). A monthly crude oil spot price forecasting model using relative inventories. *International Journal of Forecasting*, 21(3):491–501.

Yin, L. and Yang, Q. (2016). Predicting the oil prices: do technical indicators help? *Energy Economics*, 56:338–350.

Zhang, Y., Ma, F., Shi, B., and Huang, D. (2018). Forecasting the prices of crude oil: An iterated combination approach. *Energy Economics*, 70:472–483.