LCERPA
Laurier Centre for Economic Research & Policy Analysis

LCERPA Working paper No. 2023-5

December 2023

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

Amor Aniss Benmoussa
Bank of Canada

Reinhard Ellwanger
Bank of Canada

Stephen Snudden
Wilfrid Laurier University

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

Amor Aniss Benmoussa† Reinhard Ellwanger ‡ Stephen Snudden§¶

December 8, 2023

Abstract

This paper proposes techniques to include information from the underlying nominal daily series in model-based forecasts of average real series. We apply these approaches to forecasts of the real price of crude oil. Models utilizing information from daily prices yield large forecast improvements and, in some cases, almost halve the forecast error compared to current specifications. We demonstrate for the first time that model-based forecasts of the real price of crude oil can outperform the traditional random walk forecast, which is the end-of-month no-change forecast, at short forecast horizons.

JEL classification: C18, C53, Q47

Keywords: Forecasting and Prediction Methods, Temporal Aggregation, 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

‡Bank of Canada, International Economic Analysis Department

§Corresponding author; Wilfrid Laurier University, Department of Economics, 64 University Ave W, Waterloo, ON N2L 3C7; ORCID iD 0000-0002-4033-3235; Email: ssnudden@wlu.ca

¶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 This paper supersedes any related findings in "The New Benchmark for Forecasts of the Real Price of Crude Oil," an unpublished Bank of Canada Working Paper.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

Macroeconomic forecasters often target real series that are aggregated from higher-frequency nominal data. A case in point is the real price of crude oil, which plays a key role in macroeconomic projections and has attracted considerable attention among forecasters (see, e.g., Baumeister and Kilian, 2012; Alquist et al., 2013, and the survey of Filippidis et al., 2024). Notably, all existing model-based forecasts of the monthly real price of crude oil rely on monthly averages of daily oil prices, which facilitates incorporating monthly predictor variables and constructing historical price series (Alquist et al., 2013).[1] However, this form of temporal aggregation could also introduce information loss that diminishes forecasting accuracy (see, e.g., Wei, 1978; Kohn, 1982; Lütkepohl, 1986; Marcellino, 1999). This issue has recently been highlighted by Ellwanger and Snudden (2023a), who demonstrate that the traditional random-walk forecast, which is the end-of-month no-change forecast, outperforms all previous model-based forecasts of the real oil price at the one-month forecast horizon.

The goal of this paper is to investigate how daily information can be used to improve model-based forecasts of period average real series. To do so, we extend existing forecasting approaches dealing with higher-frequency information to a setting where nominal daily data are converted to real monthly average series. We henceforth refer to these approaches as *disaggregated* approaches, and apply them to popular forecasting models of the real price of crude oil.

Our main contribution is to compare the effectiveness of disaggregated forecast approaches to current forecast approaches that rely on monthly average real data. The proposed extensions provide practical solutions on how to efficiently construct forecasts when the temporal aggregate is expressed in real terms. The analysis quantifies the information loss from daily to monthly aggregation and provides a unique perspective on the forecast gains from temporal disaggregation in practice, which have thus far primarily focused on monthly to quarterly or quarterly to annual aggregation.[2]

Our second contribution is to directly compare, for the first time, the effectiveness of model-based forecasts of the monthly average real price of crude oil relative to the traditional random walk forecast, which is the end-of-month no-change forecast. In contrast, prior studies have compared

[1] For example, structural models and forecasting approaches using data prior to 1983 typically rely on historical refinery acquisition costs of crude oil, which are only available at the monthly frequency (Kilian, 2009; Alquist et al., 2013).

[2] See, for example, Zellner and Montmarquette (1971); Abraham (1982); Lütkepohl (1986); Athanasopoulos et al. (2011).

model-based forecasts against a monthly average no-change forecast, which can result in spurious predictability (Ellwanger and Snudden, 2023a). Importantly, this marks the first instance of testing forecasts of the monthly average real oil price against the null hypothesis of no-predictability.

Our empirical exercises concentrate on the most common models in oil-price forecasting, which include recursive forecasts from Autoregressive Moving Average (ARMA) models, Vector Autoregressive models (VARs), and direct forecasts from univariate and multivariate models. For each modeling class, we propose practical techniques to integrate disaggregated approaches for period average real forecasts. Together, these exercises provide a comprehensive comparison of the information loss from temporal aggregation across forecast approaches.

The first disaggregated approach that we consider is the bottom-up approach (BU). This approach consists of forecasting the underlying high-frequency series and subsequently averaging the forecasts to derive lower-frequency predictions. Theoretical results for ARMA processes show that the BU approach is preferable to models estimated with average data (Tiao, 1972; Kohn, 1982), even under parameter uncertainty (Lütkepohl, 1986). However, it remains an open question how this approach performs in practice for aggregation from daily to monthly variables and for the real price of crude oil in particular. In implementing the BU approach, we propose solutions for how to deal with the aggregation from nominal to real variables, which is a setting frequently encountered in macroeconomic forecasting.

The second disaggregated method is Period-End Point Sampling (PEPS) introduced by Ellwanger and Snudden (2023c). PEPS consists of estimating models with end-of-period data and subsequently using point forecasts as predictions for period averages. By employing end-of-period sampling during estimation, PEPS mitigates information loss caused by temporal aggregation. Moreover, this approach allows models to be estimated at the same frequency as the forecast target and is easy to implement in existing lower-frequency forecasting frameworks.

The third disaggregated method is MIxed-DAta Sampling (MIDAS), which has become a popular approach to deal with higher-frequency predictor variables (see Ghysels et al., 2007; Andreou et al., 2010). Unlike previous MIDAS applications in crude oil price forecasting, which utilize high-frequency financial variables as predictors of the monthly real price of oil (Baumeister et al., 2015; Degiannakis and Filis, 2018; Zhang and Wang, 2019), we employ this framework to forecast monthly oil prices with the underlying daily prices. To our knowledge, this is the first time the MIDAS approach is used for univariate forecasts of lower-frequency data.

Our main finding is that regardless of the approach employed, forecasting models that utilize information from daily prices substantially outperform conventional models estimated with monthly averages, especially at short forecast horizons. The disaggregated methods offer robust improvements in forecast accuracy and can reduce the mean squared forecast error (MSFE) by over 40% at the one-month horizon. Furthermore, models employing daily oil prices can provide higher directional accuracy than traditional forecasting models for forecast horizons up to 24 months ahead. These forecast gains are commensurate with those suggested in theory (Tiao, 1972; Kohn, 1982; Lütkepohl, 1986), and substantially larger than those documented in practice in prior companions of monthly to quarterly or quarterly to annual aggregation (see, e.g., Zellner and Montmarquette, 1971; Lütkepohl, 1986; Athanasopoulos et al., 2011).

Our replications find that existing model-based forecasts approaches only outperform the end-of-month no-change at long horizons, and typically only for directional accuracy. In contrast, we demonstrate for the first time that forecasting models of the real price of crude oil can outperform the end-of-month no-change forecast at short forecast horizons when disaggregated techniques are utilized. However, in such cases, the forecast gains are smaller than implied by comparisons with the monthly average no-change forecast. These findings reinforce that the real price of crude oil is more similar to asset prices than previously concluded.

An insight from our exercises is that much of the benefit of using daily oil prices stems from the information contained in the last daily price in the forecaster's information set. 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 can provide effective forecasts, as long as they incorporate information from end-of-month prices rather than monthly average prices. Although this result is consistent with the well-known fact that the last available instantaneous price contains crucial information about expectations (see, e.g., Fama, 1970), this idea is rarely exploited in forecasting applications of average data.

In examining the effectiveness of different forecasting approaches in a setting of temporally aggregated variables, our study also contributes to the broader debate about direct versus recursive forecasting approaches. Existing studies have compared direct and recursive forecast approaches within a single frequency, suggesting that recursive forecasts often outperform direct forecasts (see, e.g., Ing, 2003; Marcellino et al., 2006; Chevillon, 2007). Our findings reveal that while recursive forecasts typically outperform direct forecasts when estimated with monthly average data, they perform similarly when estimated with temporally disaggregated data.

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 instance, Ellwanger and Snudden (2023b) demonstrate that using end-of-period oil futures prices rather than period-average oil futures prices leads to significant improvements in futures-based forecasts of the real oil price. Baumeister and Kilian (2014) show that aggregating forecasts of monthly average oil prices to quarterly averages is preferable to computing forecasts with quarterly models, while Baumeister and Kilian (2012) shows that forecasting gains from data revisions primarily arise from improved information on crude oil prices. Our study expands on these findings by showing that integrating information from daily prices into model-based forecasts of the monthly real price of crude oil can substantially enhance their accuracy compared to original specifications.

2 Method

The focus of our study is real-time forecasts of the level of the monthly average real price of crude oil, which is the standard approach in the literature (see, e.g., Baumeister and Kilian, 2012; Alquist et al., 2013, and the surveys of Ellwanger and Snudden 2023a; Filippidis et al., 2024).

For Brent and WTI, the average monthly nominal price is the simple average of the daily closing prices, $\bar{S}_t = \frac{1}{n}\sum_{i=1}^{n} S_{t,i}$, where $S_{t,i}$ stands for the nominal spot price at the closing of day i in month t, and where n is the number of business-calendar days within a month.[3] Nominal prices are converted to real prices using the monthly seasonally adjusted U.S. consumer price index, denoted by CPI_t. The conventional monthly average real price of crude oil can be expressed as the monthly average of daily real prices, $R_{t,i}$:

$$\bar{R}_t = \frac{\frac{1}{n}\sum_{i=1}^{n} S_{t,i}}{CPI_t} = \frac{1}{n}\sum_{i=1}^{n}\frac{S_{t,i}}{CPI_t} = \frac{1}{n}\sum_{i=1}^{n} R_{t,i}.$$

Following the convention in the literature, we assume that the forecaster uses the available information at the end of month t to form their prediction of the monthly average real price in the h months ahead. The standard approach for forecasts of the real price of crude oil is to ignore daily oil prices and to forecast the h-step-ahead monthly average price, $\bar{R}_{t+h}$, using the history of monthly average oil prices up to $\bar{R}_t$. However, even when the goal is to forecast the monthly

[3] A different oil price series that has been used widely in forecasting applications is the monthly US refiners' acquisition costs of crude oil. We discuss the treatment of this series in Section 2.2.

average price, incorporating the information contained in higher-frequency price observations could be beneficial for the accuracy of forecasts.

For example, if the underlying data are generated by an ARMA process, relying on temporally aggregated data will result in a loss of forecast efficiency that increases with the persistence of the underlying data and the number of higher-frequency observations that are aggregated (Tiao, 1972; Kohn, 1982). Moreover, Lütkepohl (1986) demonstrated that utilizing the disaggregated data results in superior forecasts even under parameter uncertainty.[4] Based on these insights, this study examines for the first time how the series of underlying daily oil prices up to $R_{t,n}$ can be used to improve forecasts of the monthly average real price of crude oil.

2.1 Forecast Evaluation

Throughout the paper, we follow existing practice in the literature and report MSFE ratios relative to the conventional monthly average no-change forecast (see, e.g., the survey of Filippidis et al., 2024). The MSFE ratio for the h-steps-ahead forecast, $MSFE_h^{ratio}$, is calculated as the ratio of the MSFE of the model-based forecast to the MSFE of the monthly average no-change forecast:

$$MSFE_h^{ratio} = \frac{\sum_{q=1}^{Q}(\bar{R}_{q+h} - \hat{\bar{R}}_{q+h|q})^2}{\sum_{q=1}^{Q}(\bar{R}_{q+h} - \bar{R}_{q|q})^2}, \tag{1}$$

where $\hat{\bar{R}}_{q+h|q}$ is the conditional model forecast for the h-steps-ahead aggregated observation $\bar{R}_{q+h}$ for all periods of the evaluation sample, $q = 1, 2, \ldots, Q$, and $\bar{R}_{q|q}$ is the real-time monthly average no-change forecast. We also report the p-values for the test of equal predictive accuracy relative to the monthly average no-change forecasts following Diebold and Mariano (1995).[5]

The mean directional accuracy is reported using the success ratio, which indicates the proportion of times that the model-based forecasts correctly predict the direction of the change in the monthly average real price. Consistent with previous real-time exercises, the success ratio for forecast horizon h, SR_h, is:

$$SR_h = \frac{1}{Q}\sum_{q=1}^{Q} \mathbb{1}[sgn(\bar{R}_{q+h} - \bar{R}_q) = sgn(\hat{\bar{R}}_{q+h|q} - \bar{R}_{q|q})], \tag{2}$$

where $\mathbb{1}[\cdot]$ is an indicator function and $sgn(\cdot)$ is the sign function. The null hypothesis of no

[4] In contrast, Lütkepohl (1984) shows that for contemporaneous aggregation (summation over subcomponents), the gains of using aggregated versus disaggregated data are unclear under parameter uncertainty.

[5] The use of real-time data and the iterative out-of-sample forecasts does not fulfill the assumptions underlying the Diebold-Mariano test (Diebold, 2015). As in extant studies, the p-values are still reported with this caveat in mind.

directional accuracy (corresponding to a success ratio equal to one half) is tested using Pesaran and Timmermann (2009).

In addition to the standard practice of comparing forecasts against the monthly average no-change forecast, we compare, for the first time, model-based forecasts against the traditional random walk forecast, which is the end-of-month no-change forecast. The end-of-month no-change forecast is the only no-change forecast reflecting the null hypothesis that *all* future prices – both daily prices and average prices — are conditionally unpredictable and therefore corresponds to the traditional random walk forecast used in economics and finance. In contrast, forecast improvements relative to the monthly average no-change forecast are theoretically expected for all (vector) autoregressive integrated moving average representations of the daily data (Weiss, 1984; Marcellino, 1999). Moreover, Ellwanger and Snudden (2023a) shows that for the real price of crude oil, the end-of-month no-change forecast is significantly more accurate than the monthly average no-change forecasts for up to one year ahead. As such, comparisons with the end-of-month no-change forecast are necessary to provide evidence for the practical usefulness of specific model-based forecast approaches and for the predictability of the monthly average real price of crude oil.

The series of real end-of-period observations, $R_{t,n}$, is constructed using $R_{t,n} = S_{t,n}/CPI_t$, and the real-time equivalent $R_{t,n|t} = S_{t,n}/CPI_{t|t}$. The measure of real end-of-month prices is consistent with the common practice in other forecast applications, such as bilateral exchange rates (Meese and Rogoff, 1983) and primary commodity prices (West and Wong, 2014). Forecast improvements relative to the end-of-month no-change forecast are tested by replacing the conventional monthly average no-change forecast in the tests of Pesaran and Timmermann (2009) and Diebold and Mariano (1995). Throughout all tables, significant improvements over the end-of-month no-change forecast at the five percent significance level are indicated in bold.[6]

2.2 Disaggregated Data

We study model-based forecasts of the three oil-price benchmarks used in the literature: West Texas Intermediate (WTI), Brent, and US refiner acquisition cost of crude oil, imported (RAC). Crude oil prices are obtained from the U.S. Energy Information Administration (EIA). The monthly average and daily prices for WTI and Brent are available in real time. The price deflator is the seasonally adjusted U.S. consumer price index obtained from the FRASER database of the Federal Reserve

[6] Tables reporting MSFE and success ratios relative to the end-of-month no-change forecast instead of the conventional monthly average no-change forecasts are available upon request.

Bank of St. Louis and the real-time database of the Philadelphia Federal Reserve.

Unlike WTI and Brent, the RAC is a survey of the monthly average acquisition costs and is only observed at the monthly frequency. 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 daily RAC observations by applying the growth rate of the daily WTI oil price over the monthly average WTI oil price to the monthly RAC price: $S_{t,i}^{RAC} = \bar{S}_t^{RAC} \cdot (S_{t,i}^{WTI}/\bar{S}_t^{WTI})$, where $S_{t,i}$ refers to the nominal closing price on the day i and the superscripts RAC and WTI denote the nominal RAC series and the nominal WTI prices, respectively.[7] The idea of imputing RAC observations with WTI prices has ample precedent in the oil-price forecasting literature. For example, Baumeister and Kilian (2012) nowcast the average nominal RAC with the growth rate of the WTI price and use WTI futures prices to construct futures-based forecasts for RAC. The tight empirical correlation between WTI prices and the RAC implies that these nowcasts and imputations yield good results in practice.

Our exercises also replicate and extend forecasts from the Kilian and Murphy (2014) VAR model, which includes 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 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 that this is the first update to include all revisions for the entire history of each vintage.[8] Real-time data vintages start in 1991M12 and contain historical data from 1973M1 onward. 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. The results are qualitatively robust to alternative methods of nowcasting the CPI and the

[7]Very similar forecasts are obtained by imputing end-of-month RAC prices from Brent rather than WTI prices.

[8]The data appendix contains detailed descriptions of the real-time vintages, nowcasts, and data sources. This real-time database is updated monthly and available to download along with further documentation.

use of ex-post revised data.[9]

The series of daily and monthly WTI and Brent prices provided by the EIA begin in 1986M1 and 1987M5, respectively, and the estimation of models at the daily frequency begins on those dates. For forecasting models estimated at the monthly frequency, including models estimated via PEPS, prices are backcasted to 1973M1. Specifically, the 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 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 average observation. None of these backcasting choices are crucial for our results, and our main findings are robust to estimating models with data starting in 1983 or 1986.

3 Model-Based Forecasts

We focus our analysis on the most common types of model-based forecasts in the literature (see, e.g., the survey of Filippidis et al., 2024). Specifically, we examine disaggregated forecasting approaches for recursive forecasts from univariate ARMA models and from VARs. We also examine direct univariate forecasts from MIxed-DAta Sampling (MIDAS) and Linear Regression Models (LRM) with exogenous predictor variables. The selection of models comprises univariate and multivariate methods for direct and recursive forecasts. Hence, our exercises provide a comprehensive overview of the usefulness of disaggregated approaches in different settings encountered in applied forecasting. All forecasts are evaluated out-of-sample with an expanding window beginning in 1992M1and the evaluation window ends in 2021M1.

3.1 ARMA Forecasts

The most common model-based forecasts for crude oil prices are recursive forecasts from ARMA models. These univariate time-series models commonly serve as an alternative benchmark in forecast exercises and tend to be more accurate predictors of the real price of crude oil than the conventional average-price no-change forecast (Baumeister and Kilian, 2012; Funk, 2018; Snudden,

[9] This is consistent with the findings of (Baumeister and Kilian, 2012; Ellwanger and Snudden, 2023b), which show that fluctuations in the CPI deflator are generally small compared to fluctuations in nominal oil prices and tend to have minimal impact on forecasts.

2018).

The convention in the literature is to estimate ARMA models for real oil prices in log levels. Let the log level of the real price of oil be $\bar{r}_t = \ln(\bar{R}_t)$, where ln is the natural log function. The ARMA(p, q) with p autoregressive parameters, q moving-average coefficients, and innovations $\check{\epsilon}_t$ is given by:

$$\check{a}(L)\bar{r}_t \quad = \quad \check{c} + \check{b}(L)\check{\epsilon}_t \quad \forall\ t. \tag{3}$$

where $\check{c}$ is a constant and L is the lag operator such that $Ly_t = y_{t-1}$, $\check{b}(L) = (1 + \check{\alpha}_1 L + \cdots + \check{\alpha}_q L^q)$, and $a(L) = (1 - \check{\rho}_1 L - \cdots - \check{\rho}_p L^p)$. The estimated parameters are used to construct recursive model-based forecasts, $\check{\bar{r}}_{t+h|t}$, which are converted back into real prices in levels, $\check{\bar{R}}_{t+h|t} = \exp\left(\check{\bar{r}}_{t+h|t}\right)$, where exp is the exponential function.

One advantage of using monthly average data is that the prices of WTI and Brent, which are only available since the early 1980s, can be backcasted to 1973 using the RAC.[10] This is useful because the backcasting of all series back to 1973 provides a better estimate of the long-run mean, which could improve the forecast accuracy at longer horizons. However, this advantage has to be weighed against the efficiency loss from using aggregated data. Therefore, the empirical question of which method is more effective for forecasts of the real price of crude oil remains.

The first disaggregated approach that we consider for ARMA models is the BU approach. Empirical evidence for the BU approach suggests that the approach improves forecast efficiency in some, but not all cases (Zellner and Montmarquette, 1971; Abraham, 1982; Lütkepohl, 1986; Athanasopoulos et al., 2011). However, existing studies have only considered aggregation from monthly to quarterly frequency or from quarterly to annual frequency. To the best of our knowledge, this paper is the first to consider aggregation from daily to monthly variables. The information loss generally increases when aggregation occurs over more observations (Tiao, 1972). Moreover, most of the information loss occurs from aggregation over the first few observations. The daily to monthly aggregation moves from no aggregation, $n = 1$, to a substantive aggregation $n = 21$. In contrast, monthly to quarterly aggregation starts from data already aggregated, $n = 21$, to a higher degree of aggregation, $n = 63$. Thus, the information loss from daily to month aggregation is expected to be larger than monthly to quarterly aggregation.

One potential complication with constructing BU forecasts for monthly average real data is

[10]The RAC series itself is imputed for the year 1973 following Mork (1989).

that CPI is only available at the monthly frequency. This may explain why the BU approach has been overlooked in applications to real macroeconomic variables that are aggregated from a daily frequency (see, e.g., Box et al., 2015). For the baseline results, we construct forecasts of the daily nominal price of crude oil. We then average the daily forecasts to the monthly frequency and deflate the monthly average forecast using the CPI forecast. This approach mirrors the current practice of deflating forecasts of nominal monthly average oil prices into real terms using CPI forecasts in futures-based forecasts (Baumeister and Kilian, 2012; Ellwanger and Snudden, 2023b).

Since the nominal daily price is non-stationary, we estimate the model using the growth rate of nominal daily prices, $g_{t,i} = S_{t,i}/S_{t,i-1}$. The ARMA$(p, q)$ model for the growth rate of daily oil prices is given by

$$\hat{a}(L)g_{t,i} \quad = \quad \hat{c} + \hat{b}(L)\hat{\epsilon}_{t,i} \quad \forall\ i,\ t, \tag{4}$$

where $\hat{a}(L)$ and $\hat{b}(L)$ are the standard lag polynomials associated with the daily autoregressive and moving average parts, respectively. Estimated parameters are used to construct recursive model-based forecasts of the growth rate of daily nominal crude oil prices, $\hat{g}_{t,n+k|t}$, where k is the forecast horizons in days. The forecasts for the level of the nominal price on day i of month $t + h$, given month t information, are the model-implied cumulative net growth rate between day t, n and day $t + h, i$, denoted by $\hat{g}_{t+h,i|t}$:

$$\hat{S}_{t+h,i|t} = (1 + \hat{g}_{t+h,i|t}) \cdot S_{t,n}. \tag{5}$$

Then, nominal daily forecasts are averaged to the monthly frequency and converted into forecasts of real prices by deflating the monthly nominal forecasts by $C\hat{P}I_{t+h|t}$, the expected CPI deflator:

$$\hat{\bar{R}}_{t+h|t} = \frac{\frac{1}{n}\sum_{i=1}^{n} \hat{S}_{t+h,i|t}}{E_{t,n}[C\hat{P}I_{t+h|t}]}, \quad \forall\ h. \tag{6}$$

We use the standard practice of computing CPI forecasts to discount nominal prices. Following Baumeister and Kilian (2012), the CPI forecast is constructed by expanding the (nowcasted) current CPI observations with the average historical rate of inflation since 1986M7.

The second disaggregated approach for ARMA models is period-end point sampling (PEPS). This method estimates the model using point-in-time sampled end-of-period data and uses the point forecast as the forecast of the average. Ellwanger and Snudden (2023c) document that for persistent data, the forecast accuracy of the PEPS approach is both theoretically and empirically similar to that of the BU approach.

Implementing the PEPS approach is straightforward, as it merely requires replacing the time series of monthly average real prices with the time series of real end-of-period observations, $R_{t,n}$, during the model estimation. The use of end-of-month point-in-time sampling is desirable, as the last observed price reflects the latest information available to the crude oil market participants and the forecaster. On the contrary, older prices generally reflect outdated information that could deviate from current market expectations. The estimation of models with end-of-month real prices is also common for forecasts of end-of-month real prices such as bilateral exchange rates (Meese and Rogoff, 1983) and primary commodities (West and Wong, 2014).

For the PEPS approach, the model estimation closely resembles that of the traditional approach, which relies on monthly average data. Specifically, the ARMA(p, q) model is estimated at the monthly frequency using a time series of end-of-month prices in log-real levels and expressed as:

$$\tilde{a}(L)r_{t,n} \quad = \quad \tilde{c} + \tilde{b}(L)\tilde{\epsilon}_t \quad \forall\ t. \tag{7}$$

Again, the estimates are used to construct recursive model-based forecasts of the end-of-month real price of crude oil $\tilde{r}_{t+h,n|t}$ which are converted back into real prices in levels, $\tilde{R}_{t+h,n|t} = \exp\left(\tilde{r}_{t+h,n|t}\right)$. Then the end-of-month forecasts are used as forecasts of the corresponding monthly average value, $\tilde{\bar{R}}_{t+h|t} = \tilde{R}_{t+h,n|t}$.

Regarding lag length, previous studies advocate the use of 12 or 24 lags when estimating autoregressive models to predict real crude oil prices (Baumeister and Kilian, 2012). However, it is well known that temporal aggregation will always result in the introduction of moving average terms, even when the data generating process of the daily data is given by an AR process (Telser, 1967; Weiss, 1984). The in-sample analysis of both real and nominal daily crude oil prices suggests an AR model with one or two coefficients when using information criterion. This is theoretically consistent with an ARMA representation of the monthly average data and an AR(1) or AR(2) representation for the point sampled data used in the PEPS approach. For completeness, we therefore examine the forecast performance from the standard AR(12) model and also include the forecasts from the AR(2) and ARMA(1,1) models. Qualitatively, our subsequent results are not affected by alternative lag specifications considered in the literature. Instead, the information loss from aggregation turns out to be of first-order importance for the model's forecast performance relative to the choice of model parameterization.

Table 1. Real-time ARMA Forecasts of the Real Price of WTI Crude Oil

Method	Monthly Average Prices			Bottom Up			Period-End-Point 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 monthly average real price of WTI crude oil, 1992M1–2021M1. "Bottom Up" is the ex-post averaged forecast of daily prices; "Period-End-Point Sampling" use the end-of-month forecast (Ellwanger and Snudden, 2023c). Brackets report the p-values for the serial dependence robust tests of Pesaran and Timmermann (2009) for the null of no directional accuracy and of Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the end-of-month no-change forecast at the five percent significance level.

The first three columns of Table 1 present the results for the models estimated with the conventional monthly average data. For most models and forecast horizons considered, the model-based forecasts outperform the average price no-change forecasts, in line with Baumeister and Kilian (2012). Interestingly, in contrast to existing recommendations, the model estimated with two autoregressive parameters outperform the models estimated with twelve parameters in terms of the MSFE for all potential ending points of the evaluation sample since 2008 at all forecast horizons. However, the improvements against the conventional monthly average no-change forecast appear relatively modest (below 10 percent) and do not constitute significant improvements at the five percent level at any horizon. Only for directional accuracy at the 24-month horizon do they outperform the end-of-month no-change forecast at the 5 percent significance level.

The corresponding results for the BU approach, displayed in columns 4-6 of Table 1, present a very different pattern. At the one-month horizon, all three models yield about 40 percent improvements in the MSFE and a directional accuracy of 72 percent. The magnitude of the improvements at the one-month horizon is remarkable, as the forecasts are more accurate than any other model-based forecast previously reported in the literature. However, the absence of bold-faced values highlights that at the 5 percent significance level, none of these forecasts are significantly more accurate than the traditional random walk forecast, which is the end-of-month no-change forecast. In addition, the performance of the BU forecasts deteriorates quickly at longer forecast horizons.

The forecasts from the PEPS approach are evaluated in the final 3 columns of Table 1. For all horizons and models, the forecasts outperform the conventional monthly average no-change benchmark. At the one-month-ahead, the PEPS approach results in the largest improvements in forecast accuracy of up to 44 percent, which are significant at the 0.1 percent significance 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 PEPS forecasts significantly predict the direction of change in prices one-month ahead relative to the end-of-month no-change forecast at the five percent level.

The deterioration of the BU forecasts at longer horizons is because the model is estimated in growth rates, which loses information about the historical mean. For robustness, we also consider BU forecasts for the level of real daily prices. Following Froot and Ramadorai (2005), we impute a daily CPI by assigning the monthly CPI index to the last day of the month and use linear interpolation to impute a daily CPI measure, $CPI_{t,i}$. The model is then estimated using the log level of the daily real prices, and forecasts of daily real prices are averaged to the monthly frequency. The results, presented in the online appendix, show little difference in the forecast accuracy between the distinct BU approaches at short horizons. The results also show that applying the BU approach to daily prices in levels instead of growth rates slightly improves longer-horizon forecasts, albeit by less than the improvements obtained under the PEPS approach. This pattern can be explained by the fact that forecasts from models with daily prices only use data beginning in 1986M1, whereas models estimated with monthly average data and PEPS use data that start in 1973M1. The results highlight the importance of using improved information in crude oil prices, while alternative assumptions for the CPI appear to play only a minimal role.

For each oil price series, Figure 1 reports the evolution of the MSFE and the success ratios for the one-month-ahead forecast from the AR(2) models estimated with average, end-of-month, and daily data. The results suggest that the forecast performance of the BU approach closely tracks that of the end-of-month no-change forecast throughout the forecast evaluation sample, whereas the PEPS forecasts tend to outperform the end-of-month no-change forecast in the later half of the sample. For the BU and PEPS approach, the gains upon the monthly average no-change forecast at the one-month-ahead horizon are significant at the 5 percent level for WTI and RAC crude oil prices for all potential ending dates of the forecast evaluation sample between 1997M1 and 2021M1. For all series and all potential ending points of the evaluation sample, the models estimated with

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

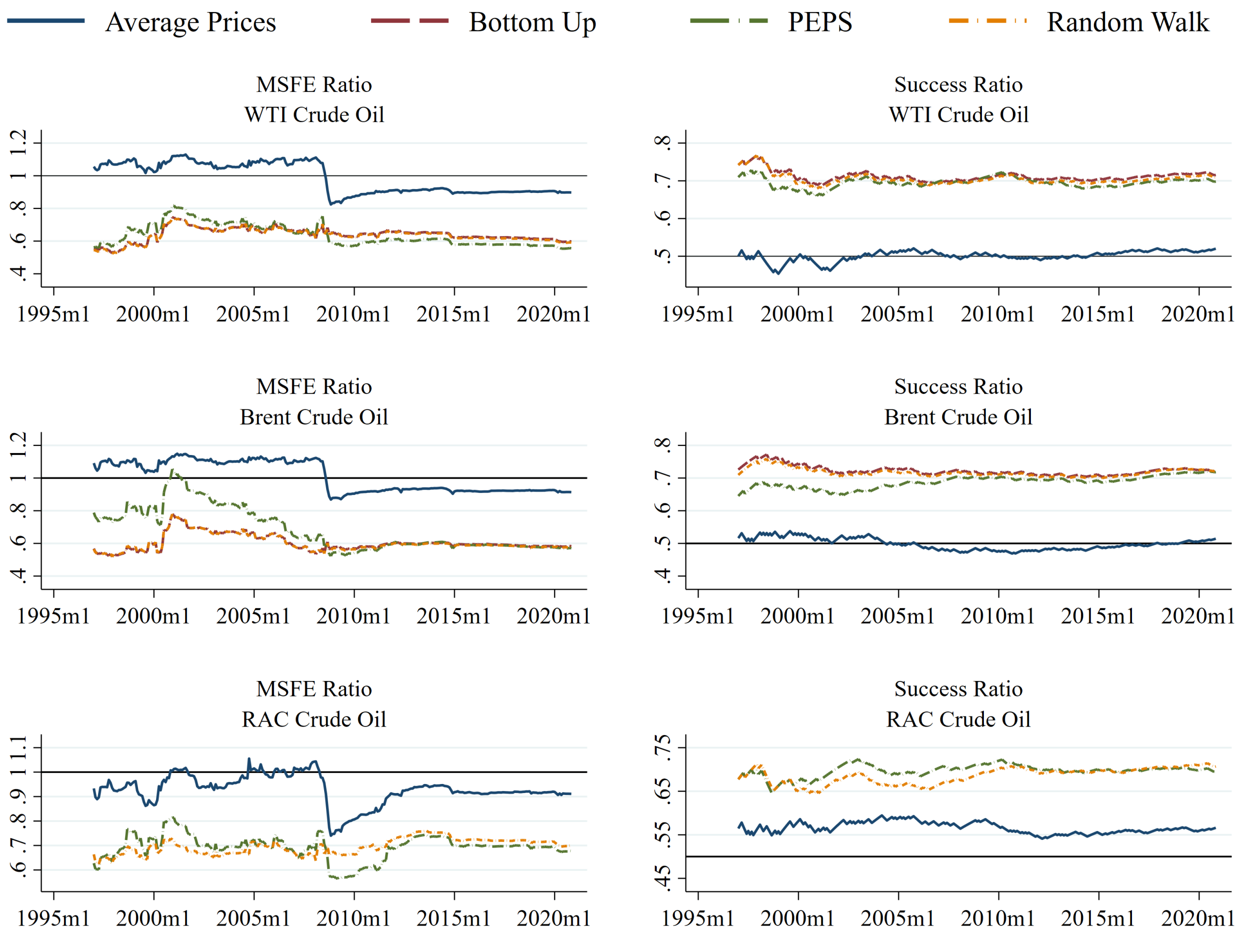

Note: One-month-ahead, out-of-sample, real-time forecasts of alternative monthly average real price of crude oil series, 1992M1–2021M1. Forecasts from AR(2) models estimated with monthly "Average Prices", the "Bottom-Up" approach, and the "PEPS" approach. "Random Walk" is the end-of-month no-change forecast. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

average prices do not outperform the end-of-month no-change forecast.

Overall, these results indicate that suitably constructed univariate time-series models that exploit disaggregated prices outperform models estimated with monthly average prices. Moreover, models estimated with disaggregated oil prices also provide statistically significant improvements in directional accuracy over the end-of-month no-change forecast at the one-month horizon. These findings are consistent with theoretical results indicating that the information gains from using daily rather than monthly average data may be substantial, particularly for short forecast horizons (Wei, 1978). The gains in our setting are much greater than the gains from using monthly series instead of quarterly aggregated series for crude oil prices (Baumeister and Kilian, 2014) and other economic variables (see, e.g., Zellner and Montmarquette, 1971; Abraham, 1982; Lütkepohl, 1986; Athanasopoulos et al., 2011).

3.2 MIDAS Forecasts

An alternative disaggregated approach that we examine is MIDAS (Ghysels et al., 2007; Andreou et al., 2010). The MIDAS approach constructs direct forecasts of lower-frequency variables with higher-frequency predictors. Previous applications of the MIDAS model to oil-price forecasts 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, we explore the relationship between daily and monthly prices of crude oil in a univariate setting. To the best of our knowledge, the MIDAS approach has only been applied for exogenous predictor variables under the assumption that the underlying higher-frequency data of the forecast variable is not observed. Our analysis thus provides a novel application of MIDAS to univariate variables with a known form of aggregation.

MIDAS forecasts of the monthly average real price of crude oil are constructed using daily oil prices via

$$\bar{R}_{t+h} = a_h + \beta_h B(L^{0/m}; \theta_h) R_{t,n} + \epsilon_{t+h}, \tag{8}$$

where $\bar{R}_{t+h}$ is the monthly average real price of crude oil, m is the maximum number of daily lags used in estimation, $L^{1/m} R_{t,n} = R_{t,n-1}$ denotes the lag operator at the daily frequency, and the restrictions on the parameters on the daily data are given by the a lag distribution $B(L^{0/m}; \theta_h)$, for instance the Beta function or the Almon Lag.

In this framework, the estimates of the restricted parameter profiles provide the weights given to the different daily prices. We estimate two of the most common restricted parameter specifications: the normalized exponential Almon lag weight function and the normalized Beta lag polynomial, both for alternative lag lengths of up to 20 daily observations.

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 daily observations beyond one lag. The finding is consistent with the results of Section 3.1, suggesting that the last observed daily oil price contains practically all the relevant information for forecasting the real price of crude oil in a univariate setting.

In the case that the only non-zero weight is on the last observation, Equation 8 simplifies to the following regression equation:

$$\bar{R}_{t+h} = \hat{\alpha}_h + \hat{\beta}_h R_{t,n} + \hat{\epsilon}_{t+h}, \tag{9}$$

i.e., only the last observed daily closing price is used to provide direct forecasts of future average prices. Equation 9 can be interpreted as a special case of the unrestricted MIDAS (UMIDAS) approach of Foroni et al. (2015, 2019), which implies estimating direct relationships between the underlying high-frequency observations and the target variable instead of estimating a parameterized polynomial function. The UMIDAS specification also has the benefit of highlighting the conceptual differences between the alternative disaggregated approaches. For the remainder of the section, we therefore work with UMIDAS forecasts of Equation 9.

Direct forecasts of the real price of oil using real end-of-month prices can also be implemented via PEPS. In this case, the model is estimated with real end-of-month prices as the dependent variable,

$$R_{t+h,n} = \tilde{\alpha}_h + \tilde{\beta}_h R_{t,n} + \tilde{\epsilon}_{t+h}. \tag{10}$$

Then the end-of-month forecast is used as a forecast of the monthly average price, $\tilde{\bar{R}}_{t+h} = \tilde{R}_{t+h,n}$.

Finally, we contrast the MIDAS and PEPS forecasts with direct forecasts estimated on average prices,

$$\bar{R}_{t+h} = \check{\alpha}_h + \check{\beta}_h \bar{R}_t + \check{\epsilon}_{t+h}. \tag{11}$$

For all forecasts, parameters are estimated with an expanding sample window using real-time data, and forecasts are computed out-of-sample.

Table 2 presents the results for the MIDAS and PEPS forecasts along with the direct forecasts constructed from the monthly average prices. It shows that the MIDAS forecasts (columns 4-6) and the 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 estimated with average prices (columns 1-3), and at short-horizons outperform them by a substantial margin. As shown in the online appendix, these forecast gains are robust across the entire forecast evaluation sample.

However, the forecast gains are very similar to those obtained from the end-of-month no-change forecasts. The bold values indicate that at the 5 percent significance level, the null hypothesis of no directional accuracy relative to the end-of-month no-change forecast can only be rejected for

Table 2. Direct Forecasts of the Real Price of Crude Oil with MIDAS and PEPS

Method	Monthly Average Prices			MIDAS			Period-End-Point 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 Timmermann (2009) for the null of no directional accuracy and of Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the end-of-month no-change forecast at the five percent significance level.

the 1-month-ahead PEPS forecasts for RAC. At the 2-year horizon, all forecasts exhibit significant directional accuracy relative to the end-of-month no-change forecast for WTI and RAC. In terms of the MSFE, we find no evidence of statistically significant improvements over the end-of-month no-change forecast for any direct forecast approach.

A key implication from the exercises presented in this section is that in practice, it makes little 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 for constructing accurate forecasts.

The analysis of the MIDAS and BU forecasts provides evidence of the relative merits of direct and recursive forecasts in the context of temporal disaggregation. Previous comparisons of direct and recursive forecasts assume that variables are observed at the same frequency (see, e.g., Ing, 2003; Marcellino et al., 2006; Chevillon, 2007). These studies generally suggest that recursive methods provide better forecasts compared to direct methods. However, the forecasts from the direct MIDAS approach and the recursive BU forecasts tend to be equally accurate at all horizons, especially when using the backcasted end-of-month data. Thus, we find very little difference between direct and recursive forecasts when the data are temporally disaggregated, and both approaches substantially improve upon the forecasts constructed using averaged data. Again, the findings confirm that accounting for temporal aggregation is of first-order importance compared to other modeling choices.

3.3 Exogenous Predictor Variables

In this section, we examine the usefulness of daily oil prices in predicting the real price of crude oil in a setting with exogenous predictor variables. The most common way to incorporate exogenous predictors is via direct forecasts from Linear Regression Models (LRMs). Such models are widely used to provide individual forecasts of the real price of crude oil (see, e.g., Baumeister and Kilian, 2012; Alquist et al., 2013; Chen, 2014) and are also key components of model averaging exercises (Baumeister et al., 2014; Wang et al., 2015; Baumeister and Kilian, 2015; Garratt et al., 2019). The predictive regressions in previous studies take on the form:

$$\check{\bar{R}}_{t+h} = \bar{R}_t \cdot (1 + \check{\alpha}_h + \check{\beta}_h X_t), \tag{12}$$

where X_t is the monthly percent change of predictor variable and $\check{\alpha}_h$ and $\check{\beta}_h$ are the least-squares estimates.

For illustrative purposes, this section focuses on a predictor variable that has been shown to yield particularly accurate forecasts of the real price of crude oil, namely the NYSE Arca (AMEX) oil index, a portfolio of oil-sensitive stock prices (Chen, 2014).[11] The predictor variable, X_t is constructed using the end-of-month AMEX stock index, $p_{t,n}$, in monthly growth rates, i.e., $X_t = p_{t,n}/p_{t-1,n} - 1$.

In addition to the standard specification of Equation 12, we propose a version of the predictive regression where the starting point of the regression is not the monthly average real price, but rather the end-of-month real price. This specification reads

$$\hat{\bar{R}}_{t+h} = R_{t,n} \cdot (1 + \hat{\alpha}_h + \hat{\beta}_h X_t), \tag{13}$$

i.e., the monthly average prices, $\bar{R}_t$, on the right-hand side of equation 13 is replaced with the end-of-month price, $R_{t,n}$. Intuitively, this means that the forecast is started from the end-of-month real price, which is the last non-averaged price available to forecasters.

In addition, we consider the PEPS specification of the LRM which produces a direct forecast

[11] In the online appendix, we document that the results are robust to an alternative predictor variable from the Commodity Research Bureau's (CRB) index of the price of industrial non-oil raw materials (used by Baumeister and Kilian, 2012; Alquist et al., 2013, among others).

of the end-of-month real price via

$$\tilde{R}_{t+h,n} = R_{t,n} \cdot (1 + \tilde{\alpha}_h + \tilde{\beta}_h X_t). \tag{14}$$

As for the other PEPS specifications, the forecast of the end-of-month price is then used as a forecast of the monthly average price, i.e., $\tilde{\bar{\bar{R}}}_{t+h} = \tilde{R}_{t+h,n}$.

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 is found to provide better forecasts for all series, especially at the two-year horizon. Including a constant has only minimal effects on forecasts below one year.

Table 3. Real-time Forecasts of the Real Price of Crude Oil Using the NYSE Arca Oil Index

Method	Montly Average Prices			End-of-Month Prices			Period-End-Point Sampling		
Price	WTI	Brent	RAC	WTI	Brent	RAC	WTI	Brent	RAC
Horizon					MSPE Ratio				
1	0.80 (0.006)	0.78 (0.004)	0.90 (0.071)	0.57 (0.000)	**0.55 (0.000)**	0.70 (0.003)	0.56 (0.000)	**0.54 (0.000)**	0.71 (0.005)
3	0.91 (0.055)	0.91 (0.045)	0.93 (0.121)	0.86 (0.025)	0.88 (0.023)	0.88 (0.044)	0.86 (0.026)	0.88 (0.027)	0.87 (0.045)
6	0.96 (0.108)	0.96 (0.099)	0.98 (0.234)	0.94 (0.040)	0.96 (0.084)	0.95 (0.086)	0.94 (0.033)	0.96 (0.077)	0.95 (0.067)
12	0.98 (0.124)	0.98 (0.094)	0.99 (0.255)	0.95 (0.036)	0.98 (0.118)	0.96 (0.052)	0.95 (0.035)	0.98 (0.115)	0.96 (0.047)
24	1.00 (0.176)	1.00 (0.225)	1.00 (0.658)	**0.98 (0.117)**	1.00 (0.332)	0.99 (0.187)	**0.98 (0.107)**	1.00 (0.346)	0.99 (0.172)
					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)	0.72 (0.000)	**0.72 (0.000)**	0.70 (0.000)
3	0.58 (0.001)	0.56 (0.030)	0.58 (0.003)	0.60 (0.000)	0.56 (0.033)	0.60 (0.000)	0.60 (0.000)	0.56 (0.022)	0.60 (0.000)
6	0.53 (0.231)	0.55 (0.082)	0.55 (0.086)	0.55 (0.092)	**0.56 (0.024)**	0.57 (0.011)	0.56 (0.030)	**0.56 (0.017)**	0.57 (0.005)
12	0.54 (0.083)	0.56 (0.016)	0.54 (0.092)	0.58 (0.003)	0.58 (0.001)	0.55 (0.048)	0.58 (0.001)	0.59 (0.000)	0.56 (0.031)
24	0.55 (0.037)	0.56 (0.014)	0.45 (0.980)	0.56 (0.023)	**0.56 (0.012)**	0.54 (0.151)	**0.56 (0.031)**	0.56 (0.016)	0.53 (0.173)

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 ofPesaran and Timmermann (2009) for the null of no directional accuracy and of Diebold and Mariano (1995) for equal MSFEs relative to the monthly average no-change forecast. Boldface values indicate statistically significant improvements over the end-of-month no-change forecast at the five percent significance level.

The forecast performance of the two approaches is reported in Table 3. Columns 1-3 present the results for models estimated with monthly average oil prices. For all three oil series, the forecasts using monthly average prices significantly outperform the average-price no-change forecast at the one-month horizon. The forecast gains are very similar to those reported in Chen (2014) and are robust over time, as shown in the online appendix. However, none of the forecasts obtained from the traditional forecasting model significantly outperform the end-of-month no-change forecast at any horizon. Again, this shows the importance of comparing forecasts against the no-change forecast from the end-of-month real price.

Columns 4-6 present the forecast performance of the disaggregated approach described by Equation 13, while the last three columns present the performance of the PEPS approach, Equation 14.

These two alternative specifications perform remarkably similarly at all horizons and outperform the forecasts computed with monthly average oil prices at the one- and three-month horizons. Forecasts from these disaggregated approaches show MSFE ratios as low as 0.54 and success ratios as high as 0.72 at the one-month horizon.[12] For Brent, the improvements in the MSFE and success ratio are also significant at the 5-percent significance level relative to the end-of-month no-change forecast at the one-month horizon. This finding is impressive, as it is the first time that any forecast of the monthly average real price of crude oil has been shown to yield significant improvements for both MSFE and the success ratio at the one-month horizon relative to the end-of-month no-change forecast. The MSFE gains are about 6 percent relative to the end-of-month no-change forecast, which is smaller than the 45 percent relative improvement implied by comparisons with the monthly average no-change forecasts. In terms of magnitude, the gains are more reminiscent of the short-horizon predictability occasionally documented for asset prices.

The regression estimates presented in Table 4 for Equations 12–14 provide useful intuition for our results. For the one-month-ahead forecast regressions using monthly average prices, the estimates of $\check{\beta}_1$ are around 0.6 and the R-squared is about 15%. On the contrary, the estimates of $\hat{\beta}_1$ and $\tilde{\beta}_1$ for the disaggregated regressions drop to 0.2–0.3, and the R-squares fall to 2-4%. Despite the decrease in the coefficients and the measure of fit, the out-of-sample forecasts of Table 3 imply that the AMEX index provides some additional predictive power that is not contained in the last available daily oil price.

Table 4. Coefficient Estimates for One-Month-Ahead Predictive Regressions Using the NYSE Arca (AMEX) Oil Index

Sampling	Monthly Average Prices			End-of-Month Prices			Period-End-Point Sampling		
Dep. Variable	WTI	Brent	RAC	WTI	Brent	RAC	WTI	Brent	RAC
AMEX	0.593	0.614	0.548	0.237	0.193	0.196	0.309	0.313	0.257
	(0.067)	(0.069)	(0.063)	(0.053)	(0.060)	(0.050)	(0.080)	(0.090)	(0.075)
Constant	0.0001	0.0004	-0.0002	-0.001	-0.001	-0.001	0.003	0.005	0.003
	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)	(0.005)	(0.006)	(0.005)
N	449	449	449	449	449	449	449	449	449
adj-R^2	0.148	0.149	0.142	0.040	0.020	0.032	0.030	0.024	0.024

Note: Coefficient estimates of $\hat{\alpha}$ and $\hat{\beta}$ from Equations 12–14 for different crude oil price series, 1992M1–2021M1. HAC standard errors in parenthesis.

The intuition for these results becomes clearer when we consider the effect of averaging on the informational content of prices. Aggregation of daily prices into monthly averages results in a loss of information because it dilutes the latest information contained in the end-of-month price with

[12] As shown in the online appendix, this result is robust to alternative evaluation sample end-point.

stale information from previous days. In contrast, end-of-month prices are used for the AMEX index, which implies that this information loss does not occur for the predictor variable. As both the AMEX and oil prices are likely to respond to similar news about the global oil market, the use of the end-of-month AMEX index helps recover some information lost by averaging the price of crude oil. The results of this section suggest that most of the forecasting gains in the original specification of Chen (2014) stem from the recovery of lost information from averaging. However, the fact that we still find some improvement in the forecast accuracy for the disaggregated approaches indicates that the AMEX index contains some additional information beyond that contained in the last price of oil.

Our results indicate that even though the use of disaggregated oil prices is of first-order importance, the addition of external predictor variables can significantly improve the forecasts of the real price of crude oil. This finding is remarkable because it is the first time that any forecast of the real price of crude oil has been shown to significantly outperform the end-of-month no-change forecasts at the one-month horizon. It is unclear whether other variables that have been shown to predict the real price of crude oil also contain predictive power in specifications relying on end-of-month oil prices, suggesting that their predictive power needs to be reevaluated.

3.4 Vector Autoregressive Models

The final model is the 4-variable VAR model of the global market for crude oil of Kilian and Murphy (2014), which was first applied for real-time forecasts by Baumeister and Kilian (2012) and has been widely used and extended in the literature (see, e.g., Funk, 2018; Snudden, 2018, among others). The model contains four variables: the monthly average 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 economic activity index (REA) as a measure of global economic activity. We focus on the original specification of Baumeister and Kilian (2012) for illustrative purposes, since it remains the benchmark for alternative VAR specifications.

Similar to the ARMA model examined in section 3.1, the convention in the literature is to estimate models with the log level of the real monthly average price of crude oil. The VAR(p) model with p autoregressive parameters can be expressed as:

$$(1 - \mathbf{a}(L))\mathbf{y}_t \quad = \quad \mathbf{e}_t \quad \forall\ t. \tag{15}$$

where $\mathbf{y}_t$, is a 4x1 vector of the variables, $\mathbf{a}(L)$ is the autoregressive parameter matrix of order p, and $\mathbf{e}_t$ is a 4x1 vector of innovations. The estimated parameters are used to construct recursive model-based forecasts of the log-real price, $\breve{\bar{r}}_{t+h|t}$ which are converted back into real prices in levels using $\breve{\bar{R}}_{t+h|t} = \exp\left(\breve{\bar{r}}_{t+h|t}\right)$.

Unfortunately, it is not possible to apply the BU approach to the oil market VAR as the three non-price variables are only observed at the monthly frequency. However, one of the practical advantages of PEPS is that this approach maintains the same model frequency as the forecast target and is easy to implement. Applying PEPS to the VAR setting only involves replacing the log level of the monthly average real price of crude oil, $\bar{r}_t$, in the vector $\mathbf{y}_t$, with the log level of the real end-of-month price, $r_{t,n}$. Again, the estimated parameters of this new specification are used to construct recursive forecasts of the end-of-month real price of crude oil $\tilde{r}_{t+h,n|t}$, which are converted back into real prices in levels, $\tilde{R}_{t+h,n|t} = \exp\left(\tilde{r}_{t+h,n|t}\right)$. Then, the end-of-month forecasts are used as forecasts of the corresponding monthly average value, $\tilde{\bar{R}}_{t+h|t} = \tilde{R}_{t+h,n|t}$.

Table 5. Forecasts of the Real Price of Crude Oil with VAR models

Estimation	Monthly Average Prices			Period-End-Point Sampling		
Horizon	WTI	RAC	Brent	WTI	RAC	Brent
			MSPE Ratio			
1	1.03 (0.582)	0.95 (0.383)	1.09 (0.770)	0.72 (0.059)	0.75 (0.104)	0.82 (0.138)
3	1.05 (0.616)	1.03 (0.571)	1.14 (0.829)	0.97 (0.423)	0.95 (0.406)	1.12 (0.763)
6	1.10 (0.849)	1.14 (0.858)	1.20 (0.956)	1.05 (0.682)	1.09 (0.766)	1.20 (0.947)
12	1.16 (0.904)	1.15 (0.876)	1.20 (0.931)	1.13 (0.845)	1.12 (0.830)	1.22 (0.928)
24	1.08 (0.687)	1.05 (0.610)	1.09 (0.706)	1.08 (0.687)	1.05 (0.620)	1.11 (0.724)
			Success Ratio			
1	0.49 (0.596)	0.56 (0.005)	0.51 (0.230)	0.67 (0.000)	0.68 (0.000)	0.66 (0.000)
3	0.53 (0.140)	0.56 (0.020)	0.52 (0.182)	0.59 (0.002)	0.60 (0.001)	0.57 (0.017)
6	0.52 (0.234)	0.54 (0.107)	0.53 (0.169)	0.58 (0.017)	**0.59 (0.004)**	0.59 (0.013)
12	0.55 (0.071)	**0.60 (0.001)**	0.57 (0.041)	**0.61 (0.002)**	0.60 (0.003)	**0.61 (0.002)**
24	0.56 (0.081)	0.58 (0.012)	**0.56 (0.074)**	0.58 (0.035)	**0.58 (0.022)**	**0.57 (0.049)**

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

The estimation follows the real-time forecasts of Baumeister and Kilian (2012). Accordingly, we estimate models with 12 autoregressive lags, although qualitatively similar results are obtained for models with 24 lags. Other than changing average prices to end-of-month prices, no other change is made, and the VAR models are estimated at the monthly frequency.

Table 5 reports the forecast performance of the VAR models for the standard specification, as well as the PEPS extension. Unlike Baumeister and Kilian (2012), 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 12 months and beyond. The deterioration in forecast accuracy in the updated sample is consistent with evidence from other studies (Funk, 2018; Snudden, 2018).

The table also shows that for most cases, the VAR models estimated with PEPS outperform the VAR models estimated with monthly average prices. This is particularly true for forecast horizons of up to 6 months. The improvements in forecast accuracy at the one-month horizon are of the order of 30 percent and, as shown in the online appendix, robust over the second half of the forecast evaluation sample. These results illustrate that model-based forecasts computed with monthly average prices are also inefficient in the VAR setting. The use of end-of-month prices in VARs helps to restore the information loss introduced from averaging, which is particularly relevant for short-horizon forecasts.

4 Conclusion

We have shown how disaggregated approaches that exploit information from the underlying daily crude oil prices can be applied to current model-based forecasts methods. Forecasts utilizing information in daily prices produce forecast improvements over the monthly average no-change forecast of up to 45% at the one-month horizon, which is unprecedented in this literature. The magnitude of the gains decreases 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 traditional random walk forecast, which is the end-of-month no-change forecast. We find that none of the current forecast practices which are estimated with monthly average prices outperform the end-of-month no-change forecasts at horizons lower than 6 months. By contrast, some disaggregated approaches displayed significantly lower MSFE ratios or improved directional accuracy relative to the end-of-period no-change forecast. In economic terms, the forecast gains relative to this benchmark are reminiscent of the short-horizon predictability documented for asset prices.

The practical implications of incorporating nominal daily data into model-based forecasts for real period-average series present an exciting avenue for further exploration. Our study shows that even straightforward applications of disaggregated methods result in substantial improvements over models relying solely on average data, reinforcing the need to utilize the valuable information contained within daily data for forecasting average series.

References

Abraham, B. (1982). Temporal aggregation and time series. *International Statistical Review*, 50(3):285–291.

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 Kourtellos, A. (2010). Regression models with mixed sampling frequencies. *Journal of Econometrics*, 158(2):246–261.

Athanasopoulos, G., Hyndman, R. J., Song, H., and Wu, D. C. (2011). The tourism forecasting competition. *International Journal of Forecasting*, 27(3):822–844.

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 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(1):S33–S43.

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.

Chevillon, G. (2007). Direct multi-step estimation and forecasting. *Journal of Economic Surveys*, 21(4):746–785.

Degiannakis, S. and Filis, G. (2018). Forecasting oil prices: High-frequency financial data are indeed useful. *Energy Economics*, 76:388–402.

Diebold, F. X. (2015). Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of diebold–mariano tests. *Journal of Business & Economic Statistics*, 33(1):1–9.

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. (2023a). Forecasts of the real price of oil revisited: Do they beat the random walk? *Journal of Banking and Finance*, 154:1–8.

Ellwanger, R. and Snudden, S. (2023b). Futures prices are useful predictors of the spot price of crude oil. *The Energy Journal*, 44(4):45–62.

Ellwanger, R. and Snudden, S. (2023c). Seize the last day: Period-end-price sampling for forecasts of temporally aggregated data. *LCERPA Working Paper No. 2023-6*.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417.

Filippidis, M., Filis, G., and Magkonis, G. (2024). Evaluating oil price forecasts: A meta-analysis. *The Energy Journal*, 45(2).

Foroni, C., Marcellino, M., and Schumacher, C. (2015). Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society. Series A*, 178(1):57–82.

Foroni, C., Marcellino, M., and Stevanovic, D. (2019). Mixed-frequency models with moving-average components. *Journal of Applied Econometrics*, 34(5):688–706.

Froot, K. A. and Ramadorai, T. (2005). Currency returns, intrinsic value, and institutional-investor flows. *The Journal of Finance*, 60(3):1535–1566.

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.

Ing, C.-K. (2003). Multistep prediction in autoregressive processes. *Econometric Theory*, 19(2):254–279.

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. (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.

Marcellino, M., Stock, J. H., and Watson, M. W. (2006). A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1-2):499–526.

Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24.

Mork, K. A. (1989). Oil and the macroeconomy when prices go up and down: an extension of hamilton's results. *Journal of Political Economy*, 97(3):740–744.

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.

Telser, L. G. (1967). Discrete samples and moving sums in stationary stochastic processes. *Journal of the American Statistical Association*, 62(318):484–499.

Tiao, G. C. (1972). Asymptotic behaviour of temporal aggregates of time series. *Biometrika*, 59(3):525–531.

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.

Weiss, A. A. (1984). Systematic sampling and temporal aggregation in time series models. *Journal of Econometrics*, 26(3):271–281.

West, K. D. and Wong, K.-F. (2014). A factor model for co-movements of commodity prices. *Journal of International Money and Finance*, 42:289–309.

Zellner, A. and Montmarquette, C. (1971). A study of some aspects of temporal aggregation problems in econometric analyses. *The Review of Economics and Statistics*, 53(4):335–342.

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

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 the 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 by 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.

Consumer Price Index: 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.

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 represent 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. The nominal crude-oil price data is obtained from the Energy Information Administration and the consumer price index is obtained from the Federal Reserve of Philadelphia.

Oil market variables used in the VAR: 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. The updated real-time data is similar to Garratt et al. (2019), but has the advantage that all historical releases, and thus data revisions, are included. This real-time database is updated monthly and is publicly available, along with further documentation, on the author's website.

We construct real-time vintages of the Real Economic Activity index using (Kilian, 2009)'s freight rates until 1984M12 and monthly values of the Baltic Dry Index obtained from Bloomberg thereafter. The index is calculated using the corrected formula of (Kilian, 2019).

Predictor variables used in the LRM: The index of prices of industrial non-oil raw materials is obtained from the Commodity Research Bureau (BVY00) and is available in real time. 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.

A2 Additional Results

A2.1 Bottom-Up Forecasts

This section considers an alternative way to implement the BU approach for the log level of the daily real price. We impute a daily CPI by assigning the monthly index to the last day of the month and applying a linear interpolation for the daily measure of CPI, $CPI_{t,i}$. Once a daily CPI is calculated, the model can be estimated in daily log real prices and daily forecasts can be averaged to monthly frequency.[1]

Specifically, the daily log real price is computed by deflating the daily nominal price by the imputed daily CPI index, $r_{t,i} = ln(R_{t,i}/CPI_{t,i})$. The ARMA$(p, q)$ model is estimated in log-real levels of the daily prices,

$$\hat{a}(L)r_{t,i} \quad = \quad \hat{c} + \hat{b}(L)\hat{\epsilon}_{t,i} \qquad \forall \; t \; , i. \tag{1}$$

The estimated parameters are used to construct recursive model-based forecasts of the daily real price $\hat{r}_{t+h,i|t}$ which are converted back into real prices in levels and averaged to the monthly frequency,

$$\hat{\hat{R}}_{t+h|t} = \frac{1}{n}\sum_{i=1}^{n} \exp(\hat{r}_{t+h,i|t}), \quad \forall \; h. \tag{2}$$

At the one-month horizon, the BU forecasts using the daily real price of crude oil results are almost identical to the baseline BU results discussed in Section 3.1. At the 3- to 12-month horizons, the BU approach using daily real prices tends to be slightly worse, especially for directional accuracy. At the 24-month horizon, the forecast performance of the BU approach using real daily prices is found to be slightly more accurate. Our key results remain unchanged. Both BU approaches outperform the forecasts constructed from monthly average prices at short forecast horizons. Furthermore, the PEPS approach results in superior forecasts compared to all the BU approaches considered.

[1]The choice of the imputation procedure has minimal effects on the results. Moreover, the results are almost identical to avoiding imputation altogether by assigning the same monthly CPI, CPI_t, to every daily observation of month t.

Table A2. Comparison of Alternative Methods to Construct BU Forecasts

Method	Nominal Daily Forecasts			Daily Real Forecasts		
Model	AR(12)	AR(2)	ARMA(1,1)	AR(12)	AR(2)	ARMA(1,1)
Horizon	MSFE Ratio					
1	0.61 (0.000)	0.60 (0.000)	0.60 (0.000)	0.61 (0.000)	0.60 (0.000)	0.61 (0.000)
3	0.93 (0.021)	0.92 (0.018)	0.94 (0.022)	0.93 (0.155)	0.97 (0.406)	0.98 (0.429)
6	1.04 (0.723)	1.03 (0.666)	1.05 (0.753)	1.04 (0.716)	1.08 (0.689)	1.08 (0.653)
12	1.16 (0.908)	1.14 (0.888)	1.16 (0.913)	1.09 (0.809)	1.17 (0.783)	1.18 (0.773)
24	1.48 (0.989)	1.42 (0.987)	1.45 (0.989)	1.16 (0.831)	1.30 (0.883)	1.33 (0.891)
	Success Ratio					
1	0.72 (0.000)	0.72 (0.000)	0.72 (0.000)	0.73 (0.000)	0.69 (0.000)	0.68 (0.000)
3	0.61 (0.000)	0.61 (0.000)	0.61 (0.000)	0.54 (0.083)	0.54 (0.078)	0.53 (0.145)
6	0.57 (0.042)	0.56 (0.068)	0.55 (0.100)	0.49 (0.514)	0.50 (0.407)	0.50 (0.399)
12	0.55 (0.143)	0.55 (0.142)	0.53 (0.408)	0.50 (0.423)	0.53 (0.168)	0.53 (0.198)
24	0.51 (0.958)	0.51 (0.718)	0.51 (0.900)	0.56 (0.071)	0.57 (0.047)	0.56 (0.070)

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

A2.2 Exogenous Predictors

This section describes the results of the LRM approach estimated with the Commodity Research Bureau's (CRB) index of the price of industrial non-oil raw materials. The predictor variable X_t is constructed using the monthly average CRB index, $\bar{p}_t$, in monthly growth rates, i.e., $X_t = \bar{p}_t/\bar{p}_{t-1} - 1$. The forecasts are constructed using equations 12–14.

Table A3. Real-time Forecasts of the Real Price of Crude Oil Using the CRB Commodity Index

Method	Montly Average Prices			End-of-Month Prices			Period-End-Point Sampling		
Price	WTI	Brent	RAC	WTI	Brent	RAC	WTI	Brent	RAC
Horizon				MSPE Ratio					
1	0.86 (0.087)	0.87 (0.072)	0.93 (0.278)	0.54 (0.001)	0.53 (0.000)	0.69 (0.028)	0.54 (0.001)	0.54 (0.000)	**0.71 (0.043)**
3	0.85 (0.102)	0.87 (0.098)	0.87 (0.148)	0.78 (0.058)	0.82 (0.059)	0.80 (0.083)	0.77 (0.064)	0.82 (0.067)	0.79 (0.090)
6	0.93 (0.174)	0.93 (0.164)	0.94 (0.218)	0.90 (0.118)	0.92 (0.133)	0.91 (0.148)	0.90 (0.121)	0.92 (0.142)	0.91 (0.153)
12	1.09 (0.830)	1.04 (0.717)	1.12 (0.877)	1.05 (0.704)	1.03 (0.636)	1.08 (0.786)	1.06 (0.722)	1.03 (0.656)	1.09 (0.802)
24	1.07 (0.814)	1.02 (0.600)	1.08 (0.792)	1.06 (0.778)	1.02 (0.589)	1.06 (0.751)	1.06 (0.788)	1.02 (0.603)	1.06 (0.761)
				Success Ratio					
1	0.54 (0.083)	0.56 (0.018)	0.70 (0.000)	**0.72 (0.000)**	0.72 (0.000)	**0.72 (0.000)**	**0.72 (0.000)**	0.72 (0.000)	**0.72 (0.000)**
3	0.56 (0.043)	**0.53 (0.211)**	**0.60 (0.000)**	0.65 (0.000)	0.59 (0.002)	**0.64 (0.000)**	0.64 (0.000)	0.60 (0.001)	**0.64 (0.000)**
6	0.55 (0.067)	**0.56 (0.024)**	**0.57 (0.011)**	0.58 (0.005)	0.60 (0.000)	**0.61 (0.000)**	0.58 (0.007)	0.60 (0.000)	**0.61 (0.000)**
12	**0.55 (0.063)**	**0.56 (0.044)**	**0.55 (0.048)**	**0.57 (0.014)**	0.58 (0.010)	0.57 (0.018)	0.57 (0.014)	**0.58 (0.008)**	0.57 (0.018)
24	0.52 (0.319)	**0.55 (0.074)**	0.54 (0.151)	0.54 (0.157)	0.57 (0.018)	0.54 (0.140)	0.54 (0.157)	0.58 (0.014)	0.54 (0.140)

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

Table A3 reports the forecast performance of the CRB index for the monthly average real price of crude oil. The forecasts estimated with monthly average prices can only significantly outperform the average-price no-change forecast at the one-month horizon for Brent and the RAC and only for the success ratio at the 5 percent significance level. Moreover, none of the forecasts significantly outperforms the end-of-month no-change forecast at any horizon.

The later columns show that the disaggregated approaches substantially improve forecast accuracy relative to the standard specification and perform very similarly at all horizons. Forecasts from these disaggregated approaches exhibit MSFE ratios as low as 0.53 and success ratios as high as 0.72 at the one-month-ahead prediction. For RAC, the improvements in the MSFE and success ratio are significant at the 5-percent level relative to the end-of-month no-change forecast at the one-month horizon. The directional accuracy is also remarkable, with significant gains upon the end-of-month no-change forecast for the one- to six-month forecast horizons.

A2.3 Robustness to sample period

The recursively updated MSFE and success ratios for the one-step-ahead forecast of the MIDAS models are shown in Figure A1. Across the entire sample, the forecast gains from using MIDAS or PEPS are very similar to those obtained from the end-of-month no-change forecast from the last observed daily price. However, for WTI and RAC, the models estimated with PEPS slightly outperform the MIDAS forecasts in terms of directional accuracy, especially towards the later half of the sample.

Figure A1. Evolution of One-Month Ahead Forecast Evaluation Criteria for Direct Forecasts with MIDAS and PEPS

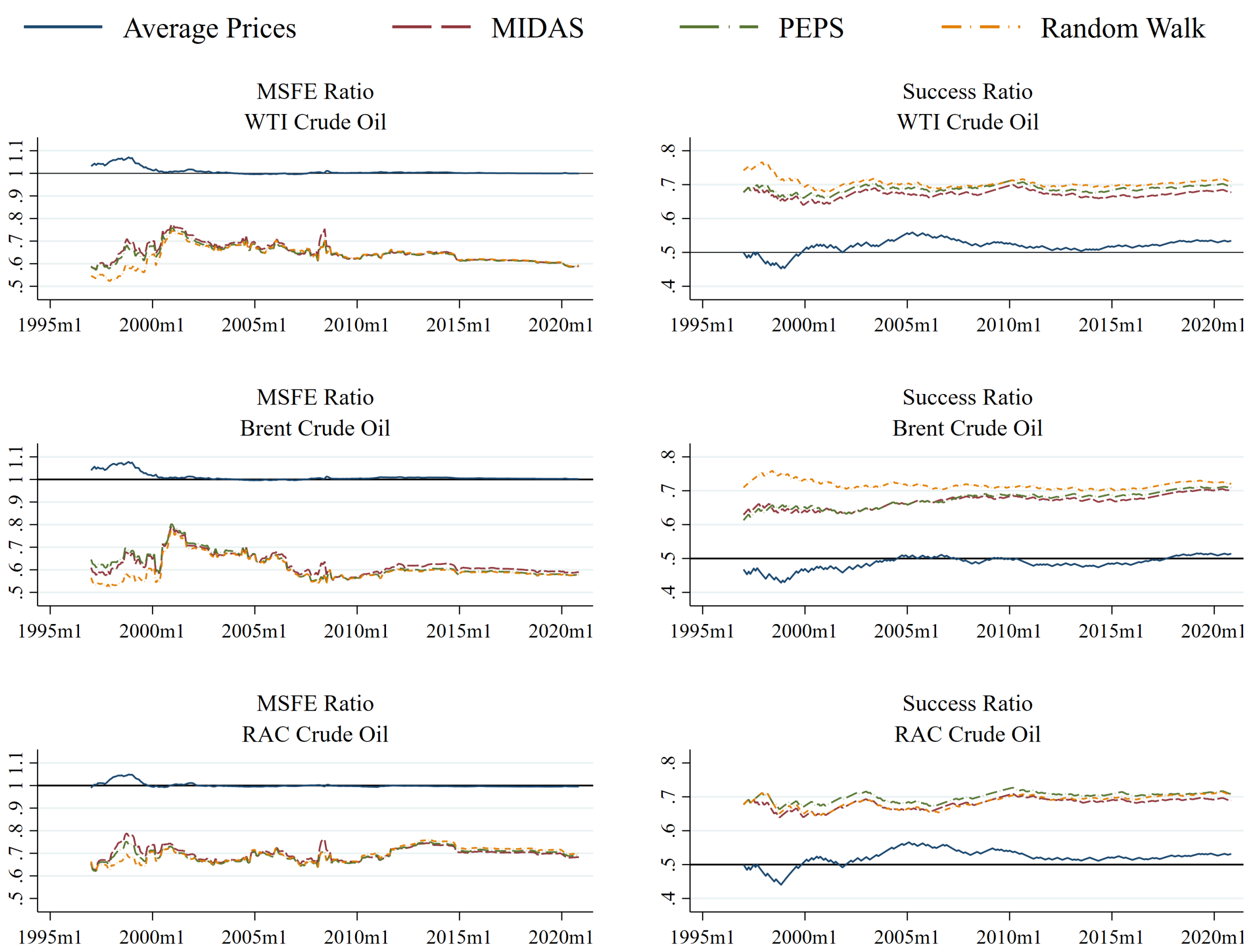

Note: One-month-ahead, real-time, out-of-sample forecasts of alternative crude oil price series, 1992M1–2021M1. MIDAS forecasts computed from Equation 9, PEPS forecasts computed from Equation 10, and direct forecasts using monthly "Average Prices" computed from Equation 11. "Random Walk" is the end-of-month no-change forecast. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

The recursively estimated MSFE and success ratios for the one-month-ahead forecasts are reported in Figure A2. The figure shows that the LRM forecasts estimated with monthly average prices following Chen (2014) do well in the late 2000s and maintain their superiority relative to the conventional monthly average no-change forecast thereafter. However, for all ending points, the forecasts constructed with monthly average prices perform worse than the end-of-month no-change forecasts. In contrast, for both Brent and WTI, the disaggregated approaches maintain small MSPE gains relative to the end-of-month no-change forecast at the one-month horizon.

Figure A2. Evolution of One-Month-Ahead Forecast Evaluation Criteria for Forecasts Using the NYSE Arca Oil Index

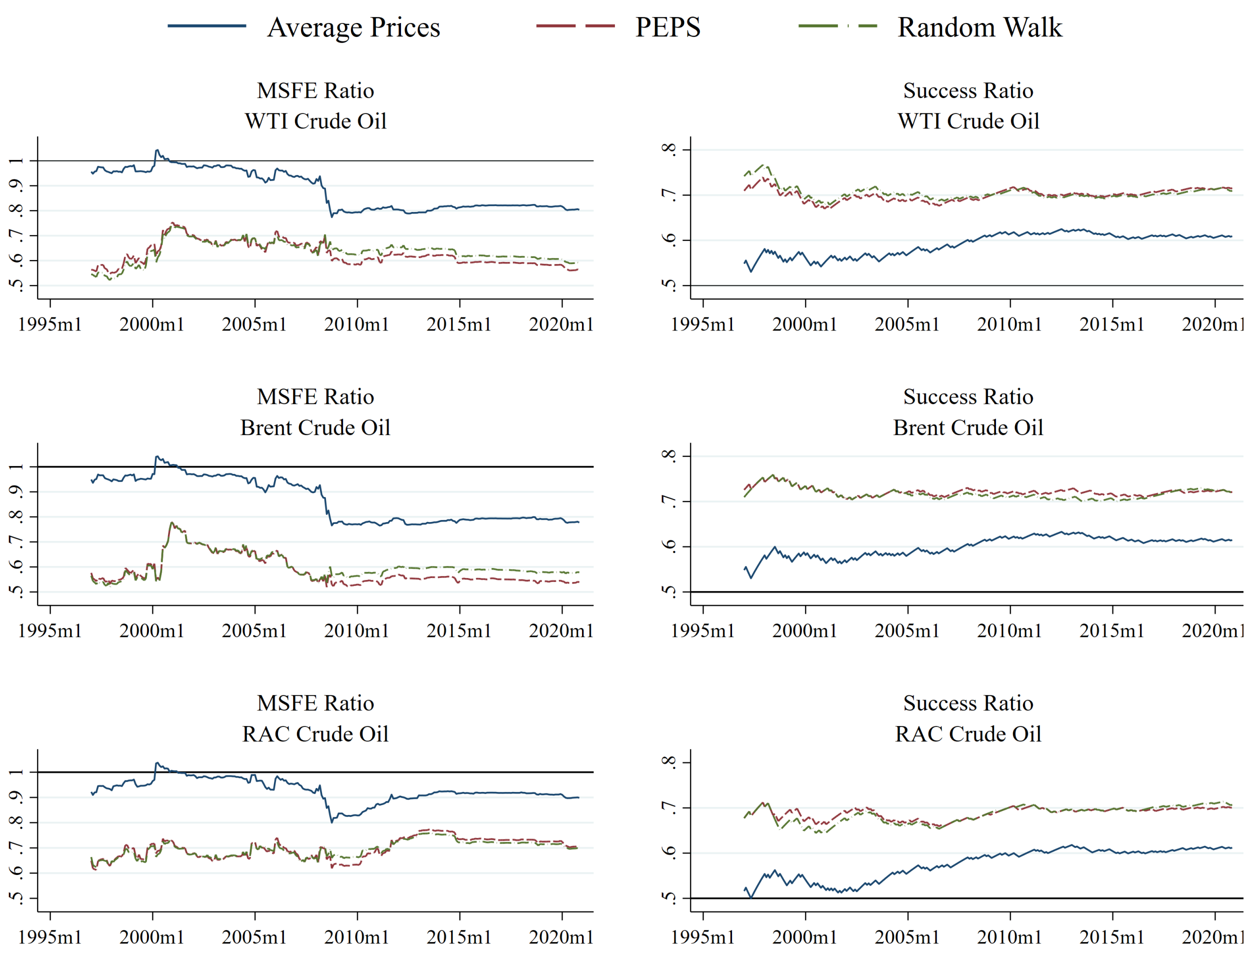

Note: One-month-ahead, real-time, out-of-sample direct forecasts of the real prices of crude oil using the NYSE Arca oil index, 1992M1–2021M1. "Random Walk" is the end-of-month no-change forecast. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted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

Figure A3 depicts the evolution of the forecast criteria for the VAR forecasts over time and shows that the deterioration of the MSFE ratios for the conventional VAR occurs primarily during the years 2010 to 2014. The figure also shows that the improvements in forecast accuracy of the

PEPS forecasts at the one-month horizon are on the order of 30 percent and, with a minor exception in the early 2000s, very robust across the sample period. In fact, between 2008 and 2012, the VAR models estimated with PEPS substantially improved upon the end-of-month no-change forecast in terms of MSFE at the one-month horizon. However, Figure A3 shows that the original VAR model only performed well at the one-month horizon during the financial crisis.

Figure A3. Evolution of One-Month-Ahead Forecast Evaluation Criteria for VAR Models

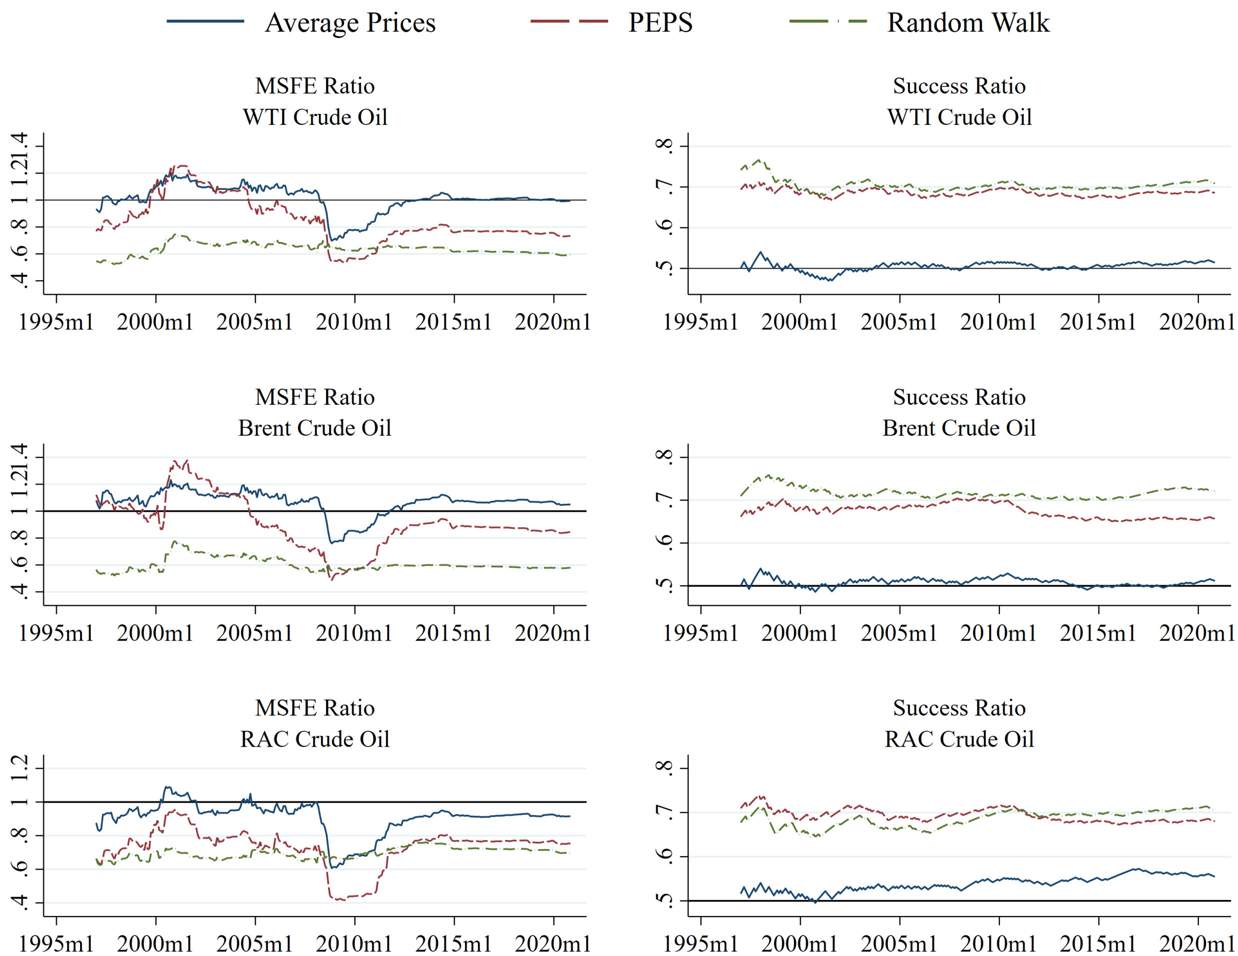

Note: One-month-ahead, real-time, out-of-sample forecasts of the real prices of crude oil, 1992M1–2021M1. VAR models are estimated with 12 lags. "Random Walk" is the end-of-month no-change forecast. MSFE ratios are expressed relative to the monthly average no-change forecast. The first 60 months are omitted to reduce starting-point effects.

References

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

Chen, S.-S. (2014). Forecasting crude oil price movements with oil-sensitive stocks. *Economic Inquiry*, 52(2):830–844.

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

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.

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. (2019). Measuring global real economic activity: Do recent critiques hold up to scrutiny? *Economics Letters*, 178:106–110.

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