

Carpe Diem: Can daily oil prices improve model-based forecasts of the real price of crude oil?*

Amor Aniss Benmoussa[†] Reinhard Ellwanger[‡] Stephen Snudden^{§¶}

March 28, 2023

Abstract

The standard approach in the literature is to compute model-based forecasts of the real price of crude oil with monthly average prices. We show how existing model-based forecasts approaches can be extended by incorporating the underlying daily oil prices. For both univariate and multivariate methods, alternative disaggregated methods are compared, including bottom-up, Period-End-Price Sampling (PEPS), and Mixed Data Sampling (MIDAS). We find that using unaveraged oil prices, especially end-of-month prices, yields large gains across a variety of forecast approaches. In some cases, forecast error is more than halved compared to existing specifications. Contrary to models estimated with monthly average prices, models that utilize the underlying daily prices can outperform the random walk forecast of daily prices at short horizons. The analysis informs how to efficiently construct forecasts for other series expressed in real terms to avoid information loss from temporal aggregation.

JEL classification: C1, C53, Q47

Keywords: Forecasting and Prediction Methods, Oil Prices

*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; Bank of Canada, International Economic Analysis Department, 234 Wellington Street W, Ottawa, Ontario, K1A 0G9, Canada; ORCID iD 0000-0002-9521-3388; Email: rellwanger@bankofcanada.ca

[§]Wilfrid Laurier University, Department of Economics

[¶]We thank Ron Alquist, Christiane Baumeister, Wing Chan, Willy Chetwin, Edouard Djeutem, Lutz Kilian, Pierre Siklos, Luis Uzeda, and Konrad Zmitrowicz, as well as seminar and conference participants at the Bank of Canada, the European Economic Association Annual Congress, the Canadian Economics Association Annual Conference, the International Association for Energy Economics International Conference, and the International Symposium on Forecasting for their discussions, comments and suggestions. This paper supersedes any related findings in “The New Benchmark for Forecasts of the Real Price of Crude Oil,” an unpublished Bank of Canada Working Paper.

1 Introduction

The real price of crude oil is a key global macroeconomic indicator that has attracted considerable attention among forecasters. A growing literature documents that a variety of different model-based forecasts of the monthly real price of oil outperform a simple monthly average no-change forecast (see, e.g. 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; Alquist et al., 2020; Baumeister et al., 2022). A common feature of these studies is that they all rely on a price of crude oil that represents a monthly average over daily prices. This form of temporal aggregation is convenient because it facilitates the inclusion of monthly predictor variables and the construction of historical prices (Alquist et al., 2013).¹ However, time-averaging also introduces a loss of information and is generally inefficient from a forecasting perspective (see, e.g., Wei, 1978; Kohn, 1982; Lütkepohl, 1986; Marcellino, 1999). Although it remains an empirical question whether estimated models for the disaggregated series fare better in practice (Lütkepohl, 1984), there is thus far no study that exploits the underlying daily oil prices for forecasting models of the real price of crude oil.

The goal of this paper is to explore alternative disaggregated approaches that incorporate daily oil prices into model-based forecasts of the monthly average real price of crude oil. The first exercise examines the forecast performance of the random walk forecast based on the last closing price relative to no-change forecasts computed from averaged data. The second exercise compares the relative effectiveness of two different disaggregated approaches for univariate time-series models. This includes Lütkepohl (1984)’s approach of estimating models for daily oil prices and aggregating the daily forecasts to the monthly frequency ex-post. We also consider forecasts via period-end price sampling (PEPS), which estimates monthly models for end-of-month real prices and uses forecasts of the end-of-month price as the forecast of the monthly average price (Ellwanger and Snudden, 2021). The third exercise evaluates the performance of mixed frequency (MIDAS) forecasts that use daily oil prices for direct forecasts of the monthly real price of oil. The fourth exercise incorporates information from daily oil prices into direct forecasts with predictor variables, while the fifth exercise extends the PEPS approach to vector autoregression (VAR) models. Together, these exercises provide a systematic assessment over how daily oil prices can be utilized in model-based forecasts

¹Structural models and forecasting approaches using pre-1980es data typically rely on historical refinery acquisition costs of crude oil, which are only available at the monthly frequency (Kilian, 2009; Alquist et al., 2013; Kilian and Murphy, 2014; Baumeister and Hamilton, 2019).

of the monthly average real price of oil.

Our key result is that irrespective of the approach, forecasting models that utilize information from daily prices are considerably more effective than standard models estimated with monthly average prices. Across several distinct forecast approaches that rely on daily prices, we obtain robust accuracy-improvements of the order of 40% at the one-month horizon. Moreover, forecasts based on daily data exhibit improvements in the success ratio for directional accuracy relative to models estimated with average data as far as 24 months ahead.

An additional insight from our exercises is that much of the benefit from using daily oil prices stems from the information contained in the last available daily price. For example, we find that in most cases, MIDAS regressions attribute the entire forecasting weight on the end-of-month price and zero weight on any other daily prices. As a consequence, forecasting models relying on monthly rather than daily data provide effective forecasts as long as they incorporate information from end-of-month prices rather than monthly average prices. This helps explain why the PEPS approach, which allows forecasters to maintain the monthly frequency for all model variables, fairs well across a wide variety of model-based applications.

A practical advantage of PEPS is that it can also be applied to forecasts of U.S. refinery acquisition cost of crude oil (RAC), which is often used in lieu of other crude oil prices because of its longer history (see, e.g., Alquist et al., 2013). Using disaggregated methods for the RAC is complicated by the fact that the RAC is published as a monthly survey. To circumvent the lack of daily data, we impute end-of-month RAC observations from WTI prices and estimate models using the PEPS approach with the imputed data. This approach is easy to implement for both ex-post and real-time data and yields large forecast gains over traditional forecasts, with over 30 percent improvements in the MSFE for real-time forecasts at the one-month-ahead horizon. Large forecast improvements are also obtained by estimating the multivariate VAR model of Kilian and Murphy (2014) via PEPS instead of the standard monthly average oil price.

The use of daily oil prices can resuscitate the benefits of some model-based forecasts over naive forecast at short horizons. This is nontrivial because at the one-month ahead prediction, all existing forecasting models estimated with monthly average data perform worse than the random walk forecast based on the last available closing price (Ellwanger and Snudden, 2022). By contrast, models that incorporate information from daily prices sometimes outperform the random walk at this horizon, albeit by a smaller margin than implied by comparisons with the monthly average no-change forecast.

Although the idea to utilize the underlying daily oil-price data in forecasts of the monthly real price of oil is new, our results are consistent with several distinct pieces of evidence provided in earlier studies. For example, Ellwanger and Snudden (2023) have recently shown that using end-of-month oil futures prices instead of monthly average oil futures prices yields large forecasting gains for futures-based forecasts of the real price of crude oil. Baumeister and Kilian (2014) document that the monthly average no-change forecasts performs better than the quarterly average no-change forecast. Baumeister and Kilian (2012) show that most forecasting gains from using revised data stem from improved information on prices, while Bork et al. (2022) and Conlon et al. (2022) document that averaging commodity prices can introduce spurious predictability in the context of return-forecasting regressions. These results are all indicative that time-averaging introduces a loss of information about future oil prices. Herein, we show that the introduction of information from daily prices into model-based forecasts can help to restore this information.

2 Forecast Methodology and Data

The focus of our empirical application is real-time forecasts of the level of the monthly average real price of crude oil, which is the standard approach in the literature (Baumeister and Kilian, 2012; Alquist et al., 2013). In this setting, the forecaster uses the available information at the end of each month to form their prediction for the following months. The goal is to forecast the average real price, $R_{t,a} = p_{t,a}/CPI_t$, where R_t stands for the real price on month t , $p_{t,a}$ stands for the nominal average price, and CPI_t is the U.S. consumer price index. We study forecasts of the three most common oil-price benchmarks in the literature: West Texas Intermediate (WTI), Brent, and US refiner acquisition cost of crude oil, imported (RAC). For WTI and Brent, $p_{t,a}$ is the monthly average of daily nominal closing prices, whereas the RAC is a monthly survey price.

All forecasts are computed in real time and evaluated out-of-sample with an expanding window beginning in 1992M1. Unless otherwise indicated, the evaluation window ends in 2021M1. Throughout the paper, we follow existing practices and report MSFE ratios relative to that of the conventional monthly average no-change forecast, along with the p-values for the null hypothesis of equal predictability following Diebold and Mariano (1995).² Similarly, we report success ratios for the direction of the change in the monthly average price along with the test of Pesaran and

²The use of real-time data and the iterative out-of-sample forecasts does not fulfill the assumptions underlying the Diebold-Mariano test (Kilian, 2015). As in many pre-existing studies, the p-values are still reported with this warning in mind for those forecasts.

Timmermann (2009). Reporting these statistics facilitates the comparison with results from the existing literature.

Additionally, our tables indicate statistically significant improvements over the random walk forecast. The random walk forecast is the no-change forecast based on the last observed closing price contained in the forecaster’s information set: $R_{t,n}^x = p_{t,n}^x / CPI_t$ where n denotes the last day of month t .³ While assessing improvements in terms of the MSFE-accuracy relative to the random walk forecast is straightforward, the assessment of the success ratio for directional accuracy contains an important subtlety. To test whether a particular model outperforms the random walk in terms of directional accuracy, the direction of the forecast is computed as the difference between the model-implied forecast and the random walk forecast, $R_{t,n}^x$. Unlike the MSFE ratios, the success ratios are not transitive across the comparison against the different no-change forecasts. For example, the forecast with the highest success ratio for the direction of the change relative to the monthly average price is not necessarily the forecast with the highest success ratio for the direction of the change relative to the last observed closing price.

Crude oil prices are obtained from the U.S. Energy Information Administration (EIA). For Brent and WTI, the EIA reports the average monthly price as simple averages of the daily closing price. Both average and daily prices for these series are available in real time. Implementing forecast approaches relying on disaggregated data for the RAC is complicated by the fact that daily RAC data does not exist. For our empirical analysis, we impute end-of-period RAC observations by applying the growth rate of the end-of-month WTI oil price over the monthly average WTI oil price to the monthly RAC price: $p_{t,n}^{RAC} = p_{t,a}^{RAC} \cdot (p_{t,n}^{WTI} / p_{t,a}^{WTI})$, where $p_{t,n}$ refers to the nominal closing price on the last day of the month and superscripts RAC and WTI denote the nominal RAC series and nominal WTI prices, respectively. The idea to impute RAC observations with WTI prices has ample precedent in the oil-price forecasting literature. For example, it is standard to nowcast the average nominal RAC via the growth rate of WTI prices and to replace missing RAC futures prices with WTI futures prices (Baumeister and Kilian, 2012). Although the RAC technically differs from the price of WTI, the tight empirical correlation between the series implies that this imputation works well for forecasting. Our empirical results show that this applies not only to monthly averages, but also to end-of-period observations.

³The no-change forecast based on the last observed instantaneous price is the optimal no-change forecast of the monthly average data under the null hypothesis that *all* future prices are unpredictable (Working, 1960). As such, predictability in the traditional sense can only be established for with the random walk forecast based on closing prices, even for real prices (Ellwanger and Snudden, 2022)

Nominal prices are converted to real prices using the seasonally adjusted U.S. consumer price index. The price index is the seasonally adjusted U.S. consumer price index obtained from the FRASER database of the Federal Reserve Bank of St. Louis and the real-time database of the Philadelphia Federal Reserve.

Our exercises also present forecasts from the Kilian and Murphy (2014) VAR model, which includes oil prices, the growth rate of global crude oil production, a proxy for the change in global crude oil inventories, and an indicator of real economic activity. An additional contribution of this paper is to extend the real-time data of Baumeister and Kilian (2012) using historical data vintages from the EIA’s Monthly Energy Review and Short-Term Energy Outlook. Unlike earlier updates of this dataset provided in Garratt et al. (2019), we were able to collect all monthly vintages after 2010M12, which means this is the first update to include all revisions for the entire history of each vintage.⁴ Real-time data vintages start in 1991M12 and contain historical data starting from 1973M1. Real-time data on U.S. crude-oil inventories, U.S. petroleum inventories and OECD petroleum inventories are obtained from historical releases in the EIA’s Monthly Energy Review or the International Data Browser. The real-time version of Kilian (2009)’s real economic activity index is computed using the corrected formula (Kilian, 2019).

All nowcasts follow Baumeister and Kilian (2012). For example, the monthly average RAC is nowcasted using the month-over-month growth rate of the monthly WTI series. Any missing real-time observations for the consumer price index are nowcasted using the average historical growth rate. To nowcast the imputed end-of-period observations for the RAC, we use the historical vintage of the nowcasted RAC prices at the time of the forecast. As discussed in Section 4, none of the main results are affected by different nowcasting choices.

The series of daily and monthly WTI and Brent prices provided by the EIA begin in 1986M1 and 1987M5, respectively. Estimation for models estimated at the daily frequency begin on those dates. For forecasting models estimated at the monthly frequency, including models estimated via PEPS, prices are backcasted to 1973M1. Specifically, monthly average and end-of-month prices for WTI and Brent are backcasted to 1983M4 using the growth rate of the monthly average and end-of-month price of the front WTI futures contract, respectively. Then, following the standard practice, all monthly average prices before 1983M4 are backcasted using the growth rate of the monthly average RAC. For the backcast of end-of-month prices before 1984M4, we apply the same growth rate to the last available end-of-month observation rather than the last observed monthly

⁴See Appendix A1 for details.

average observation. None of these backcasting choices are crucial for our results, as our main findings are robust to estimating models with data starting in 1983 or 1986.

3 Forecasting the Real Price of Crude Oil with Disaggregated Oil Prices

3.1 No-Change Forecasts

This section compares the relative performance of no-change forecasts that are averaged at different frequencies. To compute no-change forecasts for the h -months-ahead real price of crude, $R_{t+h,a}$, backward averages of nominal daily prices in period t are deflated by the monthly CPI,

$$\hat{R}_{t+h,a} = \frac{\frac{1}{s} \sum_{j=n-s+1}^n p_{t,j}}{CPI_t}, \quad (1)$$

where $p_{t,i}$ refers to the nominal price on day i of month t , n indicates the number of total days in month t , and s indicates the number of end-of-month daily observations that are used for the average. Ellwanger and Snudden (2022) show that the random walk forecast, $s = 1$, is a more effective predictor of the monthly average real price of crude oil than the conventional monthly average no-change forecast, $s = n$. The first three columns of Table 1 replicate this result and report the MSFEs and success ratios for the random-walk forecasts relative to the monthly average no-change forecasts. The random-walk forecasts are significantly more accurate than the monthly average no-change forecasts, with accuracy-improvements of over 40 percent for Brent and WTI at the one-month-ahead-horizon. The gains decrease as the forecasting horizon increases, but they never vanish completely and are still statistically significant at the 5 percent significance level for up to 24-month-ahead forecasts. Similar results are obtained for the random walk forecast of the RAC, which relies on the imputed end-of-month price. The forecast gains are stable across the forecast evaluation sample (Ellwanger and Snudden, 2022).

We now examine whether shorter averages over daily data provides additional gains in the performance of the end-of-period no-change forecasts. Price-averaging could be potentially useful if the daily series is affected by random noise, which could arise from measurement errors or market-microstructure dynamics. Table 1 reports the no-change forecasts that are computed from averaging daily closing prices over the last two trading days, $s = 2$, the last week of trading, $s = 5$, and the last two weeks of trading before the end of the month, $s = 10$, for the case of the real price of

Table 1. Alternative No-Change Forecasts of the Real Price of Crude Oil

Price	Brent	WTI	RAC	WTI	WTI	WTI
Sampling	EoM	EoM	WTI EoM	EoM 2 Day Avg.	EoM 1 Week Avg.	EoM 2 Week Avg.
Horizon	MSFE Ratio					
1	0.58 (0.000)	0.59 (0.000)	0.70 (0.001)	0.60 (0.000)	0.67 (0.000)	0.78 (0.001)
3	0.92 (0.017)	0.89 (0.012)	0.89 (0.010)	0.89 (0.017)	0.92 (0.025)	0.95 (0.071)
6	0.98 (0.123)	0.95 (0.029)	0.95 (0.027)	0.96 (0.044)	0.98 (0.155)	0.99 (0.297)
12	0.98 (0.170)	0.96 (0.028)	0.96 (0.018)	0.98 (0.116)	1.00 (0.473)	1.01 (0.714)
24	1.00 (0.443)	0.99 (0.189)	0.98 (0.123)	1.00 (0.440)	1.01 (0.838)	1.01 (0.936)
Horizon	Success Ratio					
1	0.72 (0.000)	0.71 (0.000)	0.71 (0.000)	0.70 (0.000)	0.70 (0.000)	0.67 (0.000)
3	0.57 (0.005)	0.61 (0.000)	0.61 (0.000)	0.61 (0.000)	0.59 (0.000)	0.57 (0.006)
6	0.56 (0.013)	0.55 (0.057)	0.58 (0.006)	0.56 (0.039)	0.55 (0.043)	0.55 (0.047)
12	0.59 (0.000)	0.57 (0.007)	0.57 (0.009)	0.58 (0.003)	0.56 (0.017)	0.52 (0.322)
24	0.56 (0.015)	0.52 (0.292)	0.54 (0.153)	0.51 (0.415)	0.49 (0.737)	0.47 (0.880)

Notes: Forecast performance of alternative no-change forecasts relative to the monthly average no-change forecasts for alternative crude oil price series, 1992M1–2021M1. *EoM* stands for end-of-month observations, *Avg.* stands for averages. Brackets report the p-values for serial dependence robust tests of Pesaran and Timmermann (2009) for the null of no directional accuracy and Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast.

WTI crude oil. Very similar results are obtained for the real price of Brent crude oil. These results indicate that there is no benefit from averaging, as the no-change forecasts tend to perform more poorly as the length of the averaging window increases. This shows that, for forecasting the real price of crude oil, the informational benefits from using end-of-month prices clearly outweigh the concerns over increasing the volatility relative to average prices.⁵

Table 1 highlights that even the simplest possible forecast that relies on disaggregated data yields accuracy-gains over the conventional monthly average no-change forecast used in the literature. The magnitude of these gains is impressive, in particular for short forecast horizons. In fact, as shown in Ellwanger and Snudden (2022), all existing forecast models relying on monthly average data have higher MSFEs than the random walk forecast at the one-step-ahead prediction. One goal of this paper is to examine whether model-based forecasts that rely on disaggregated prices can improve upon the random walk forecast at this horizon.

⁵The empirical variances of the monthly closing prices of crude oil are statistically indistinguishable from the variances of the monthly average prices of crude oil (see Appendix A1). This suggests that the improvements from using closing prices is driven by a smaller bias.

3.2 ARIMA Forecasts

This section compares the forecast performance of two different forecast approaches that incorporate disaggregated data in the context of autoregressive-moving-average (ARMA) models (Box et al., 2015). These univariate time-series models are a regular feature of forecasting exercises and are often found to be more accurate predictors of the real price of crude oil than the conventional average-price no-change forecast (Baumeister and Kilian, 2012).

The first approach estimates models using daily closing price data and then aggregates the daily forecasts to the monthly frequency, as proposed by Lütkepohl (1984). We henceforth refer to this approach as bottom-up approach. Lütkepohl (1986) showed that forecasts from the bottom-up approach are usually superior to forecasts obtained from the aggregated process, even when the process is unknown and parameters have to be estimated.

The second approach is period-end price sampling (PEPS), which estimates models at the monthly frequency but replaces monthly average prices by end-of-month prices, as proposed by Ellwanger and Snudden (2021). The end-of-month forecasts values are then used as forecasts of the average value of the corresponding month. Ellwanger and Snudden (2021) document that for persistent data, the forecast-accuracy of the PEPS approach are often of similar magnitude to that of bottom-up approach.

We examine of the forecast performance of both the bottom up and PEPS approach, as well as their relative performance, by applying these methods for the first time to the monthly forecasts of the real price of oil. Forecast via ARMA models estimated with monthly average and end-of-month real prices are estimated for the log real price. Forecasts are converted to real prices in levels using

$$\hat{R}_{t+h|t} = \exp(\hat{r}_{t+h|t}), \quad (2)$$

where $\hat{R}_{t+h|t}$ is the forecast of the monthly average real price of oil in levels, $\hat{r}_{t+h|t}$ is the forecast of the log real price, and \exp is the exponential function.

For the bottom up approach, we estimate ARMA models using the growth rate of nominal daily prices, $g_{t,i} = p_{t,i}/p_{t,i-1}$. Forecasts for the level of the nominal price on day i of month $t+h$ given month t information are based on the model-implied cumulative net growth rate between day t, n and day $t+h, i$, denoted by $\hat{g}_{t+h,i|t}$:

$$\hat{p}_{t+h,i|t} = (1 + \hat{g}_{t+h,i|t}) \cdot p_{t,n}. \quad (3)$$

Then, nominal daily forecasts are averaged to the monthly frequency, as in Lütkepohl (1984), and converted into forecasts of real prices by deflating the monthly nominal forecasts by the expected CPI deflator. This conversion of the nominal forecast to real terms is akin to the adjustment made to nominal futures prices in forecasts of the real price of crude oil (see, e.g., Baumeister and Kilian, 2012). Following these authors, we compute the expected CPI by extending the CPI using the historical inflation rate beginning in 1986M7.⁶

Table 2. Forecasts of the Real Price of WTI Crude Oil Based on ARMA Models

Method	Monthly Average Prices			Bottom Up			Period-End-Price Sampling		
Model	AR(12)	AR(2)	ARMA(1,1)	AR(12)	AR(2)	ARMA(1,1)	AR(12)	AR(2)	ARMA(1,1)
Horizon	MSFE Ratio								
1	0.93 (0.239)	0.90 (0.139)	0.91 (0.130)	0.61 (0.000)	0.60 (0.000)	0.60 (0.000)	0.58 (0.002)	0.56 (0.001)	0.56 (0.001)
3	0.98 (0.440)	0.93 (0.275)	0.94 (0.227)	0.93 (0.021)	0.92 (0.018)	0.94 (0.022)	0.87 (0.126)	0.84 (0.113)	0.84 (0.088)
6	1.02 (0.565)	0.94 (0.334)	0.95 (0.287)	1.04 (0.723)	1.03 (0.666)	1.05 (0.753)	0.94 (0.271)	0.89 (0.215)	0.89 (0.180)
12	1.05 (0.632)	0.94 (0.354)	0.93 (0.265)	1.16 (0.908)	1.14 (0.888)	1.16 (0.913)	0.96 (0.362)	0.90 (0.250)	0.89 (0.190)
24	1.13 (0.728)	0.97 (0.441)	0.92 (0.302)	1.48 (0.989)	1.42 (0.987)	1.45 (0.989)	0.99 (0.470)	0.95 (0.385)	0.92 (0.308)
	Success Ratio								
1	0.50 (0.519)	0.52 (0.294)	0.54 (0.066)	0.72 (0.000)	0.72 (0.000)	0.72 (0.000)	0.72 (0.000)	0.71 (0.000)	0.72 (0.000)
3	0.48 (0.745)	0.50 (0.599)	0.53 (0.088)	0.61 (0.000)	0.61 (0.000)	0.61 (0.000)	0.57 (0.008)	0.58 (0.002)	0.59 (0.000)
6	0.49 (0.641)	0.50 (0.519)	0.50 (0.403)	0.57 (0.042)	0.56 (0.068)	0.55 (0.100)	0.54 (0.191)	0.53 (0.276)	0.54 (0.096)
12	0.55 (0.175)	0.54 (0.263)	0.54 (0.155)	0.55 (0.143)	0.55 (0.142)	0.53 (0.408)	0.55 (0.196)	0.54 (0.261)	0.56 (0.064)
24	0.59 (0.079)	0.59 (0.059)	0.56 (0.111)	0.51 (0.958)	0.51 (0.718)	0.51 (0.900)	0.59 (0.060)	0.61 (0.035)	0.56 (0.086)

Note: Real-time, out-of-sample forecasts for the real price of WTI crude oil, 1992M1–2021M1. “Bottom Up” is the ex-post averaged forecast based on daily prices (Lütkepohl, 1984); “Period-End-Price Sampling” is the forecast based on monthly end-of-month prices (Ellwanger and Snudden, 2021). Brackets report the p-values for serial dependence robust tests of Pesaran and Pesaran and Timmermann (2009) for the null of no directional accuracy and Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the random walk forecast at the five percent significance level.

Table 2 shows the result for the different forecasting approaches for AR(2), AR(12), and ARMA(1,1) models. Similar results are obtained for alternative lag specifications considered in the literature.⁷

The first three columns of Table 2 display the results for the models estimated with the conventional monthly average data. For the majority of the models and forecast horizons considered, they outperform the average price no-change forecasts, in line with Baumeister and Kilian (2012). However, the improvements appear relatively modest (below 10 percent) and do not constitute significant improvements against the conventional monthly average no-change forecast at the five percent level at any horizon. Only for directional accuracy at the 24 month horizon do they out-

⁶The CPI forecasts thus correspond directly to the standard inflation forecasts computed to discount nominal prices in the literature (see, e.g., Baumeister and Kilian, 2015; Garratt et al., 2019; Ellwanger and Snudden, 2023).

⁷Previous studies advocate for the use of 12 or 24 lags when estimating autoregressive model to predict crude oil prices (Alquist and Kilian, 2010; Baumeister and Kilian, 2012). However, our results show that for the models estimated at a monthly frequency, the AR(2) outperforms the AR(12) at most forecasting horizons.

perform the random walk forecast at the 5 percent level.

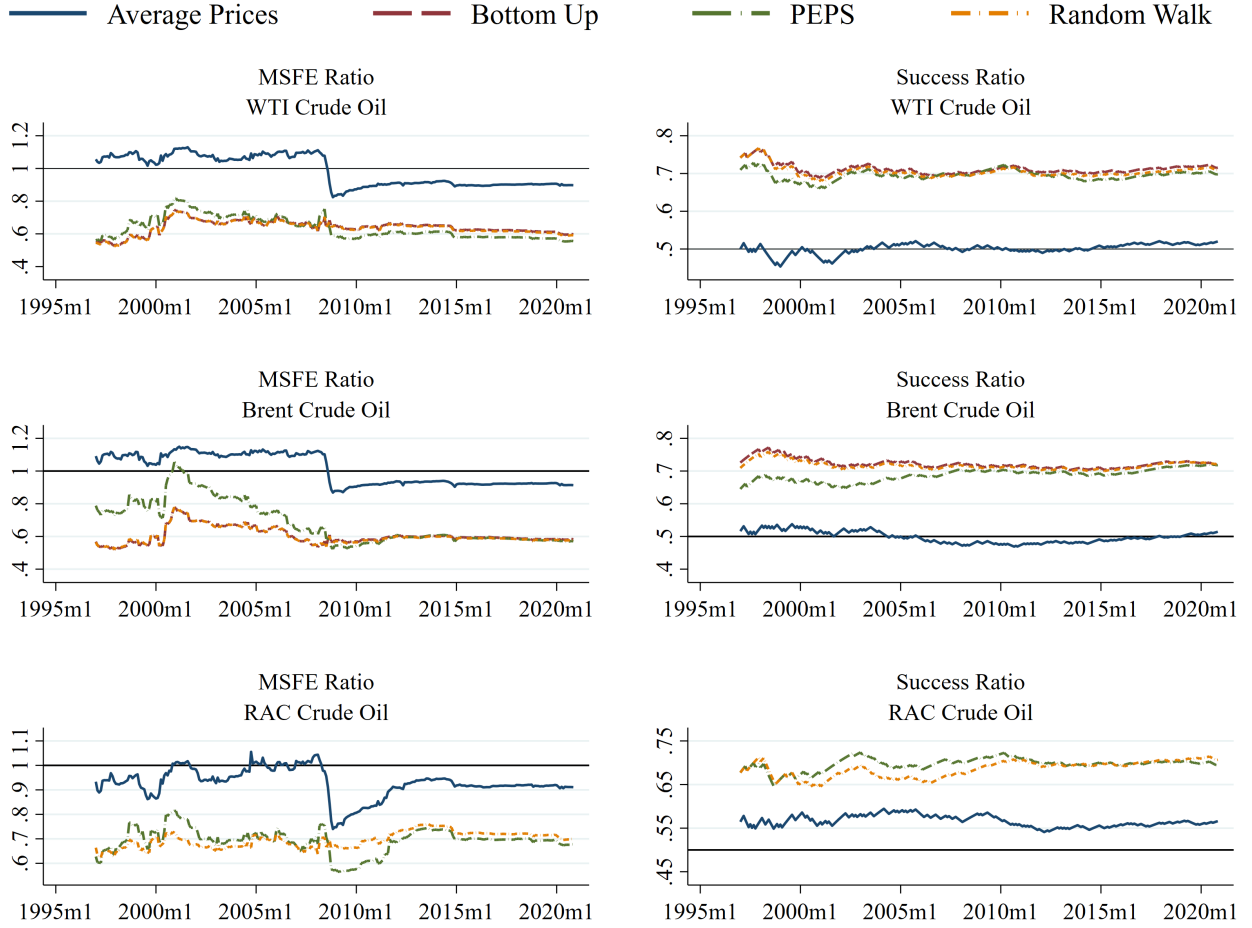
The corresponding results for the bottom-up approach, displayed in columns 4-6 of Table 2, display a very different pattern. At the one-month horizon, all three models exhibit about 40 percent improvements in the MSFE and a directional accuracy of 72 percent. The forecasts are more accurate than any other one-month-ahead forecasts previously reported in the literature, as summarized by Ellwanger and Snudden (2022). However, the absence of bold-faced values highlights that at a 5 percent significance level, none of these forecasts are significantly more accurate than the random walk forecast. Moreover, the performance of the bottom-up forecast deteriorates quickly for longer forecast horizons. For 6-months-horizons and beyond, the MSFE-accuracy of the bottom-up approach is worse than even the monthly average no-change forecast.

The forecasts based on the PEPS approach are evaluated in the final 3 columns of Table 2. For all horizons and models, the forecasts outperform the conventional monthly average no-change benchmark. At the one-step-ahead prediction, the PEPS approach results in the largest improvements in forecast accuracy of up to 44 percent, which are significant at the 0.1 percent level. For all three models, forecast criteria, and for all horizons, the models estimated with end-of-period prices provide at least equally good forecasts of average prices than models estimated with average prices. Moreover, the bold values show that the AR-models significantly predict the direction of change in prices at the one-month-ahead relative to the random walk forecast at the five percent level.

Similar forecast gains are also obtained for RAC and Brent crude oil. Figure 1 reports the evolution of the MSFE and success ratios for the one-month-ahead forecast for the AR(2) models estimated with average, end-of-month, and daily data. The results suggest that the forecast performance of the bottom-up approach closely tracks that of the random-walk forecast throughout the forecast evaluation sample, whereas the PEPS forecasts tend to outperform the random walk forecast in the later half of the sample. For the bottom-up and PEPS approach, the gains upon the monthly average no-change forecast at the one-month-ahead horizon are significant at the five percent level for WTI and RAC crude oil prices for all potential ending dates of the forecast evaluation sample between 1997M1 and 2020m1. For all series and for all potential ending points of the evaluation sample, the models estimated with average prices fail to outperform the random walk.

Overall, these results show that suitably constructed univariate time-series models that exploit disaggregated prices outperform models estimated with monthly average prices. Moreover, models based on disaggregated oil prices can also provide improvements over the random walk forecast at

Figure 1. Evolution of One-Month-Ahead Forecast Evaluation Criteria for AR(2) Models



Note: One-Month-Ahead, real-time, out-of-sample forecasts of alternative crude oil price series, 1992M1–2021M1. Forecasts from AR(2) models estimated with Monthly “Average Prices”, the “Bottom-Up” approach of (Lütkepohl, 1984), and the “PEPS” approach of (Ellwanger and Snudden, 2021). “Random Walk” is the random walk forecast based on the last observed closing price. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

the one-step-ahead prediction.

3.3 MIDAS Forecasts

Another approach that relies on disaggregated data is mixed-data sampling (MIDAS, see Ghysels et al., 2007; Andreou et al., 2010), which allows forecasters to construct direct forecasts of lower-frequency variables with higher frequency predictors. Previous applications of the MIDAS model to oil-price forecasts in have focused on the role of high-frequency financial variables as predictors (Baumeister et al., 2015; Degiannakis and Filis, 2018; Zhang and Wang, 2019). Instead, this section uses the MIDAS setting to explore the relationship between daily oil prices and the monthly real

price of crude oil.

To investigate the predictive content of daily oil prices for the real price of crude oil, we estimate MIDAS models of the form

$$R_{t+h,a} = a_h + \beta_h B(L^{1/n}; \theta_h) R_{t,a} + \epsilon_{t+h}, \quad (4)$$

where $R_{t+h,a}$ is the monthly average real price of crude oil, $B(L^{1/n}; \theta_h) = \sum_{i=0}^N c(i; \theta) L^{i/m}$, and $C(1; \theta) = 1$. We explore both the normalized exponential Almon lag weight function and the normalized beta lag polynomial for alternative lag lengths of up to 20 days of daily data. Interestingly, the full-sample estimates are conclusive that for WTI and RAC, the coefficients on all daily prices except the last daily price are zero. A similar pattern is obtained for Brent, except for the 3- and 6-month-ahead forecasts, for which there are some small non-zero coefficients on longer lags. The finding is consistent with the results from Section 3.1 and supports the idea that the last observed daily oil price contains practically all the relevant information for forecasting.

Thus, in almost all cases, the MIDAS regressions of Equation 4 simplify to the following regression equation:

$$R_{t+h,a} = \alpha_h + \beta_h R_{t,n} + \epsilon_{t+h}, \quad (5)$$

i.e., only the last observed daily closing price is included in the MIDAS forecast and used to provide direct forecasts of future average prices. For the remainder of the section, we therefore work with MIDAS forecasts using the parsimonious specification of Equation 5, which has the additional benefit of facilitating the comparison across different forecast approaches.

The MIDAS forecasts of Equation 5 represent direct forecast of the real price of oil using real end-of-month closing prices, which can also implemented via PEPS. In this case, the model is estimated for real end-of-month closing prices as the dependent variable,

$$R_{t+h,n} = \alpha_h + \beta_h R_{t,n} + \epsilon_{t+h}. \quad (6)$$

Then, the end-of-month price is used as the forecast of the monthly average price. Finally, we contrast the MIDAS and PEPS forecasts with conventional direct forecasts based on average prices,

$$R_{t+h,a} = \alpha_h + \beta_h R_{t,a} + \epsilon_{t+h}. \quad (7)$$

For all forecasts, parameters are estimated recursively using real-time data and forecasts are com-

puted out-of-sample.

Table 3. Direct Forecasts of the Real Price of Crude Oil with MIDAS and PEPS

Method	Monthly Average Prices			MIDAS			Period-End-Price Sampling		
Series	WTI	Brent	RAC	WTI	Brent	RAC	WTI	Brent	RAC
Horizon	MSFE Ratio								
1	1.00 (0.443)	1.00 (0.624)	1.00 (0.172)	0.59 (0.000)	0.59 (0.000)	0.68 (0.001)	0.59 (0.000)	0.58 (0.000)	0.69 (0.001)
3	1.00 (0.544)	1.01 (0.736)	1.00 (0.464)	0.88 (0.043)	0.92 (0.099)	0.87 (0.026)	0.88 (0.022)	0.92 (0.049)	0.88 (0.013)
6	0.99 (0.383)	1.01 (0.581)	0.99 (0.403)	0.93 (0.130)	0.98 (0.341)	0.93 (0.130)	0.93 (0.097)	0.98 (0.319)	0.93 (0.096)
12	0.94 (0.249)	0.97 (0.374)	0.96 (0.295)	0.90 (0.133)	0.96 (0.302)	0.91 (0.142)	0.91 (0.116)	0.96 (0.277)	0.92 (0.124)
24	0.90 (0.172)	0.92 (0.182)	0.91 (0.150)	0.89 (0.139)	0.91 (0.172)	0.89 (0.109)	0.89 (0.133)	0.91 (0.150)	0.89 (0.101)
	Success Ratio								
1	0.53 (0.265)	0.51 (0.494)	0.53 (0.392)	0.68 (0.000)	0.70 (0.000)	0.69 (0.000)	0.70 (0.000)	0.71 (0.000)	0.71 (0.000)
3	0.51 (0.575)	0.51 (0.552)	0.49 (0.740)	0.58 (0.010)	0.53 (0.288)	0.55 (0.111)	0.57 (0.031)	0.55 (0.089)	0.56 (0.062)
6	0.51 (0.558)	0.54 (0.235)	0.53 (0.318)	0.55 (0.188)	0.57 (0.049)	0.58 (0.039)	0.56 (0.105)	0.56 (0.074)	0.59 (0.014)
12	0.53 (0.343)	0.55 (0.178)	0.56 (0.147)	0.55 (0.209)	0.55 (0.131)	0.59 (0.032)	0.56 (0.161)	0.56 (0.108)	0.59 (0.037)
24	0.59 (0.068)	0.61 (0.033)	0.62 (0.027)	0.60 (0.038)	0.64 (0.007)	0.63 (0.017)	0.60 (0.051)	0.64 (0.007)	0.62 (0.019)

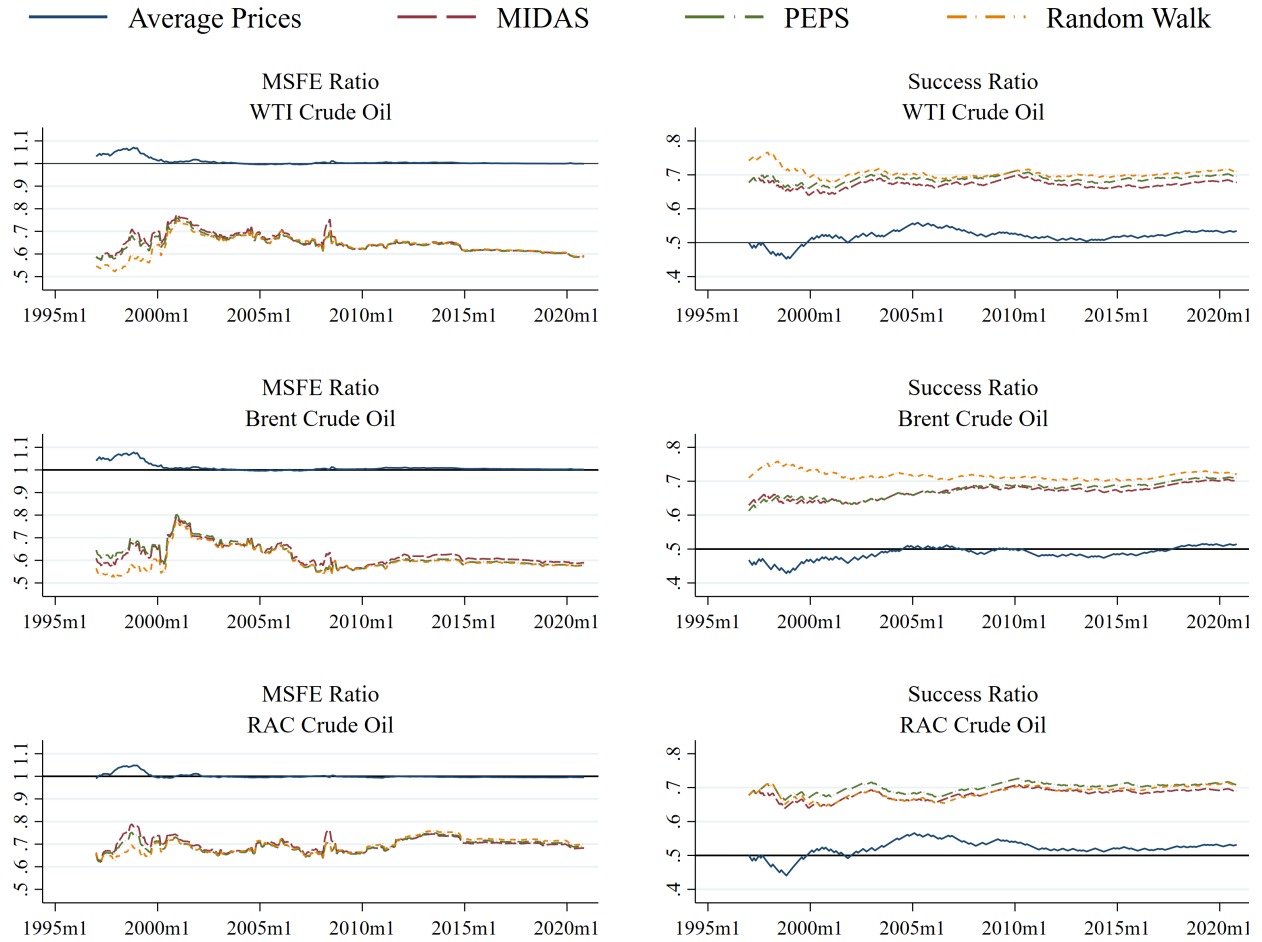
Note: Real-time, out-of-sample-forecasts of the real price of crude oil, 1992M1–2021M1. Brackets report the p-values for serial dependence robust tests of Pesaran and Pesaran and Timmermann (2009) for the null of no directional accuracy and Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the random walk forecast at the five percent significance level.

Table 3 presents the results for the MIDAS and PEPS forecasts alongside the direct forecasts constructed from the monthly average prices. It shows that the MIDAS forecasts (columns 4-6) and PEPS forecasts (columns 7-9) are very similar for all three oil-price series and across all forecast horizons. For all series and forecast horizons, these forecasts are at least as accurate as the direct forecasts based on average prices (columns 1-3), and in many cases outperform them by a substantial margin. However, these gains are very similar to those obtained from the random walk forecasts based on last observed daily prices. The bold values indicate that at the 5 percent significance level, the null hypothesis of no directional accuracy relative to the random walk forecast can only be rejected for the 1-month-ahead PEPS forecasts for RAC. At the 2-year horizon, all forecasts exhibit significant directional accuracy relative to the random walk for WTI and RAC. In terms of the MSPE, we find no evidence for statistical significant improvements over the random walk for any forecast.

The recursively updated MSPE and success ratios for the one-step-ahead forecast are shown in Figure 3. Across the entire sample, the forecast gains from using MIDAS or PEPS are very similar to those obtained from the random walk forecast from the last observed daily price. However, for WTI and RAC, the models estimated with PEPS slightly outperform in terms of directional accuracy compared to the MIDAS forecasts, especially towards the later half of the sample.

A key implication from the exercises presented in this section is that in practice, it makes little

Figure 2. Evolution of One-Month Ahead Forecast Evaluation Criteria for Direct Forecasts with MIDAS and PEPS



Note: One-month-ahead, real-time, out-of-sample forecasts of alternative crude oil price series, 1992M1–2021M1. MIDAS forecasts computed from Equation 5, PEPS forecasts computed from Equation 6, and direct forecasts using monthly “Average Prices” computed from Equation 7. “Random Walk” is the random walk forecast based on the last observed closing price. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

difference whether the forecaster uses the real monthly average prices or the real end-of-month price as the target variable in the estimation of the forecasting model. By contrast, using the information contained in disaggregated oil prices and in particular the last observed oil price in the forecaster’s information set is crucial to construct accurate forecasts.

3.4 Forecasts with Exogenous Predictors

This section examines the usefulness of disaggregated oil prices in direct forecasts that employ exogenous predictors. These models are widely used in the literature and a key component of

model averaging exercises (Baumeister et al., 2014; Wang et al., 2015; Baumeister and Kilian, 2015; Garratt et al., 2019). Following the literature, predictive regressions take on the form

$$\hat{R}_{t+h,a} = R_{t,a} \cdot (1 + \hat{\alpha}_h + \hat{\beta}_h X_t), \quad (8)$$

where X_t is the monthly predictor variable and $\hat{\alpha}_h$ and $\hat{\beta}_h$ are the least-squares estimates of a regression of the percent change in the monthly real price of oil on the predictor. In addition to this standard specification, we propose a version of the predictive regression where the monthly average is replaced with the end-of-month real price,

$$\hat{R}_{t+h,n} = R_{t,n} \cdot (1 + \hat{\alpha}_{n,h} + \hat{\beta}_{n,h} X_t). \quad (9)$$

Following the PEPS approach, forecasts for $R_{t+h,n}$ are then used to predict $R_{t+h,a}$.

The monthly predictor variable, X_t , examined for this application, is a portfolio of oil-sensitive stock prices, the NYSE Arca (AMEX) oil index. This index has been used to predict oil prices by Chen (2014) and Baumeister et al. (2015), among others. For all forecasts, parameters are estimated recursively using real-time data and forecasts are computed out-of-sample. We set $\alpha_h = 0$, which is common in the literature and provides better forecasts for all series, especially at the two-year horizon.⁸

The forecast performance of the two approaches is reported in Table 4. Columns 1-3 present the results for models estimated with monthly average oil prices. For all three oil series, the forecasts based on monthly average prices significantly outperform the average-price no-change forecast at the one-month horizon. However, the forecast performance deteriorates quickly with the forecast horizon and is often worse for forecasts beyond the 3-month-ahead prediction. Moreover, none of the forecasts obtained from the traditional forecasting model significantly outperforms the random walk forecasts of the last daily price.

Columns 4-6 present the results for models estimated with PEPS. They show that for all oil price series and forecast horizons, the forecasts computed with PEPS outperform the forecasts computed with monthly average oil prices. Forecasts via PEPS exhibit MSFE ratios as low as 0.54 and success ratios as high as 0.72. For Brent, the improvements in the MSFE and success ratio are also significant at the 5-percent significance level relative to the random walk of the last daily price

⁸The forecast results for horizons less than one year are unchanged in terms of the economic and statistical significance when the constant is included.

Table 4. Forecasts of the Real Price of Crude Oil Based on the NYSE Arca Oil Index

Method	Monthly Average Prices			Period-End-Price Sampling		
Price	WTI	Brent	RAC	WTI	Brent	RAC
Horizon	MSPE Ratio					
1	0.80 (0.006)	0.78 (0.004)	0.90 (0.071)	0.56 (0.000)	0.54 (0.000)	0.71 (0.005)
3	0.95 (0.150)	0.96 (0.156)	0.97 (0.314)	0.86 (0.035)	0.89 (0.047)	0.89 (0.074)
6	0.98 (0.059)	1.03 (0.751)	1.00 (0.459)	0.93 (0.017)	0.96 (0.065)	0.95 (0.040)
12	1.00 (0.415)	1.12 (0.934)	1.04 (0.937)	0.95 (0.049)	0.98 (0.192)	1.00 (0.440)
24	1.01 (0.571)	1.19 (0.911)	1.09 (0.996)	1.00 (0.499)	1.01 (0.582)	1.08 (0.991)
	Success Ratio					
1	0.61 (0.000)	0.62 (0.000)	0.61 (0.000)	0.72 (0.000)	0.72 (0.000)	0.70 (0.000)
3	0.49 (0.703)	0.54 (0.273)	0.51 (0.543)	0.60 (0.000)	0.56 (0.023)	0.58 (0.002)
6	0.55 (0.102)	0.56 (0.166)	0.48 (0.699)	0.55 (0.066)	0.58 (0.002)	0.55 (0.097)
12	0.53 (0.175)	0.52 (0.380)	0.51 (0.281)	0.60 (0.001)	0.57 (0.012)	0.56 (0.029)
24	0.44 (0.909)	0.51 (0.564)	0.38 (0.992)	0.51 (0.464)	0.54 (0.153)	0.47 (0.772)

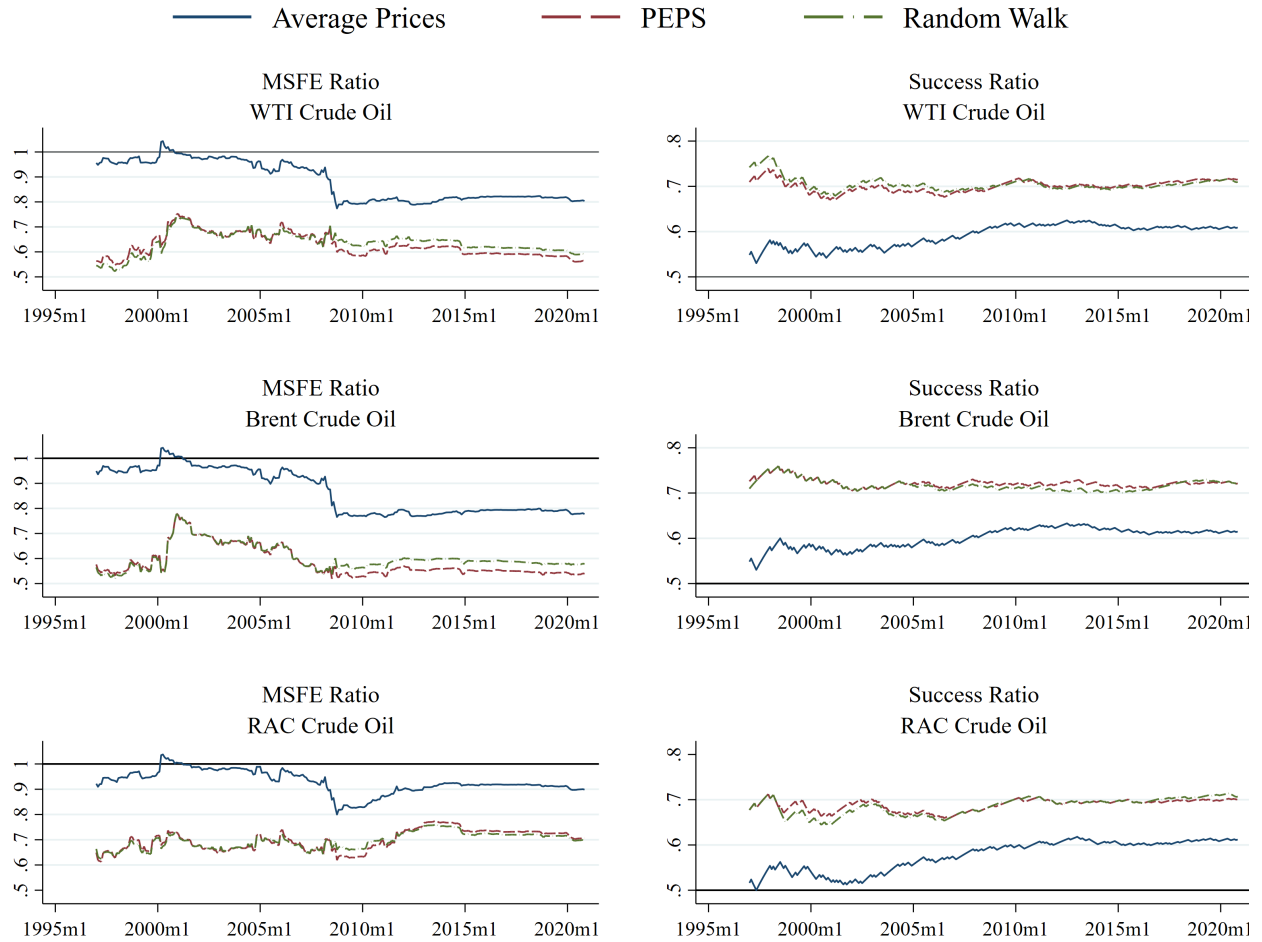
Note: Real-time, out-of-sample forecasts of the real price of crude oil, 1992M1–2021M1. Regression-based direct forecasts using the NYSE Arca oil index. Brackets report the p-values for serial dependence robust tests of Pesaran and Pesaran and Timmermann (2009) for the null of no directional accuracy and Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the random walk forecast at the five percent significance level.

at the one-month horizon. This finding is impressive, as no existing model-based forecast computed with monthly average oil prices outperforms the random walk forecast at this horizon (Ellwanger and Snudden, 2022). The MSFE gains are about 6 percent relative to the random walk, which is smaller than the 45% improvement implied by comparisons with the average-price no-change forecasts and more reminiscent of the short-horizon predictability documented for asset prices.

The recursively estimated MSFE and success ratios for the one-month-ahead forecasts are reported in Figure 3. It shows that models estimated with monthly average prices did particularly well during the 2000s and maintained gains relative to the conventional monthly average no-change forecast after 2008. However, for all ending points, the forecasts constructed with monthly average prices performed worse than the random-walk forecasts based on closing prices. Moreover, forecasts constructed with monthly average prices performed worse than models estimated with PEPS throughout the entire sample.

The regression estimates for Equations 8 and 9 provide useful intuition for the forecasting results presented in this section. Table 5 reports the coefficient estimates and robust standard errors of the predictive regression models for WTI, Brent and the RAC. For regressions based on monthly average prices, the estimates of β for all series are around 0.6 and the R-squared is about 15%. For regressions based on PEPS, the estimates for β drops to 0.3, and the R-squared falls to 3% or

Figure 3. Evolution of One-Month-Ahead Forecast Evaluation Criteria for Forecasts Based on the NYSE Arca Oil Index



Note: One-month-ahead, real-time, out-of-sample direct forecasts of the real prices of crude oil using the NYSE Arca oil index, 1992M1–2021M1. “Random Walk” is the random walk forecast based on the last observed closing price. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

below. This suggests that the majority of the forecasting gains from using the index in predictive regressions with average prices stem from recovering information that is already contained in the last available oil closing price and lost through averaging. Still, the results also imply the NYSE Arca oil index contains some additional predictive power that is not contained in the last observed closing price of oil.

3.5 Forecasts with Vector Autoregressive Models

A practical advantage of using PEPS for forecasts of the real price of crude oil is that it maintains the monthly frequency and thus can easily be extended to multivariate models. This section illustrates

Table 5. Coefficient Estimates for One-Month-Ahead Predictive Regressions Based on the NYSE Arca (AMEX) Oil Index

Sampling	Monthly Average Prices			End-of-Month Prices		
Dep. Variable	WTI	Brent	RAC	WTI	Brent	RAC
AMEX	0.593 (0.067)	0.614 (0.069)	0.548 (0.063)	0.309 (0.080)	0.313 (0.090)	0.257 (0.075)
Constant	0.0001 (0.004)	0.0004 (0.004)	-0.0002 (0.004)	0.003 (0.005)	0.005 (0.006)	0.003 (0.005)
N	449	449	449	449	449	449
R ²	0.148	0.149	0.142	0.030	0.024	0.024

Note: Coefficient estimates of $\hat{\alpha}_1, \hat{\alpha}_{1,n}$ and $\hat{\beta}_1, \hat{\beta}_{1,n}$ from Equations 8 and 9 for different crude oil price series, 1992M1–2021M1. HAC standard errors in parenthesis.

the point for forecasts with vector-autoregressive (VAR) models.

Our application examines the effect of introducing PEPS into the Kilian and Murphy (2014) 4-variable VAR model, which has been widely used in the literature. The model contains four variables: the real price of crude oil, the growth rate of global crude oil production, a proxy for the change in global crude oil inventories, and the real economics activity index (REA) as a measure of global economic activity. The original model has been refined in many ways in recent studies, for example, by using alternative economic activity indexes and inventory measures and by allowing for stochastic volatility (see, e.g., Funk, 2018; Snudden, 2018; Baumeister et al., 2020, among others). We focus on the original specification for illustrative purposes since it remains a benchmark for alternative VAR models.

The estimation follows the real-time forecasts of Baumeister and Kilian (2012). Accordingly, we estimate models with 12 autoregressive lags and with log real prices.⁹ Forecasts are converted to real price in levels. Other than the crude oil price series used for estimation, no other change is made, and all models are estimated at the monthly frequency.

Table 6 reports the forecast performance of the VAR models for the different series of the real price of crude oil with the updated forecast evaluation sample from 1992M1 to 2021M1. The models estimated with average prices only significantly improve upon the monthly average-price no-change forecast in terms of directional accuracy at horizons of 6 months and beyond. The deterioration in forecast accuracy for the MSFE at short horizons is consistent with evidence from other recent studies (Funk, 2018; Snudden, 2018; Baumeister et al., 2020). Figure 4 depicts the evolution of the forecast criteria over time and shows that deterioration in the MSFE ratios occurs mostly during

⁹Similar results are obtained for models with 24 lags, which are also frequently used in the literature.

Table 6. Forecasts of the Real Price of Crude Oil with VAR models

Estimation	Monthly Average Prices			Period-End-Price Sampling		
Horizon	WTI	RAC	Brent	WTI	RAC	Brent
MSPE Ratio						
1	1.00 (0.488)	0.92 (0.318)	1.05 (0.639)	0.73 (0.080)	0.75 (0.132)	0.84 (0.196)
3	0.99 (0.473)	0.98 (0.465)	1.08 (0.660)	0.91 (0.341)	0.91 (0.353)	1.07 (0.635)
6	1.04 (0.593)	1.09 (0.677)	1.15 (0.800)	0.99 (0.463)	1.04 (0.589)	1.16 (0.806)
12	1.14 (0.853)	1.17 (0.864)	1.22 (0.912)	1.11 (0.793)	1.14 (0.829)	1.24 (0.917)
24	1.03 (0.560)	1.00 (0.508)	1.05 (0.612)	1.03 (0.570)	1.02 (0.532)	1.07 (0.639)
Success Ratio						
1	0.51 (0.182)	0.56 (0.007)	0.51 (0.203)	0.69 (0.000)	0.68 (0.000)	0.66 (0.000)
3	0.55 (0.037)	0.55 (0.047)	0.54 (0.063)	0.58 (0.006)	0.58 (0.007)	0.56 (0.034)
6	0.59 (0.005)	0.57 (0.017)	0.58 (0.012)	0.61 (0.001)	0.59 (0.003)	0.60 (0.004)
12	0.59 (0.014)	0.63 (0.000)	0.57 (0.036)	0.61 (0.002)	0.61 (0.002)	0.60 (0.010)
24	0.60 (0.010)	0.59 (0.011)	0.59 (0.022)	0.59 (0.019)	0.60 (0.014)	0.59 (0.032)

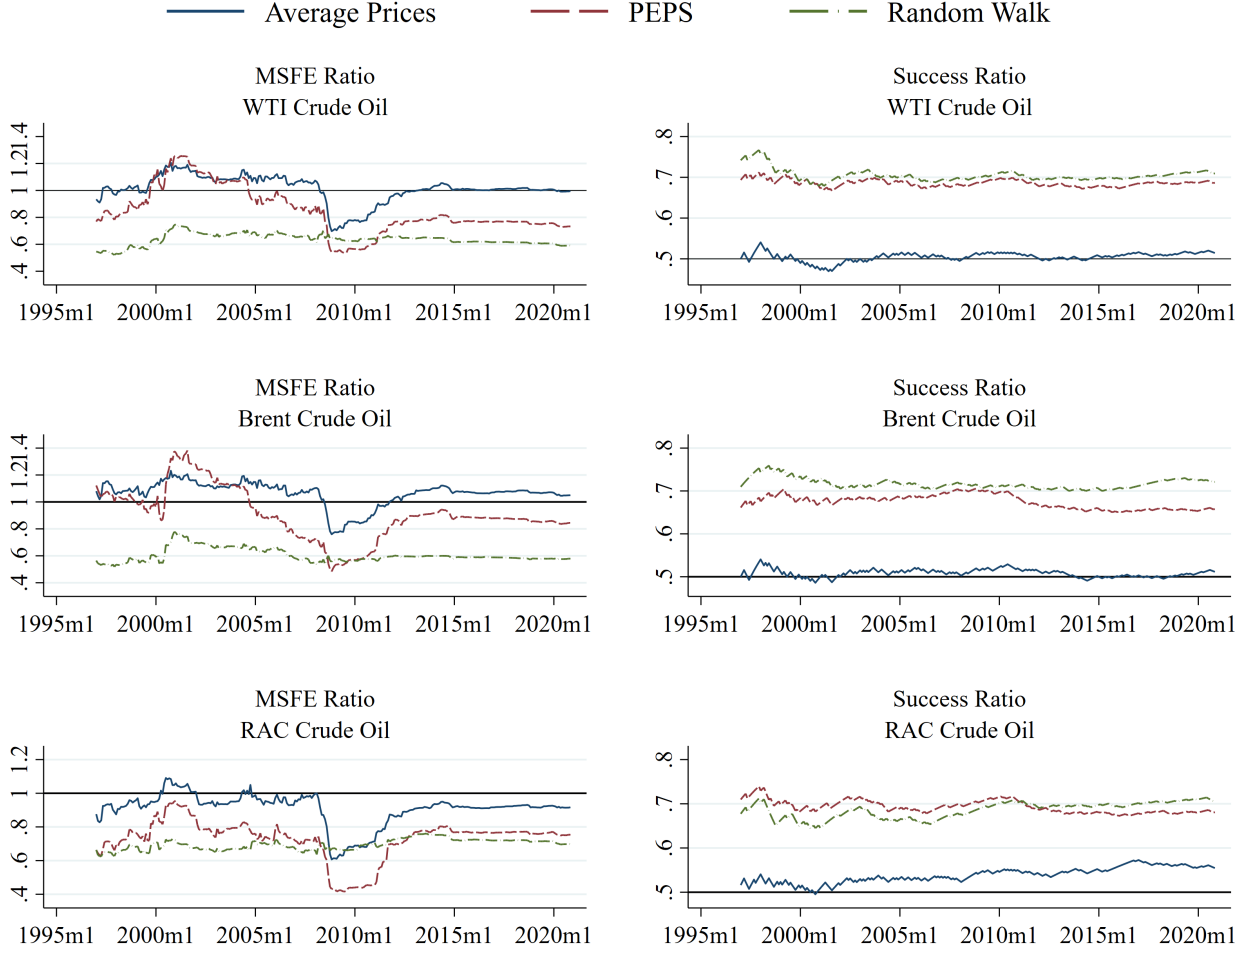
Note: Real-time, out-of-sample-forecasts of the real price of crude oil, 1992M1–2021M1. VAR models estimated with 12 lags. Brackets report the p-values for serial dependence robust tests of Pesaran and Pesaran and Timmermann (2009) for the null of no directional accuracy and Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the random walk forecast at the five percent significance level.

the years 2010 to 2014.

Table 6 shows that for most cases, VAR models estimated with PEPS outperform models estimated with monthly average prices. This is particularly true for forecasts horizons of up to 6 months. As shown in Figure 4, the improvements in forecast-accuracy at the one-step-ahead prediction forecast are of the order of 30 percent and, with a small exception in the early 2000s, very robust across the sample period. Even during the peak performance of the VAR model around 2008, forecast improvements as large as 40 percent could have been realized by using PEPS rather than estimating the VAR models with monthly average prices. The results for RAC show that estimating models with the imputed end-of-month observations instead of the average RAC yields forecast improvements of over 30 percent, which are robust for all sample end dates between 1995 and 2021. This highlights once more that even in the case where the end-of-periods observations are not available, approximating such observations with simple methods can produce large and robust forecast gains.

The results of this section show that information from end-of-month prices via PEPS can easily be incorporated into multivariate forecasting models and improves in particular short-horizon forecasts of such models. Undoubtedly, the VAR models presented in this section can be further refined. However, given the plethora of modifications proposed in the literature, a complete examination of

Figure 4. Evolution of One-Month-Ahead Forecast Evaluation Criteria for VAR Models



Note: One-month-ahead, real-time, out-of-sample forecasts of the real prices of crude oil, 1992M1–2021M1. VAR models estimated with 12 lags using either monthly average or end-of-month (PEPS) prices. “Random Walk” is the random walk forecast based on the last observed closing price. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

alternative VAR models is beyond the scope of this paper and left for future research.

4 Robustness

The results presented in this paper are remarkably robust to alternative modelling choices. For example, it can be shown that using ex-post revised data instead of real-time data does not affect any of our conclusions. Moreover, very similar forecasts are obtained by imputing end-of-month RAC prices from Brent rather than WTI prices.

Our results also hold for alternative methods of nowcasting the CPI and for forecasts exercises which using the nominal closing price instead of nominal monthly average prices. This robustness

is expected, as fluctuations in the CPI deflator are generally small compared to the fluctuations of nominal oil prices (Baumeister and Kilian, 2012; Alquist et al., 2013).

Finally, it can be shown that model-based forecast using disaggregated approaches are also superior to model estimated with quarterly and annual average price, which are of primary interest to policymakers (Baumeister and Kilian, 2014, 2015). Consistent with theoretical results, the forecast improvements are even larger at the quarterly and annual frequency compared to the monthly frequency.

5 Conclusion

The standard approach in the literature is to compute model-based forecasts of the real price of crude oil from monthly average prices. We have shown how approaches that exploit information from the underlying daily oil prices can be applied to existing models. Forecasts relying on end-of-month prices yield forecast improvements over the standard no-change forecast of up to 45% at the one-step-ahead prediction, which is unprecedented in this literature. The magnitude of the gains decrease with the forecast horizon, yet models estimated with daily prices still improve upon traditional forecasts as far ahead as the 24-month horizon.

We also directly tested, for the first time, whether model-based forecasts of the monthly real price of crude oil outperform the random walk forecast based on monthly closing prices. We find that none of the forecasts estimated with monthly average prices outperform the random walk at horizons lower than 6 months. By contrast, some approaches relying on daily prices displayed significantly lower MSFE ratios or improved directional accuracy relative to the random walk forecast. In economic terms, the forecast gains relative to the random walk are reminiscent of the short-horizon predictability documented for asset prices.

We have proposed practical solutions for how to incorporate information contained in daily oil prices into several types of model-based forecasts that have previously only been estimated with average oil prices. This includes intuitive methods for imputing end-of-month observations for the RAC and for backcasting end-of-month prices to estimate models with longer historical data. Undoubtedly, the refinement of the models and methods considered in this paper opens up a promising avenue for further research. The fact that even simple applications of disaggregated approaches provide large improvements over models estimated with average oil prices only strengthens our insight that forecasts of the real price of oil should exploit the information contained in daily oil

prices.

References

- Alquist, R., Bhattarai, S., and Coibion, O. (2020). Commodity-price comovement and global economic activity. *Journal of Monetary Economics*, 112:41–56.
- 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.
- Andreou, E., Ghysels, E., and Kourtellis, A. (2010). Regression models with mixed sampling frequencies. *Journal of Econometrics*, 158(2):246–261.
- Baumeister, C., Guérin, P., and Kilian, L. (2015). Do high-frequency financial data help forecast oil prices? The midas touch at work. *International Journal of Forecasting*, 31(2):238–252.
- Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.
- 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., Kilian, L., and Lee, T. K. (2014). Are there gains from pooling real-time oil price forecasts? *Energy Economics*, 46:S33–S43.
- Baumeister, C., Korobilis, D., and Lee, T. K. (2020). Energy markets and global economic conditions. *NBER Working Paper No. 27001*.
- Baumeister, C., Korobilis, D., and Lee, T. K. (2022). Energy markets and global economic conditions. *The Review of Economics and Statistics*, 104(4).

- Bork, L., Kaltwasser, P. R., and Sercu, P. (2022). Aggregation bias in tests of the commodity currency hypothesis. *Journal of Banking & Finance*, 135:106392.
- Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time series analysis: forecasting and control*. John Wiley & Sons.
- Chen, S.-S. (2014). Forecasting crude oil price movements with oil-sensitive stocks. *Economic Inquiry*, 52(2):830–844.
- Conlon, T., Cotter, J., and Eyiah-Donkor, E. (2022). The illusion of oil return predictability: The choice of data matters! *Journal of Banking & Finance*, 134:106331.
- Degiannakis, S. and Filis, G. (2018). Forecasting oil prices: High-frequency financial data are indeed useful. *Energy Economics*, 76:388–402.
- 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*.
- Ellwanger, R. and Snudden, S. (2022). Forecasts of the real price of oil revisited: Do they beat the random walk? *LCERPA Working Paper No. 2022-X*.
- Ellwanger, R. and Snudden, S. (2023). Futures prices are useful predictors of the spot price of crude oil. *The Energy Journal*, 44(4):45–62.
- 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.
- Ghysels, E., Sinko, A., and Valkanov, R. (2007). Midas regressions: Further results and new directions. *Econometric reviews*, 26(1):53–90.
- 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.

- Kilian, L. (2015). Comment on Francis X. Diebold’s “Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of Diebold–Mariano tests”. *Journal of Business and Economic Statistics*, 33(1):13–17.
- Kilian, L. (2019). Measuring global real economic activity: Do recent critiques hold up to scrutiny? *Economics Letters*, 178:106–110.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.
- Kohn, R. (1982). When is an aggregate of a time series efficiently forecast by its past? *Journal of Econometrics*, 18(3):337–349.
- Lütkepohl, H. (1984). Forecasting contemporaneously aggregated vector ARMA processes. *Journal of Business & Economic Statistics*, 2(3):201–214.
- Lütkepohl, H. (1986). Forecasting temporally aggregated vector ARMA processes. *Journal of Forecasting*, 5(2):85–95.
- Marcellino, M. (1999). Some consequences of temporal aggregation in empirical analysis. *Journal of Business & Economic Statistics*, 17(1):129–136.
- Pesaran, M. H. and Timmermann, A. (2009). Testing dependence among serially correlated multicategory variables. *Journal of the American Statistical Association*, 104(485):325–337.
- 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.
- Wei, W. W. (1978). Some consequences of temporal aggregation in seasonal time series models. In *Seasonal analysis of economic time series*, pages 433–448. NBER.
- Working, H. (1960). Note on the correlation of first differences of averages in a random chain. *Econometrica*, 28(4):916–918.
- 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.
- Zhang, Y.-J. and Wang, J.-L. (2019). Do high-frequency stock market data help forecast crude oil prices? Evidence from the midas models. *Energy Economics*, 78:192–201.

Online Appendix (Not for publication)

A1 Real-Time Data

This section describes the construction of the real-time data used in the empirical exercises.

Crude oil prices: Crude-oil price data for WTI and Brent are obtained from the [EIA Petroleum and other liquids](#). Daily closing prices for WTI and Brent are used to calculate both the monthly average and end-of-month prices. This data is not subject to revisions and is observed in real time. As reported by the EIA and standard in applied work, the monthly average price is the simple average of daily closing prices. The end-of-month price is the closing price on the last trading day of the month. The daily series of WTI and Brent prices begin in 1986M1 and 1987M5, respectively. For models estimated at the daily frequency, the negative WTI price on April 20, 2020, is replaced with an imputed price for that day using the daily growth rate of Brent prices. The monthly WTI and Brent price series are backcasted to 1983M4 using the growth rate in the monthly price of the futures front contract. Monthly prices before 1983M4 are backcasted using the growth rate of the RAC (imported). None of these backcasting choices are crucial for our results, as our main findings are robust to estimating models with data starting in 1983 or 1986.

Table A1. Descriptive Statistics for the Real Prices of Crude Oil

Monthly Average	Date Range	Mean	Std. Dev.	Min	Max
U.S. Refiner Acquisition Cost, Imported	1973m1 - 2021m1	24.03	12.31	5.95	60.85
Brent	1973m1 - 2021m1	23.50	10.67	6.08	61.56
West Texas Intermediate	1973m1 - 2021m1	22.63	11.09	5.71	58.34
End of Month	Date Range	Mean	Std. Dev.	Min	Max
U.S. Refiner Acquisition Cost, Imported	1973m1 - 2021m1	24.07	12.37	5.76	63.64
Brent	1973m1 - 2021m1	22.27	10.23	5.02	64.36
West Texas Intermediate	1973m1 - 2021m1	21.39	10.51	5.13	60.68

Note: All series are the level of real prices from 1973M1–2021M5. 2021M5 data vintage. The end-of-month price for the U.S. refiner acquisition cost (imported) is imputed using the WTI price. Nominal crude-oil price data obtained from the Energy Information Administration and the consumer price index is obtained from the Federal Reserve of Philadelphia.

Oil market variables: The real-time data for the nominal U.S. refiner acquisition cost for crude-oil imports, world crude-oil production, U.S. crude-oil inventories, U.S. petroleum inventories, and OECD petroleum inventories are obtained from the EIA. Real-time vintages start in 1973M1 and use the vintages of Baumeister and Kilian (2012) from 1991M12 to 2010M12. Vintages from 2011M01 onwards are collected using the same data sources. This real-time database is updated monthly and is publicly available, along with further documentation on [Stephen Snudden’s website](#). The

updated real-time data is similar to Garratt et al. (2019) but differs primarily in that all historical releases, and thus data revisions, are included.

Other economic variables: For the construction of the Real Economic Activity index we use Kilian (2009)’s freight-rate data until 1984M12 and, thereafter, monthly values of the Baltic Dry Index obtained from Bloomberg. The index is computed using the corrected formula (Kilian, 2019). As in Baumeister and Kilian (2012), the real-time data for the monthly seasonally adjusted U.S. consumer price index for all urban consumers is obtained from the Economic Indicators published by the Council of Economic Advisers from the FRASER database of the Federal Reserve Bank of St. Louis and from the macroeconomic real-time database of the Federal Reserve Bank of Philadelphia. The AMEX Oil Index is obtained from Yahoo Finance (XOI) and is available in real time.

Nowcasting

The CPI series, the global crude-oil market variables used in the VARs, and the refiners import price of crude oil are subject to historical revisions and are reported with a lag. For each vintage, observations that are missing due to reporting lags are nowcasted following Baumeister and Kilian (2012):

- Missing observations for the U.S. CPI, and U.S. crude-oil inventories are nowcasted by extending the series by using the average of the historical growth rate at the respective point in time.
- Missing observations for the ratio of OECD petroleum inventories to U.S. petroleum inventories are kept constant at the last available value for this ratio.
- Monthly nominal U.S. crude oil imported acquisition cost by refiners are extrapolated with the growth rate of the monthly average of the nominal WTI price.