International Journal of Forecasting 31 (2015) 238–252

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Do high-frequency financial data help forecast oil prices? The MIDAS touch at work

Christiane Baumeister [a], Pierre Guérin [a], Lutz Kilian [b,c,*]

[a] Bank of Canada, Canada

[b] University of Michigan, United States

[c] CEPR

ARTICLE INFO

Keywords:
Mixed frequency
Real-time data
Oil price
Forecasts

ABSTRACT

In recent years there has been an increased interest in the link between financial markets and oil markets, including the question of whether financial market information helps to forecast the real price of oil in physical markets. An obvious advantage of financial data in forecasting monthly oil prices is their availability in real time on a daily or weekly basis. We investigate the predictive content of these data using mixed-frequency models. We show that, among a range of alternative high-frequency predictors, cumulative changes in US crude oil inventories in particular produce substantial and statistically significant real-time improvements in forecast accuracy. The preferred MIDAS model reduces the MSPE by as much as 28% compared with the no-change forecast and has a statistically significant directional accuracy as high as 73%. This MIDAS forecast is also more accurate than a mixed-frequency real-time VAR forecast, but is not systematically more accurate than the corresponding forecast based on monthly inventories. We conclude that there is not typically much lost by ignoring high-frequency financial data in forecasting the monthly real price of oil.

© 2014 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.

1. Introduction

The substantial variation in the real price of oil since 2003 has renewed interest in the question as to how monthly and quarterly oil prices should be forecast.[1] The links between financial markets and the price of oil have received particular attention, including the question of whether financial market information may help forecast the price of oil in physical markets (e.g., Fattouh, Kilian, & Mahadeva, 2013). One obvious advantage of financial data is their availability in real time at high frequency. Financial data are not subject to revisions and are available on a daily or weekly basis. Existing forecasting models for the monthly real price of oil do not take advantage of these rich data sets. Our objective is to assess whether there is useful predictive information for the real price of oil in high-frequency data from financial and energy markets, and to identify which predictors are most useful. The incorporation of daily or weekly data into monthly oil price forecasts requires the use of models for mixed-frequency data.

The development of models for variables sampled at different frequencies has attracted a substantial amount of interest in recent years. A comprehensive review is provided by Foroni, Ghysels, and Marcellino (2013). A large and growing body of literature has documented the

* Correspondence to: Department of Economics, University of Michigan, 309 Lorch Hall, 611 Tappan Street, Ann Arbor, MI 48104, United States.

E-mail address: lkilian@umich.edu (L. Kilian).

[1] A comprehensive review of this body of literature is provided in the handbook chapter by Alquist, Kilian, and Vigfusson (2013). More recent contributions not covered in that review include the studies by Baumeister and Kilian (2014a, 2014b), Baumeister, Kilian, and Zhou (2013), Bernard, Khalaf, Kichian, and Yelou (2013), and Chen (2014).

http://dx.doi.org/10.1016/j.ijforecast.2014.06.005
0169-2070/© 2014 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.

benefits of combining data of different frequencies for forecasting macroeconomic variables such as real GDP growth and inflation. One approach has been to construct mixed-frequency vector autoregressive (MF-VAR) forecasting models (e.g., Schorfheide & Song, 2014). An alternative approach involves the use of univariate mixed-data sampling (MIDAS) models (e.g., Andreou, Ghysels, & Kourtellos, 2011). The MIDAS model employs distributed lag polynomials to ensure a parsimonious model specification, while allowing for the use of data sampled at different frequencies. The original MIDAS model requires nonlinear least squares estimation (see Andreou, Ghysels, & Kourtellos, 2010). Foroni, Marcellino, and Schumacher (in press) propose a simplified version of the MIDAS model (referred to as unrestricted MIDAS or U-MIDAS) that may be estimated by ordinary least squares, and it has been shown to produce highly accurate out-of-sample forecasts in many applications, provided that the data frequencies to be combined are not too different.

Numerous studies have documented the ability of MIDAS regressions to improve the accuracy of quarterly macroeconomic forecasts based on monthly predictors, and the accuracy of monthly forecasts based on daily or weekly predictors (e.g., Andreou, Ghysels, & Kourtellos, 2013; Armesto, Engemann, & Owyang, 2010; Clements & Galvao, 2008, 2009; Ghysels & Wright, 2009; Hamilton, 2008). In practice, the use of high-frequency financial data is of particular interest, because financial asset prices embody forward-looking information. Another reason for this interest is that financial data are measured accurately and are available in real time, while lower-frequency macroeconomic data tend to be subject to revisions and are available only with a delay.

These differences in informational structure are particularly evident when forecasting oil prices. Commonly used predictors of the real price of oil, such as global oil production, global oil inventories, global real activity, or the US refiners' acquisition cost for crude oil, only become available with considerable delays and are subject to potentially large, but unpredictable, revisions that may persist for up to two years (see Baumeister & Kilian, 2012). Despite these drawbacks, several recent studies have shown that it is possible to systematically beat the no-change forecast of the monthly real price of oil in real time (e.g., Baumeister & Kilian, 2012, 2014a, 2014b).

The current paper investigates whether the accuracy of oil price forecasts can be improved by utilizing high-frequency information from financial markets and from US energy markets. The set of high-frequency predictors includes (1) the spread between the spot prices of gasoline and crude oil; (2) the spread between the oil futures price and the spot price of crude oil; cumulative percentage changes in (3) the Commodity Research Bureau (CRB) index of the price of industrial raw materials, (4) US crude oil inventories, and (5) the Baltic Dry Index (BDI); (6) returns and excess returns on oil company stocks; (7) cumulative changes in US nominal interest rates (LIBOR, Fed funds rate); and (8) cumulative percentage changes in the US trade-weighted nominal exchange rate.

Our starting point is a MIDAS model for the monthly real price of oil. For reasons discussed in Section 2, we focus initially on predictors measured at weekly intervals and constructed from daily observations. As is standard in the oil price forecasting literature, we assess all forecasts based on their mean-squared prediction errors and directional accuracy. We consider forecast horizons, h, ranging from 1 month to 24 months. Our MIDAS models nest the no-change forecast of the real price of oil, allowing us to compare the accuracy of MIDAS regressions with those of competing models evaluated against the same benchmark. We also compare the MIDAS model forecasts to real-time forecasts from the corresponding model based on the same predictors measured at monthly frequency.

Our results reinforce and strengthen recent evidence that the monthly real price of oil can be forecast in real time. We find that the most accurate h-month-ahead forecasts are obtained based on the percentage change in US crude oil inventories over the preceding h months. For example, the preferred MIDAS forecast has a statistically significant directional accuracy as high as 56% at the 12-month horizon, and as high as 69% at the 24-month horizon. It also produces mean-squared prediction error (MSPE) reductions relative to the no-change forecast of 8% at the 12-month horizon and of 28% at the 24-month horizon. These improvements in forecast accuracy are large by the standard of previous work on forecasting oil prices. However, at horizons shorter than 12 months, the MSPE reductions of this MIDAS model are quite modest or nonexistent.

The way in which the MIDAS model is implemented matters to some extent. While there is typically little difference in accuracy between the MIDAS model with equal weights and the MIDAS model with estimated weights, the unrestricted MIDAS model tends to be slightly less accurate than the other specifications. The success of these MIDAS forecasts based on US crude oil inventories prompted us to also investigate the accuracy of the MF-VAR model obtained by including the same weekly inventory data in a monthly oil market VAR forecasting model of the type examined by Baumeister and Kilian (2012). We found that the latter specification did not perform systematically better than the original VAR model, and was clearly worse than the MIDAS model. The MIDAS model for US crude oil inventories does not have systematically lower MSPEs than the corresponding forecasting model based on monthly US inventory data, however, and has comparable directional accuracy.

While the improvements in forecast accuracy are less substantial for other weekly financial predictors, the pattern of results is similar. Although MIDAS models often significantly outperform the no-change forecast, the corresponding forecasts from models based on monthly financial predictors do too, and there is little to choose between these models. Examples include models based on oil futures prices, returns on oil company stocks and gasoline price spreads. In some cases, the MIDAS model forecasts are actually inferior to the forecasts from the corresponding monthly model, or fail to improve on the no-change forecast.

These conclusions are robust to whether the MIDAS models are estimated based on daily or weekly data. Even when MIDAS models work well, therefore, not much is lost by ignoring high-frequency financial data in forecasting the monthly real price of oil. Not only is this finding important for applied oil price forecasters, it is also interesting from a methodological point of view. It reminds us that, despite the intuitive appeal of MIDAS models, it is by no means a foregone conclusion that the use of daily or weekly predictors will improve the accuracy of monthly forecasts: the answer depends on whether the additional signal contained in the high-frequency data compensates for the additional noise. Different empirical applications may produce different results.

The remainder of the paper is organized as follows. In Section 2, we review our data sources and the conventions used in transforming the daily data to weekly frequency. Section 3 provides a brief overview of the mixed-frequency forecasting models. Section 4 explains our reasons for selecting the high-frequency predictors and contains the main empirical results. We also show that our results are robust to changes in the data frequency and to the use of forecast combinations. The concluding remarks are in Section 5.

2. Data

Our objective is to compare the real-time out-of-sample forecast accuracies for the monthly real price of oil of a set of models that include high-frequency data from financial and energy markets. We focus on forecasts of the real US refiners' acquisition cost of crude oil imports, which is a widely used proxy for the global price of oil (see Alquist et al., 2013). The refiners' acquisition cost measures what refiners actually pay for the crude oil they purchase. We deflate this price by the US consumer price index for all urban consumers.

2.1. Data construction

For the time being, even if daily data are available, we focus on data measured at the weekly frequency, for two reasons. First, in the early part of the sample there are gaps in the daily data for some of the time series that we consider. By relying on weekly data, we are able to construct internally consistent time series for longer time spans. Second, some of our data are available only at weekly frequency, and the choice of weekly data facilitates comparisons across forecasting models.

One complication that arises with weekly data is that some months consist of five weeks instead of four. We adapt the approach proposed by Hamilton and Wu (2014) in order to generate a balanced weekly data set where each month consists of four weeks. Week 1 ends on the 5th business day of the month, week 2 ends on the 10th business day of the month, week 3 ends on the day when the near-term contract expires, which is approximately on the 15th business day of the month, and week 4 ends on the last business day of the month, which is the date on which the forecasts are considered to be generated.[2] Our weekly predictors correspond to the log-level, the weekly growth rate or the cumulative growth rate of the variable of interest, observed on the last trading day of the week. If no data are available for a given day, we use the preceding daily observation. Cumulative growth rates over h months are defined as the percentage change between the current daily observation and the corresponding daily observation exactly h months earlier. Monthly variables are constructed as averages of daily data over the month (and then transformed as appropriate), consistent with the construction of the US Energy Information Administration (EIA) oil price data.

2.2. Data sources

The daily West Texas Intermediate (WTI) spot oil price is obtained from the Wall Street Journal, and the corresponding daily NYMEX oil futures prices for maturities of 1–18 months are obtained from Bloomberg.[3] Daily data for the spot price of regular gasoline for delivery in New York Harbor are available from the EIA for the period June 1986 to March 2013.[4] The daily spot price index for non-oil industrial raw materials from the CRB is available from June 1981 onwards. Daily data for the BDI are obtained from Bloomberg starting in January 1985. Data for US crude oil inventories are reported from August 1982 onwards in the *Weekly Petroleum Status Report* issued by the EIA, but consistent weekly time series could only be constructed back to January 1984, due to gaps in the earlier data. Our analysis takes account of the fact that this report is issued every Wednesday and contains data extending to the preceding Friday. The closing price of the price-weighted NYSE Arca Oil Index is available from *Yahoo! Finance* from September 1983 onwards. This index is designed to measure the performance of the oil industry through changes in the stock prices of a cross-section of widely-held corporations involved in the exploration, production, and development of petroleum.[5] Daily data for the closing price of the NYSE composite index, which measures the performances of all common stocks listed on the New York Stock Exchange, are obtained from *Yahoo! Finance* for the period January 1966 to March 2013. Weekly data for the federal funds rate, the 3-month LIBOR rate and the nominal trade-weighted US dollar index for major currencies are available from the FRED database from July 1954, January 1986 and January 1973, respectively, onwards.

The monthly real-time data for world oil production, the Kilian (2009) index of global real economic activity, the nominal refiners' acquisition cost of imported crude oil, the US consumer price index for all urban consumers, and the proxy for global crude oil inventories are taken from the real-time database developed by Baumeister and Kilian (2012), which contains vintages from January 1991 to March 2013.

[2] For a Bayesian approach to the modeling of irregularly-spaced data, see Chiu, Eraker, Foerster, Kim, and Seoane (2014). It is unlikely that there would be gains from having one additional weekly observation at irregular intervals in our models, because several alternative timing conventions that we considered generated very similar results.

[3] The spot price data start in January 1985, the oil futures price data for maturities 1–9 months start in June 1984, those for the 12-month maturity start in December 1988, those for the 15-month maturity in June 1989, and those for the 18-month maturity in October 1989.

[4] The gasoline spot price is reported in US dollars per gallon, and is converted to US dollars per barrel by multiplying the price by 42 gallons/barrel to make it compatible with the crude oil price (see Baumeister et al., 2013).

[5] The index is composed of the following companies: Anadarko Petroleum, BP plc, ConocoPhillips, Chevron, Hess, Marathon Oil, Occidental Petroleum, Petr, Phillips 66, Total SA, Valero Energy, and Exxon Mobil.

3. Real-time forecasting models

In this section, we review the forecasting models considered in Section 4. The objective is to forecast the monthly real price of oil using weekly predictors. For expository purposes, it is useful to focus on mixed-frequency VAR (MF-VAR) models first, before discussing MIDAS models.

3.1. MF-VAR forecasts

There are two approaches to estimating the MF-VAR model. One is to estimate the model in state space representation (see, e.g., Schorfheide & Song, 2014), while the other is to stack the weekly predictors in a vector, depending on the timing of their release (see Ghysels, 2012). The main difference between the two is that there are no missing observations in the latter, as the model is estimated at the monthly frequency, and standard estimation methods can be used. We therefore focus on this approach.

3.1.1. MF-VAR model represented as a stacked-vector system

Denote the releases of the weekly variables in the first, second, third and fourth week of each month t by x_t^1, x_t^2, x_t^3 and x_t^4. Define $z_t = [x_t^{w\prime}, x_t^{m\prime}]'$ where $x_t^w = [x_t^{1\prime}, x_t^{2\prime}, x_t^{3\prime}, x_t^{4\prime}]'$, where x_t^m is the vector of monthly variables, including the log of the real price of oil. Then, the variables in the system evolve according to the monthly VAR model

$$A(L)z_t = u_t, \tag{1}$$

where u_t is white noise and $A(L)$ denotes the autoregressive lag order polynomial. The model in Eq. (1) can be estimated by least squares methods, as in the case of a single-frequency VAR model. Forecasts of the real price of oil at monthly horizons $h = 1, \ldots, 24$ may be generated by iterating the recursively estimated VAR model forward, conditional on the date t information set, and converting the forecast of the monthly real price of oil from log-levels to levels.

3.2. Univariate mixed-frequency forecasts

A more parsimonious approach to dealing with mixed-frequency data involves specifying a univariate MIDAS regression. There are three alternative MIDAS representations. Let X_t^w denote a predictor observed in week $w \in \{1, 2, 3, 4\}$ of month t. The weekly predictor may depend on the horizon h of the forecast, in which case we add an additional superscript h. For example, we may define $X_t^{h,w}$ as the cumulative change in X_t^w between the last day of the current week and the last day of the same week h months ago. If the weekly predictor does not depend on h, the superscript h is dropped.

3.2.1. MIDAS regression with estimated weights

The MIDAS model for combining weekly financial predictors with monthly oil price observations is defined as

$$R_{t+h} = R_t \left(1 + \beta B(L^{1/w}; \theta) X_t^w\right) + \varepsilon_{t+h}, \tag{2}$$

where R_t is the current level of the monthly real price of oil. The MIDAS lag polynomial $B(L^{1/w}; \theta)$ is an exponential Almon lag weight function

$$B(L^{1/w}, \theta) = \sum_{j=0}^{3} b(j; \theta) L^{j/w},$$

where the lag operator is defined as

$$L^{j/w}(X_t^w) = X_{t-j/w}^w,$$

and $\theta \equiv \{\theta_1, \theta_2\}$, such that

$$b(j; \theta) = \frac{\exp(\theta_1(j+1) + \theta_2(j+1)^2)}{\sum_{j=0}^{3} \exp(\theta_1(j+1) + \theta_2(j+1)^2)}.$$

Our results are not sensitive to the choice of the exponential Almon lag polynomial. Similar results would be obtained with a beta lag polynomial. The model parameters β and θ are estimated recursively by the method of nonlinear least squares, and forecasts are generated as:

$$R_{t+h|t} = R_t \left(1 + \widehat{\beta} B(L^{1/w}; \widehat{\theta}) X_t^w\right).$$

In some cases, there will be a priori reasons to restrict β to unity, in which case only θ has to be estimated.[6] Restricting β to unity makes sense, for example, when using the oil futures spread to predict changes in the nominal price of oil (see, e.g., Baumeister & Kilian, 2012). This restriction amounts to imposing the absence of a time-varying risk premium.

3.2.2. Equal-weighted MIDAS regressions

An even more parsimonious representation imposes equal weights on the weekly data, resulting in the MIDAS model:

$$R_{t+h} = R_t \left(1 + \beta \sum_{i=0}^{3} \frac{1}{4} X_{t-i/4}^w\right) + \varepsilon_{t+h}. \tag{3}$$

In this case, no estimation is required except for the parameter β. The model is linear in β and may be estimated recursively by ordinary least squares. If β is known, no regression is required and the MSPE of this model may be evaluated using the Diebold and Mariano (1995) test. Our reasons for using equal weights may be explained by appealing to the classical bias–variance tradeoff in forecasting. In small samples, the reduction in the forecast variance from imposing parameters in the MIDAS polynomial may easily outweigh the effects of any misspecification bias. Thus, it makes sense to compare this specification to less restrictive MIDAS specifications.

3.2.3. Unrestricted MIDAS regressions

The issue of whether the added parsimony of the equal-weighted MIDAS model reduces the MSPE is an empirical question. An alternative approach is to relax the restrictions implied by the original MIDAS model. This yields the

6 Note that the MIDAS model does not include an intercept. This fact allows us to nest the random walk forecast without drift. It can be shown that the inclusion of an intercept would systematically lower the forecast accuracy of our MIDAS models.

unrestricted MIDAS (or U-MIDAS) model:

$$R_{t+h} = R_t \left(1 + \sum_{i=0}^{3} \alpha_i X^{w}_{t-i/4}\right) + \varepsilon_{t+h}, \quad (4)$$

which is linear in α_i and can be estimated recursively by ordinary least squares.

3.3. Monthly forecasts

For comparison, we also report results for forecasts from the corresponding monthly forecasting model of the form

$$R_{t+h} = R_t \left(1 + \beta X_t^m\right) + \varepsilon_{t+h},$$

where X_t^m denotes the monthly predictor corresponding to X_t^w. The parameter β is estimated by recursive ordinary least squares. As before, β may be restricted to unity.

4. Empirical results

All forecasts are constructed subject to real-time data constraints. Unknown model parameters are estimated recursively. The estimation period starts as early as data availability allows, and, as a result, differs between models. The earliest starting date of the estimation period is February 1973 and the latest starting date is October 1989. The initial estimation period ends in December 1991, such that, for example, the initial one-month forecast is for January 1992 and the initial 12-month forecast is for December 1992. The estimation period is updated recursively on a monthly basis. The forecast evaluation period ends in September 2012. The use of such a long evaluation period minimizes the danger of spurious forecast successes.

The real oil price forecasts are evaluated in levels against the value of the real price of oil realized in the March 2013 vintage of the real-time data set. We discard the last six observations of the oil price data, as they are still subject to revisions. All forecasts are evaluated based on their MSPE relative to that of the monthly no-change forecast of the level of the real price of oil. MSPE ratios below one indicate that the model in question is more accurate than the no-change forecast. We also report the directional accuracy of the forecasts in the form of the success ratio, defined as the proportion of times that the model in question predicts correctly whether the real price of oil rises or falls. Under the null hypothesis of no directional accuracy, one would expect a success ratio of 0.5. Higher ratios indicate an improvement on the no-change forecast.

While there is no valid test for the statistical significance of the real-time MSPE reductions from models based on estimated MIDAS or U-MIDAS weights, the equal-weighted MIDAS specification with $\beta = 1$ imposed does not suffer from parameter estimation uncertainty, thus allowing the use of the conventional *DM* test of equal MSPEs (see Diebold & Mariano, 1995).[7] The statistical significance of any gains in directional accuracy is evaluated using the test of Pesaran and Timmermann (2009).

4.1. MIDAS results

The set of high-frequency predictors includes (1) the spread between the spot prices of gasoline and crude oil; (2) the spread between the oil futures price and the spot price of crude oil; cumulative percentage changes in (3) the CRB index of the price of industrial raw materials, (4) US crude oil inventories, and (5) the Baltic Dry Index; (6) returns and excess returns on oil company stocks; (7) cumulative changes in US nominal interest rates (LIBOR, Fed funds rate); and (8) cumulative percentage changes in the US trade-weighted nominal exchange rate.

4.1.1. Oil futures prices

Forecasting models based on oil futures prices form a good starting point. In the absence of a risk premium, arbitrage implies that the oil futures price is the conditional expectation of the spot price of oil (see Alquist & Kilian, 2010). Equivalently, in logs, this means that

$$E_t(\Delta s_{t+h}) = f_t^h - s_t, \quad (5)$$

where h is the forecast horizon and the maturity of the futures contract in months. For our sample period, the maximum maturity for which continuous weekly time series of WTI oil futures and spot prices are available is 18 months. Eq. (5) suggests that we express the MIDAS forecasting model for horizon h as a polynomial in $X_t^{h,w} = f_t^{h,w} - s_t^w$, where the spread is measured on the last day of week $w = 1, 2, 3, 4$ of a given month t. We also make an adjustment for expected inflation, which is approximated by the average inflation rate since July 1986, following Baumeister et al. (2013).

Table 1 shows that the equal-weighted MIDAS forecast has lower MSPEs than the no-change forecast at every horizon between one month and 18 months. The gains in accuracy are negligible at horizons under 12 months, but more substantial at longer horizons. The largest reduction in the MSPE is 17% at horizon 15. The MSPE reductions at horizons 12, 15, and 18 are statistically significant based on the *DM* test. There are no statistically significant gains in directional accuracy at short horizons. In fact, some of the success ratios are well below 0.5. Significant improvements in directional accuracy are observed at horizons 9, 12, 15, and 18. The largest success ratio is 63%. Similar results are obtained for the model based on estimated MIDAS weights, and only slightly less accurate results are obtained for the unrestricted MIDAS model.

[7] There are two reasons why we can only assess the statistical significance of the directional accuracy statistics, not of the MSPE reductions. One problem is that all standard tests of equal MSPEs are based on the population MSPE, not the actual out-of-sample MSPE. This means that these tests are not appropriate for our purpose. This point was first made by Inoue and Kilian (2004), and has been accepted widely in recent years. If one uses these tests anyway, one will reject the null of equal MSPEs too often. This point has been illustrated, for example, by Alquist et al. (2013). Clark and McCracken (2012) are working to try to address this issue, but their solutions do not apply in our context. The second problem is that standard tests for equal predictive accuracy do not apply when using real-time data. Clark and McCracken (2009) show how this problem may be overcome in the context of standard tests of no predictability in population. They focus on special cases under additional assumptions, but their analysis does not cover our forecast settings, nor does it address the first problem above.

Table 1
Forecasting the monthly real price of oil using the oil futures spread. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS Equal weights MSPE ratio	MIDAS Equal weights Success ratio	MIDAS Estimated weights MSPE ratio	MIDAS Estimated weights Success ratio	MIDAS Unrestricted MSPE ratio	MIDAS Unrestricted Success ratio	Monthly model MSPE ratio	Monthly model Success ratio
1	**0.996**	0.478	1.000	0.466	1.014	0.466	**0.997**	0.462
3	**0.965**	**0.530**	**0.954**	**0.563**	**0.941**	**0.571**	**0.974**	0.498
6	**0.975**	0.488	**0.964**	**0.508**	**0.980**	0.488	**0.975**	**0.512**
9	**0.938**	**0.568***	**0.922**	**0.564**	**0.939**	**0.568**	**0.944**	**0.589***
12	**0.872***	**0.592***	**0.857**	**0.601***	**0.878**	**0.601***	**0.886****	**0.613***
15	**0.829***	**0.621***	**0.829**	**0.617***	**0.890**	**0.638***	**0.860****	**0.634***
18	**0.848***	**0.629***	**0.854**	**0.625***	**0.962**	**0.625***	**0.906**	**0.621***

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, equal weights
- $R_{t+h|t} = R_t(1 + B(L^{1/4}; \hat{\theta})(X_t^{h,w}) - E_t(\pi_t^h))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, unrestricted
- $R_{t+h|t} = R_t(1 + X_t^h - E_t(\pi_t^h))$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the difference between the log of the oil futures price for maturity h and the log of the spot price of oil in week w of month t, X_t^h is the difference between the log of the oil futures price for maturity h and the log of the spot price of oil in month t, and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test, and for the equal-weighted MIDAS model and the monthly model, are indicated by asterisks, as arestatistically significant reductions in the MSPE according to the Diebold–Mariano test.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

Although the MIDAS model compares favorably with the no-change forecast, so do traditional models based on the most recent monthly oil futures spread. The last two columns of Table 1 show the corresponding results based on the monthly oil futures model, as implemented by Baumeister and Kilian (2012). That model generates broadly similar results, in that the MSPE reductions are statistically significant at horizons 12 and 15, and directional accuracy at horizons 9, 12, 15, and 18. While the equal-weighted MIDAS model has slightly lower MSPEs at all horizons, the monthly forecasting model has slightly higher and more statistically significant directional accuracy at longer horizons. Overall, there is little to choose between these models.

4.1.2. Gasoline price spreads

Petroleum products such as gasoline and heating oil are produced by refining crude oil. Many oil market analysts and financial analysts believe that the prices for these petroleum products contain useful information about the future evolution of the price of crude oil. In particular, changes in the product price spread – defined as the extent to which today's price of gasoline or heating oil deviates from today's price of crude oil – are widely viewed as predictors of changes in the spot price of crude oil. For example, in April 2013 Goldman Sachs cut its oil price forecast, citing significant downward pressure on product price spreads, which it interpreted as an indication of a reduced final demand for products, and hence an expectation of falling crude oil prices (see Strumpf, 2013).

This forecasting approach was formalized and evaluated recently by Baumeister et al. (2013) using monthly data. Their analysis demonstrates that models of the gasoline price spread with an intercept of zero but a freely estimated slope parameter are reasonably successful at predicting the real price of oil at horizons of up to 24 months. In the analysis below, we impose the same restriction. A preliminary analysis with alternative models confirmed that all other specifications are inferior.

Table 2 considers the MIDAS analogue of the model proposed by Baumeister et al. (2013), with X_t^w denoting the spread between the spot price of gasoline and the WTI spot price of crude oil, measured on the last trading day of week $w = 1, 2, 3, 4$ of a given month t. The parameter β is estimated freely. Table 2 shows that this equal-weighted MIDAS model has lower MSPEs than the no-change forecast at every horizon from 1 month to 24 months, but, with a few exceptions, the MSPE reductions are modest. There are no statistically significant gains in directional accuracy. Similar results hold when estimating the MIDAS weights. The unrestricted MIDAS model is somewhat less accurate.

Because of the presence of parameter estimation uncertainty, a proper assessment of the statistical significance of the MSPE reductions in Table 2 is not possible, but we can compare these results with those obtained for the corresponding monthly model, building on the work of Baumeister et al. (2013). The latter model has slightly lower MSPEs at eight of the nine horizons. Both models' directional accuracy is statistically insignificant and erratic. There is no reason to favor either of these models. As was the case with oil futures, there are no clear advantages in the use of the MIDAS model.

4.1.3. CRB index of the spot price of industrial raw materials

There is a long tradition of modelling oil prices jointly with other industrial commodities (e.g., Barsky & Kilian, 2002; Frankel, 2008). The CRB provides a widely used index of the spot price of industrial raw materials excluding crude oil. Alquist et al. (2013) first made the case that cumulative percentage changes in this CRB price index can be viewed as a proxy for the expected cumulative percentage change in the price of oil. The rationale for this statement is that, often, fluctuations in industrial commodity prices

Table 2
Forecasting the monthly real price of oil using the gasoline-crude oil spot price spread. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.993**	**0.578**	**0.998**	**0.590**	1.071	**0.554**	**0.989**	**0.562**
3	**0.996**	**0.583**	1.004	**0.583**	1.019	**0.534**	**0.990**	**0.583**
6	**0.991**	**0.574**	**0.984**	**0.582**	**0.997**	**0.533**	**0.978**	**0.545**
9	**0.984**	0.490	**0.987**	0.494	1.011	0.485	**0.963**	0.436
12	**0.963**	0.441	**0.961**	0.483	**0.964**	**0.555**	**0.934**	**0.521**
15	**0.956**	**0.532**	**0.950**	**0.540**	**0.945**	**0.591**	**0.931**	**0.516**
18	**0.973**	**0.504**	**0.970**	**0.543**	**0.966**	**0.582****	**0.971**	0.470
21	**0.976**	**0.541**	**0.972**	**0.563**	1.003	**0.546**	**0.986**	0.454
24	**0.935**	**0.588**	**0.927**	**0.566**	**0.953**	**0.540**	**0.934**	**0.500**

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}) - E_t(\pi_t^h))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h - E_t(\pi_t^h))$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the difference between the log of the gasoline spot price and the log of the spot price of oil in week w of month t, X_t^h is the difference between the log of the gasoline spot price and the log of the spot price of oil in month t, and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

Table 3
Forecasting the monthly real price of oil using the CRB spot price index of industrial raw materials. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.929**	**0.558****	**0.927**	**0.562**	**0.978**	**0.546**	**0.934**	**0.546****
3	**0.862**	**0.628***	**0.831**	**0.636***	**0.861**	**0.632***	**0.863**	**0.628***
6	1.113	**0.611***	1.112	**0.623***	1.085	**0.570***	1.107	**0.598***
9	1.163	**0.573**	1.158	**0.564**	1.085	0.469	1.143	**0.593***
12	1.132	**0.546**	1.131	**0.546**	1.131	0.454	1.100	**0.592***
15	1.150	**0.574****	1.144	**0.574**	1.131	0.451	1.118	**0.617***
18	1.254	**0.539**	1.252	**0.539**	1.154	0.418	1.232	**0.578***
21	1.382	**0.528**	1.382	**0.528**	1.139	0.445	1.376	**0.528**
24	1.377	**0.513**	1.380	**0.509**	1.172	0.451	1.394	0.443

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, equal weights
- $R_{t+h|t} = R_t(1 + B(L^{1/4}; \hat{\theta})(X_t^{h,w}) - E_t(\pi_t^h))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}) - E_t(\pi_t^h))$, unrestricted
- $R_{t+h|t} = R_t(1 + X_t^h - E_t(\pi_t^h))$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percentage change in the CRB spot price index of industrial raw materials over the preceding h months in week w of month t, X_t^h is the percentage change in the CRB spot price index of industrial raw materials over the preceding h months in month t, and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test, and for the equal-weighted MIDAS model and the monthly model are indicated by asterisks, as are statistically significant reductions in the MSPE according to the Diebold–Mariano test.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

are driven by persistent, and hence predictable, variation in global real economic activity. Several studies have elaborated on this insight and demonstrated that such models have statistically significant directional accuracy and yield statistically significant MSPE reductions for the real price of oil (see Baumeister & Kilian, 2012, 2014a, 2014b).

The CRB index is also available on a daily basis, which allows us to incorporate weekly observations of the cumulative percentage change in this index into a MIDAS model. Consistent with the analysis of Alquist et al. (2013), the MIDAS model is estimated with $\beta = 1$ imposed. Table 3 shows that the equal-weighted MIDAS model has directional accuracy at all horizons and statistically significant directional accuracy at some horizons. This model also reduces the MSPE at short horizons by as much as 14%, but the reductions are never statistically significant based on the *DM* test. At longer horizons, there are no reductions in the MSPE. Similar results are obtained for the MIDAS model

Table 4
Forecasting the monthly real price of oil using the Baltic Dry Index. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS							
	Equal weights		Estimated weights		Unrestricted		Monthly model	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.950**	**0.502**	**0.947**	0.498	**0.991**	0.442	**0.952**	**0.546**[**]
3	1.049	0.470	1.079	0.466	1.078	0.482	1.109	0.462
6	1.015	**0.504**	1.023	**0.512**	1.083	**0.520**	1.056	0.492
9	1.030	**0.502**	1.025	0.498	1.033	0.490	1.124	**0.548**
12	1.087	0.445	1.094	0.445	1.166	0.441	1.447	**0.500**
15	1.123	0.383	1.136	0.387	1.203	0.391	1.544	0.426
18	1.297	0.435	1.308	0.414	1.327	0.397	2.112	0.474
21	1.399	0.341	1.393	0.332	1.397	0.358	2.214	0.411
24	1.391	0.363	1.407	0.363	1.464	0.442	2.185	0.327

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}))$, equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percentage change in the BDI over the preceding h months in week w of month t, and X_t^h is the percentage change in the BDI over the preceding h months in month t. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
[*] Denotes significance at the 5% level.
[**] Denotes significance at the 10% level.

with estimated weights. The unrestricted MIDAS model is somewhat less accurate.

The last entries in Table 3 allow us to compare the performance of the MIDAS model with that of the corresponding model based on the monthly CRB predictor. The MSPE results are very similar and again statistically insignificant, but, overall, the monthly model has a somewhat higher and more statistically significant directional accuracy. We conclude that, in this case, there is no gain from switching to MIDAS models, and the monthly model is preferred.

4.1.4. Baltic Dry Index

The central idea behind using the CRB spot price index for industrial raw materials in forecasting the price of oil is that the real price of oil is predictable, to the extent that the global business cycle is predictable. This is also the motivation for the inclusion of measures of global real economic activity such as the Kilian (2009) index in VAR oil price forecasting models. One limitation of the latter index, and of all other measures of global real economic activity, is that it is not available at daily frequency. While there are daily real-time indices of US real economic activity, such as the business cycle conditions index of Aruoba, Diebold, and Scotti (2009), there are no similar indices with the same global coverage as the monthly Kilian (2009) index.

An alternative business cycle indicator that is used widely by practitioners is the Baltic Dry Index (BDI), which is quoted on a daily basis by Bloomberg. This index is available starting in 1985. The name of this index derives from the fact that it is maintained by the Baltic Exchange in London. The BDI measures the cost of moving bulk dry cargo on representative ocean shipping routes in the world. Because dry bulk cargo consists primarily of materials that serve as industrial raw materials, such as coal, steel, cement, and iron ore, this index is seen in the business world to be an indicator of future industrial production. In short, the BDI is viewed as a real-time leading indicator for the world economy, and is used to predict future economic activity (e.g., Bakshi, Panayotov, & Skoulakis, 2011). This fact also makes it a potentially useful predictor for the real price of oil.

Despite its popularity among practitioners, the BDI differs from other measures of real economic activity based on dry cargo shipping rates, such as the Kilian (2009) index, in several ways. Without further transformations, the BDI is at best a crude proxy for changes in global real economic activity. For the purpose of exploring its predictive content within the MIDAS framework, we focus on the percentage change in the BDI over the last h months, rather than transforming the BDI into a business cycle index. The β parameter is estimated freely.

Table 4 shows that there is little gain in accuracy from including the BDI data. Apart from a negligible reduction in the MSPE at the 1-month horizon, the first two MIDAS models tend to have higher MSPEs than the random walk, and lack directional accuracy at all horizons. The unrestricted MIDAS model is even less accurate. We conclude that there does not appear to be useful predictive information in the BDI data. This result is confirmed by the corresponding monthly regression models. Our findings underscore the importance of transforming the BDI data prior to constructing oil price forecasts.

4.1.5. US crude oil inventories

Economic theory suggests that changes in expectations about the real price of oil, all else equal, are reflected in changes in crude oil inventories (see Alquist & Kilian, 2010). This line of reasoning has led to the development of structural oil market models that model changes in global crude oil inventories explicitly (see Kilian & Lee, 2014, Kilian & Murphy, 2014, Knittel & Pindyck, 2013). Monthly changes in global crude oil inventories have also been shown to have predictive power for the real price of oil (see Alquist et al., 2013). Although such data are not available at weekly frequency, US crude oil inventories are. This fact suggests the inclusion of percentage changes in weekly US crude oil inventories over the most recent h months in

Table 5
Forecasting the monthly real price of oil using US crude oil inventories. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS						Monthly model	
	Equal weights		Estimated weights		Unrestricted			
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	1.000	**0.530**	**0.998**	**0.550***	1.003	0.478	1.001	0.414
3	1.004	**0.599***	1.008	**0.587***	1.018	**0.579***	**0.998**	**0.575****
6	1.007	0.463	1.010	0.463	1.021	0.451	1.018	**0.537**
9	**0.964**	**0.506**	**0.961**	**0.523**	**0.985**	**0.523**	**0.981**	**0.519**
12	**0.922**	**0.559**	**0.910**	**0.584**	**0.929**	**0.588**	**0.926**	**0.534**
15	**0.886**	**0.609****	**0.881**	**0.613****	**0.884**	**0.566**	**0.886**	**0.630****
18	**0.835**	**0.621***	**0.828**	**0.625***	**0.836**	**0.599****	**0.835**	**0.629****
21	**0.688**	**0.712***	**0.686**	**0.729***	**0.690**	**0.734***	**0.681**	**0.716***
24	**0.720**	**0.686***	**0.706**	**0.695***	**0.714**	**0.712***	**0.695**	**0.708***

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}))$, equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta} B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, enrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta} X_t^h)$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percentage change in US crude oil inventories over the preceding h months in week w of month t, and X_t^h is the percentage change in US crude oil inventories over the preceding h months in month t. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

a MIDAS forecasting model for the real price of oil. This approach can be shown to generate slightly more accurate forecasts than expressing crude oil inventories as a fraction of world crude oil production, as per Hamilton (2009), and much more accurate forecasts than constructing the deviation of inventories from a time series trend, as per Ye, Zyren, and Shore (2005).

Table 5 summarizes the results. The MIDAS model based on equal weights with β estimated freely is essentially tied with the no-change forecast at horizons 1, 3 and 6, but reduces the MSPE by up to 28% compared with the no-change forecast at longer horizons. Very similar, but marginally more accurate, results are obtained when the MIDAS weights are estimated. The unrestricted MIDAS model also performs well. Moreover, all MIDAS models have high and statistically significant directional accuracy, especially at longer horizons. The directional accuracy may be as high as 73% in some cases. We conclude that MIDAS models based on weekly observations of cumulative changes in US oil inventories are promising tools for applied oil price forecasters, relative to the no-change forecast.

Compared with the corresponding models based on monthly US inventory data, however, the conclusion is less clear.[8] Table 5 shows that the MIDAS model has slightly higher or slightly lower MSPEs than the monthly model, depending on the horizon. Likewise, there is little to choose between the monthly model and the MIDAS model when it comes to directional accuracy. Both models perform quite well, especially at longer horizons. It is clear that the improved forecast accuracy of the MIDAS model at longer horizons has less to do with the imposition of the MIDAS structure than with the choice of predictor.

4.1.6. Oil-company stock prices

Chen (2014) recently showed that oil-sensitive stock price indices, particularly the stock prices of oil companies, can help to forecast the real price of crude oil at short horizons. Such information is available readily at a daily frequency. Building on the work of Chen (2014), we explore this insight using a MIDAS regression, with X_t^w denoting the weekly return on the NYSE Arca Oil Index, measured on the last day of week $w = 1, 2, 3, 4$ of a given month t. This index includes 13 major international oil and natural gas companies. The parameter β is estimated freely.

The upper panel of Table 6 shows that the MIDAS model with equal weights systematically reduces the MSPE relative to the no-change forecast for horizons of up to 15 months. The largest MSPE reduction is 6% at the one-month horizon. There is also some evidence of directional accuracy, but only the one-month-ahead success ratio is statistically significant. However, when estimating the weights and when estimating the MIDAS model in its unrestricted form, the MSPE ratios deteriorate. Although the MIDAS model with equal weights performs better than the no-change forecast, it is not systematically more accurate than the monthly real-time forecast.[9] There is no reason to prefer one specification over the other.

The lower panel of Table 6 shows that the same ranking of models applies when defining X_t^w as the weekly

8 The monthly forecasting models are estimated recursively on the same estimation period as the MIDAS models.

9 These reductions in the MSPE are considerably lower than those reported by Chen (2014). For example, Chen reported a 22% MSPE reduction at the one-month horizon. These results can be traced to a number of differences. First and most importantly, we are forecasting the real US refiners' acquisition cost for crude oil imports, which is subject to real-time delays and revisions, whereas Chen (2014) focused on the real WTI price, which for the most part is not. This accounts for about two-thirds of the difference in results. The remainder is accounted for largely by the fact that we focus on the monthly average price, as reported by the US Energy Information Administration, rather than the end-of-month price that Chen focuses on.

Table 6
Forecasting the monthly real price of oil using returns on oil stocks. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS Equal weights MSPE ratio	Success ratio	Estimated weights MSPE ratio	Success ratio	Unrestricted MSPE ratio	Success ratio	Monthly model MSPE ratio	Success ratio
	Returns on the NYSE Oil Index							
1	**0.943**	**0.586***	**0.987**	**0.570****	**0.999**	**0.590**	**0.945**	**0.518**
3	**0.952**	**0.567**	**0.970**	**0.575****	**0.972**	**0.567**	**0.951**	**0.547**
6	**0.986**	**0.529**	**0.991**	**0.545**	**0.998**	**0.537**	**0.984**	**0.504**
9	**0.986**	**0.523**	1.000	**0.531**	1.022	**0.560**	**0.989**	**0.531**
12	**0.986**	**0.576**	1.004	**0.571****	1.032	**0.563**	**0.983**	**0.588***
15	**0.991**	**0.515**	**0.999**	**0.528**	1.024	**0.536**	**0.990**	**0.506**
18	1.004	0.496	1.008	0.435	1.026	0.453	1.018	0.483
21	1.003	0.476	1.015	0.463	1.017	0.463	1.007	0.459
24	**0.994**	0.447	1.007	**0.509**	**0.995**	0.496	1.002	0.465
	Excess returns of the NYSE Oil Index relative to the NYSE Composite Index							
1	**0.968**	**0.554***	1.007	**0.538****	1.010	**0.530****	**0.973**	**0.530**
3	**0.982**	**0.518**	**0.998**	**0.518**	1.001	**0.522**	**0.985**	**0.526**
6	**0.993**	0.496	**0.999**	**0.537**	1.003	**0.520**	**0.996**	**0.508**
9	1.002	0.469	1.023	**0.502**	1.033	**0.535**	1.002	0.486
12	1.000	**0.500**	1.019	**0.534**	1.046	**0.521**	**0.998**	**0.517**
15	**0.999**	0.485	1.011	**0.532**	1.048	0.489	1.001	**0.502**
18	1.004	0.478	1.015	0.483	1.037	0.427	1.001	**0.500**
21	1.000	**0.502**	1.015	0.441	1.026	0.450	**0.997**	**0.520****
24	1.003	0.482	1.019	0.491	1.019	0.434	1.001	0.447

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}))$, equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta}B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta}X_t^h)$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the 1-week return (or excess return) on the NYSE Oil Index in week w of month t, and X_t^h is the 1-month return (or excess return) on the NYSE Oil Index in month t. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

excess return on the NYSE Arca Oil Index relative to the NYSE Composite Index, except that the reductions in the MSPE and the improvements in directional accuracy are negligible.

4.1.7. US interest rates

There is an impression among many observers that lower interest rates are associated with looser economic policies, and hence a higher demand for crude oil, and possibly also a lower supply of crude oil. Either way, this argument suggests a predictive relationship between changes in interest rates and changes in the price of oil. This perception has been boosted by studies suggesting that low real interest rates lead to high real commodity prices (see, e.g., Barsky & Kilian, 2002; Frankel, 2008).[10] We investigate this proposition by fitting a MIDAS model for the difference between the interest rate on the last day of the current week and the interest rate h months earlier. We consider two alternative measures of US interest rates: the US federal funds rate and the LIBOR rate. The parameter β is estimated freely.

Table 7 indicates that the approach yields modest MSPE reductions at horizons of 6–18 months for all MIDAS specifications involving the federal funds rate, but typically lacks directional accuracy. The corresponding results for the LIBOR rate are even less favorable, regardless of the specification. A comparison with the corresponding monthly forecasting model shows that very similar or worse results are obtained using monthly data only. Neither forecasting approach appears to be superior to the no-change forecast. This evidence reinforces the skepticism regarding the empirical content of models linking oil price fluctuations to variations in US interest rates. While there is no doubt about the theoretical link in question, its quantitative importance has yet to be established.

4.1.8. Trade-weighted US exchange rate

Another popular view is that fluctuations in the value of the dollar relative to other currencies predict changes in the real price of oil, as it becomes more or less expensive for importers of crude oil abroad to purchase crude oil. Previous studies of this question have found no evidence in monthly data to support this view (see Alquist et al., 2013). Here, we return to this question using MIDAS regression specifications that allow the use of high-frequency measures of cumulative percentage changes in the trade-weighted US nominal exchange rate.

Table 8 shows that none of the MIDAS models produce reductions in the MSPE, although there is some evidence of directional accuracy at selected horizons. Exactly the same

[10] This argument is distinct from the implications of the Hotelling (1931) model of exhaustible resources that the price of oil should grow at the rate of interest. The latter proposition was evaluated and rejected by Alquist et al. (2013).

Table 7
Forecasting the monthly real price of oil using US interest rates. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS						Monthly model	
	Equal weights		Estimated weights		Unrestricted			
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
	Federal funds rate							
1	**0.998**	**0.510**	**0.999**	**0.534****	1.001	**0.502**	**0.998**	0.470
3	1.004	**0.530**	1.004	**0.538**	1.005	**0.526**	1.004	**0.530**
6	**0.969**	0.459	**0.971**	**0.504**	**0.967**	**0.520**	**0.966**	0.459
9	**0.960**	**0.506**	**0.963**	**0.510**	**0.964**	**0.515**	**0.953**	**0.506**
12	**0.952**	**0.504**	**0.947**	0.475	**0.946**	0.483	**0.952**	0.496
15	**0.961**	**0.515**	**0.954**	**0.502**	**0.946**	**0.502**	**0.963**	0.498
18	**0.986**	0.491	**0.982**	0.487	**0.977**	0.487	**0.987**	**0.500**
21	1.011	0.480	1.009	0.472	**0.997**	0.480	1.012	0.489
24	1.032	0.434	1.032	0.442	1.024	0.438	1.032	0.434
	LIBOR							
1	1.006	**0.522**	1.010	**0.526**	1.013	**0.530**	1.037	**0.534****
3	1.017	**0.538****	1.018	**0.506**	1.018	**0.571***	1.023	**0.547**
6	**0.996**	0.463	**0.996**	0.475	1.033	0.496	1.014	0.385
9	**0.994**	0.436	**0.992**	0.461	**0.992**	0.461	1.086	0.486
12	**0.980**	0.458	**0.979**	0.454	**0.986**	0.483	1.050	0.382
15	**0.995**	0.485	**0.994**	0.481	**0.993**	0.485	1.033	0.430
18	1.011	0.457	1.011	0.461	1.008	0.470	1.050	0.457
21	1.033	0.459	1.034	0.454	1.034	0.480	1.083	0.389
24	1.058	0.429	1.060	0.434	1.064	0.460	1.088	0.358

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \hat{\beta}\sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}))$, equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta}B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta}X_t^h)$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the change in the interest rate over the preceding h months in week w of month t, and X_t^h is the change in the interest rate over the preceding h months in month t. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

Table 8
Forecasting the monthly real price of oil using the nominal trade-weighted US exchange rate. Evaluation period: 1992.1–2012.9.

Horizon (months)	MIDAS						Monthly model	
	Equal weights		Estimated weights		Unrestricted			
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	1.005	0.466	1.006	**0.514**	1.018	**0.514**	1.007	0.466
3	1.081	**0.502**	1.078	0.486	1.084	0.490	1.068	0.494
6	1.006	0.426	1.016	0.418	1.038	0.434	**1.000**	0.480
9	1.061	**0.622***	1.070	**0.548**	1.097	**0.523**	1.069	**0.618***
12	1.174	**0.618***	1.188	**0.613***	1.199	**0.592***	1.176	**0.626***
15	1.149	**0.591***	1.147	**0.600***	1.176	**0.600***	1.146	**0.600****
18	1.157	**0.565**	1.163	**0.547**	1.175	**0.543**	1.153	**0.560**
21	1.143	0.459	1.146	0.472	1.163	0.472	1.140	0.463
24	1.079	0.482	1.079	0.451	1.078	0.465	1.078	0.478

Notes: The forecasts are constructed as:
- $R_{t+h|t} = R_t(1 + \hat{\beta}\sum_{i=0}^{3} \frac{1}{4}(X_{t-i/4}^{h,w}))$, equal weights
- $R_{t+h|t} = R_t(1 + \hat{\beta}B(L^{1/4}; \hat{\theta})(X_t^{h,w}))$, estimated weights
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{3} \hat{\alpha}_i(X_{t-i/4}^{h,w}))$, unrestricted
- $R_{t+h|t} = R_t(1 + \hat{\beta}X_t^h)$, monthly model

where R_t is the real price of oil, $X_{t-i/4}^{h,w}$ is the percentage change in the exchange rate over the preceding h months in week w of month t, and X_t^h is the percentage change in the exchange rate over the preceding h months in month t. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

pattern applies to the corresponding monthly model in Table 8. There is some evidence of modest statistically significant directional accuracy at intermediate horizons, but again the MIDAS model has no advantage over the monthly model. We conclude that these models are effectively indistinguishable.

Table 9
Forecasting the monthly real price of oil using daily predictors. Evaluation period: 1992.1–2012.9.

Horizon (months)	Equal-weighted daily MIDAS models							
	Oil futures spread		Gasoline-crude oil spot spread		CRB spot price		Baltic Dry Index	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.997**	**0.522**	**0.995**	**0.578**	**0.934**	**0.554****	**0.995**	**0.514**
3	**0.976**	**0.534**	**0.997**	**0.583**	**0.864**	**0.623***	1.045	0.466
6	**0.978**	0.492	**0.993**	**0.557**	1.119	**0.607***	1.017	**0.504**
9	**0.938**	**0.564**	**0.985**	0.473	1.165	**0.568**	1.031	**0.502**
12	**0.870***	**0.580***	**0.964**	0.450	1.129	**0.546**	1.089	0.437
15	**0.829***	**0.621***	**0.958**	**0.528**	1.147	**0.583**	1.125	0.391
18	**0.846***	**0.638***	**0.975**	**0.530**	1.253	**0.543**	1.300	0.435
21	NA	NA	**0.976**	**0.559**	1.381	**0.528**	1.401	0.336
24	NA	NA	**0.936**	**0.580**	1.378	**0.509**	1.391	0.367
Horizon (months)	Returns on the NYSE Oil Index		Federal funds rate		LIBOR		Nominal trade-weighted US exchange rate	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.896**	**0.614***	**0.998**	**0.514****	1.006	**0.522**	1.009	0.466
3	**0.928**	**0.571**	1.004	**0.567***	1.016	**0.538****	1.061	0.498
6	**0.973**	**0.549**	**0.973**	0.471	**0.996**	0.471	1.012	0.430
9	**0.980**	**0.527**	**0.962**	**0.519**	**0.994**	0.423	1.032	**0.602***
12	**0.979**	**0.576****	**0.952**	0.492	**0.981**	0.454	1.106	**0.630***
15	**0.987**	0.485	**0.962**	**0.502**	**0.995**	0.485	1.088	**0.596***
18	1.002	0.496	**0.987**	0.496	1.011	0.453	1.091	**0.560**
21	1.000	0.450	1.012	0.493	1.034	0.450	1.084	0.467
24	**0.989**	**0.518**	1.032	0.429	1.059	0.425	1.044	0.478

Notes: By analogy to Tables 1–8, the forecasts are constructed from the models:
- $R_{t+h|t} = R_t(1 + \sum_{i=0}^{19} \frac{1}{20}(X_{t-i/20}^{h,d}) - E_t(\pi_t^h))$
- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^{19} \frac{1}{20}(X_{t-i/20}^{h,d}) - E_t(\pi_t^h))$
- $R_{t+h|t} = R_t(1 + \hat{\beta} \sum_{i=0}^{19} \frac{1}{20}(X_{t-i/20}^{h,d}))$

where R_t is the real price of oil, the daily predictor $X_{t-i/20}^{h,d}$ is defined by analogy to the predictor $X_{t-i/4}^{h,w}$ for the last day of the week, and $E_t(\pi_t^h)$ denotes the expected inflation rate over h periods. The benchmark model is the monthly no-change forecast. Bold entries indicate improvements on the no-change forecast. Statistically significant reductions in the MSPE according to the Diebold–Mariano test, where appropriate, and statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.

* Denotes significance at the 5% level.

** Denotes significance at the 10% level.

Moreover, neither model can be recommended for forecasting oil prices, especially compared with some of the models discussed earlier. This result reinforces the conclusions of Alquist et al. (2013) about the lack of predictive content of exchange rates for oil prices. The notion that fluctuations in the trade-weighted US exchange rate lead fluctuations in the real price of oil lacks empirical support.

4.2. Sensitivity analysis

We now show that our main results for the MIDAS model based on predictors measured at weekly frequency are robust to a number of extensions and modifications.

4.2.1. MIDAS models based on daily predictors

With the exception of the US inventory data, many of the predictors used in this paper are also available at daily frequency. A natural question, therefore, is whether our results are robust to applying the MIDAS framework to daily data rather than weekly data. Consider the example of forecasting the price of oil based on cumulative changes in the BDI. The key difference is that the daily MIDAS model is based on cumulative percentage changes in the BDI, measured on each of the 20 business days of a given month, whereas the weekly MIDAS models in Section 3 are based on cumulative percentage changes in the BDI, measured on the last trading day of each week of the month. Table 9 demonstrates that our results are remarkably robust to this change in the MIDAS model specification. For expository purposes, we focus on the equal-weighted MIDAS model. There is no evidence that including all daily observations rather than only the daily observations at the end of each week provides a systematic improvement in forecast accuracy.

4.2.2. Pooling MIDAS forecasts based on weekly data

So far, we have focused on the performance of individual MIDAS models one model at a time. One might expect forecast pooling to provide a further increase in the accuracy of the MIDAS approach. Table 10 illustrates that this is not the case in general. The table summarizes the forecast accuracy of equal-weighted combinations of equal-weighted MIDAS models and of MIDAS models with estimated weights. We focus on equal-weighted forecast combinations, because an additional analysis showed that equal weights generate systematically more accurate pooled forecasts than inverse MSPE weights that are based on recent forecast performances.

Table 10 shows that pooled forecasts generate systematic MSPE reductions relative to the benchmark model at horizons of up to 18 months, but the MSPE reductions are usually quite small. More importantly, the MSPE reductions do not exceed those of the best individual MIDAS

Table 10
Pooled MIDAS forecasts of the monthly real price of oil. Evaluation period: 1992.1–2012.9.

Horizon (months)	Equal-weighted forecast combination			
	Equal-weights MIDAS models		Estimated-weights MIDAS models	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.949**	**0.570***	**0.952**	**0.562****
3	**0.962**	**0.615***	**0.959**	**0.611***
6	**0.983**	**0.516**	**0.982**	**0.557**
9	**0.981**	**0.560**	**0.980**	**0.544**
12	**0.971**	**0.563**	**0.969**	**0.563**
15	**0.963**	**0.545**	**0.961**	**0.553**
18	**0.986**	**0.603***	**0.985**	**0.599***
21	1.018	**0.528**	1.017	**0.524**
24	1.023	0.478	1.022	0.482

Notes: The forecasts underlying the forecast combination are obtained from Tables 1–8. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

models systematically. For example, the equal-weighted MIDAS forecast based on returns on oil company stocks in Table 6 tends to be slightly more accurate than the corresponding pooled forecast in Table 10 at most horizons. Finally, the pooled forecast lacks forecast accuracy at longer horizons. Beyond a horizon of nine months, the MIDAS model based on US inventories is systematically more accurate than the pooled forecast. We conclude that forecast pooling is of limited use in this context.

4.3. MF-VAR results

Despite the availability of numerous high-frequency predictors of the real price of oil, we conclude that only the weekly data on US crude oil inventories stand out as useful predictors of the real price of oil. The surprisingly good performance of the MIDAS model based on US crude oil inventories raises the question of whether even more accurate real-time forecasts could be obtained by incorporating the same weekly inventory data into an MF-VAR model.

Our baseline VAR model includes the percentage change in global crude oil production, a measure of global real activity that was proposed by Kilian (2009), the real price of oil, and the change in global crude oil inventories. This choice of variables is motivated by economic theory (see Kilian & Lee, 2014, Kilian & Murphy, 2014). The model specification is identical to that employed by Baumeister and Kilian (2012), except that the lag order is restricted to two lags, compared to 12 lags in the original analysis. The reason for this is that the MF-VAR model becomes computationally intractable for higher lag orders. By construction, in the MF-VAR(2) model there will be two months' worth of lags of the weekly predictor.

The results shown in Table 11 are obtained based on the stacked vector representation of the mixed-frequency VAR model. Estimating the state space representation of the model as per Schorfheide and Song (2014) yields similar results (which are not shown here, to conserve space). Table 11 illustrates that including weekly US crude oil inventory data in the VAR(2) model does not improve the accuracy of the recursive real-time VAR forecast. In fact, the MF-VAR(2) forecast is slightly less accurate than the original VAR(2) forecast. Either way, the MSPE reductions relative to the no-change forecast are small and do not extend beyond the one-month horizon.

This evidence may seem to suggest that the information conveyed by the US inventory data is already contained in the baseline VAR because of the inclusion of monthly global crude oil inventories. However, the corresponding MIDAS model in Table 5 which does not contain information about global crude oil inventories is much more accurate than the VAR(2) model, especially at longer horizons, which indicates that the more parsimonious MIDAS model structure is what makes the difference. In fact, regardless of which high-frequency predictor is included in the MF-VAR(2) model, the MF-VAR(2) forecasts rarely outperform those of the random walk even at horizon 1, and never beyond horizon 3.[11] Our results demonstrate that MF-VAR models are systematically less accurate than MIDAS models in forecasting the real price of oil in real time.

5. Conclusion

We conclude that the best way of modelling mixed-frequency data in our context involves the use of MIDAS models rather than MF-VAR models. In general, the equal-weighted MIDAS model and the MIDAS model with estimated weights generate the most accurate real-time forecasts based on mixed-frequency data. We found no evidence that unrestricted MIDAS model forecasts are as accurate as or more accurate than forecasts from other MIDAS specifications.

Based on these MIDAS models, we reviewed a wide range of high-frequency financial predictors of the real price of oil. The results can be classified as follows:

- In many cases, the equal-weighted MIDAS model forecasts improve on the no-change forecast, but so does the corresponding forecast from a model including only lagged monthly data, and there is little to choose between the MIDAS model forecast and the forecast from the monthly model. Examples include models that incorporate weekly oil futures spreads, weekly gasoline product spreads, weekly returns on oil company stocks, and weekly US crude oil inventories.

[11] These results are not shown here, to conserve space.

Table 11
VAR and MF-VAR forecasts of the monthly real price of oil. Evaluation period: 1992.1–2012.9.

Horizon (months)	VAR(2)		MF-VAR(2) with weekly US crude oil inventories	
	MSPE ratio	Success ratio	MSPE ratio	Success ratio
1	**0.915**	**0.566***	**0.950**	**0.530**
3	1.007	**0.543**	1.090	**0.522**
6	1.108	**0.553**	1.244	0.459
9	1.224	**0.539**	1.479	0.436
12	1.309	**0.563**	1.630	0.458
15	1.362	**0.549**	1.735	0.447
18	1.426	**0.539**	1.908	0.427
21	1.487	**0.533**	2.064	0.450
24	1.482	**0.518**	2.071	0.465

Notes: The four variables in the VAR model are the growth rate of world oil production, the log of the real price of oil, the Kilian (2009) global real economic activity index, and the change in global crude oil inventories. The weekly US crude oil inventories are expressed as the percentage change over the preceding *h* months. The benchmark model is the monthly no-change forecast. Statistically significant improvements in directional accuracy according to the Pesaran–Timmermann test are indicated by asterisks.
* Denotes significance at the 5% level.
** Denotes significance at the 10% level.

- In some cases, the MIDAS forecast improves on the no-change forecast somewhat, but is inferior in its turn to the corresponding monthly real time forecast. One example is the model incorporating cumulative percentage changes in the weekly CRB spot price index for non-oil industrial raw materials.
- In yet other cases, the MIDAS forecast is about as accurate as the corresponding monthly forecast, but neither is systematically more accurate than the no-change forecast. Examples include models based on cumulative percentage changes in the trade-weighted nominal US exchange rate, in US interest rates, or in the Baltic Dry Index.

Although many MIDAS models improve on the no-change forecast, the only case in which we have documented large and systematic improvements in forecast accuracy involves the inclusion of weekly data on US crude oil inventories in the MIDAS model. The latter specification yields not only impressive reductions in the MSPE at horizons of between 12 and 24 months, but also an unusually high directional accuracy. The largest reduction in the MSPE we observed was 28%, and the largest success ratio was 73%. These gains in real-time forecast accuracy are large compared with those reported in any previous study on forecasting oil prices.

While our analysis has produced strong new evidence that the monthly real price of oil is predictable at horizons beyond one year, this success cannot be attributed to the use of the MIDAS model, because the corresponding forecasting model based on monthly US crude oil inventory data produces similar gains in accuracy. Our analysis suggests that, unlike in many other studies, typically not much will be lost by ignoring high-frequency financial data when forecasting the monthly real price of oil. This is true whether one relies on daily or weekly predictors in the MIDAS model, and even when using forecast combinations.

Throughout the paper, we focused on MIDAS models for one high-frequency predictor at a time. An alternative strategy would have been to impose a factor structure on the set of high-frequency financial predictors, as per Andreou et al. (2013). The latter approach is natural in the context of macroeconomic forecasting, but less appealing in our context, given the much smaller number of potential predictors that can be proposed on economic grounds. The reason for this is that the real price of oil is determined in global oil markets, and the set of relevant global predictors is much smaller.

There are a number of potential extensions of our analysis. For example, although we focused on monthly oil price forecasts, it would have been straightforward to extend our analysis to quarterly horizons. Baumeister and Kilian (2014a, 2014b) show that the best way to generate quarterly forecasts is usually to average monthly forecasts by quarter. One could also extend the analysis to include other oil price measures such as the WTI price. Doing so would raise additional complications, given the instability in the relationship between global oil prices and the WTI price in recent years (see Baumeister & Kilian, 2014a). We therefore focused on the real US refiners' acquisition costs for crude oil imports in this paper, because that price is a widely used proxy for the global price of oil.

Acknowledgments

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada. We thank Frank Schorfheide for providing access to his MF-VAR code, Cynthia Wu for explaining the conventions used in constructing the weekly data used by Hamilton and Wu (2014), and Dmitri Tchebotarev for excellent research assistance. We also thank the editor and three anonymous referees for comments.

References

Alquist, R., & Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, *25*, 539–573.
Alquist, R., Kilian, L., & Vigfusson, R. J. (2013). Forecasting the price of oil. In G. Elliott & A. Timmermann (Eds.), *Handbook of economic forecasting. Vol. 2* (pp. 427–507). Amsterdam: North-Holland.
Andreou, E., Ghysels, E., & Kourtellos, A. (2010). Regression models with mixed sampling frequencies. *Journal of Econometrics*, *158*(2), 246–261.
Andreou, E., Ghysels, E., & Kourtellos, A. (2011). Forecasting with mixed-frequency data. In M. P. Clements & D. F. Hendry (Eds.), *Oxford handbook of economic forecasting* (pp. 225–245). Oxford University Press.
Andreou, E., Ghysels, E., & Kourtellos, A. (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business and Economic Statistics*, *31*, 240–251.

Armesto, M. T., Engemann, K. M., & Owyang, M. T. (2010). Forecasting with mixed frequencies. *Federal Reserve Bank of St. Louis Review*, *92*, 521–536.
Aruoba, S. B., Diebold, F. X., & Scotti, C. (2009). Real-time measurement of business conditions. *Journal of Business and Economic Statistics*, *27*, 417–427.
Bakshi, G., Panayotov, G., & Skoulakis, G. (2011). *The Baltic Dry Index as a predictor of global stock returns, commodity returns, and global economic activity*. Mimeo, University of Maryland.
Barsky, R. B., & Kilian, L. (2002). Do we really know that oil caused the great stagflation? A monetary alternative. In B. Bernanke & K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001* (pp. 137–183).
Baumeister, C., & Kilian, L. (2012). Real-time forecasts of the real price of oil. *Journal of Business and Economic Statistics*, *30*, 326–336.
Baumeister, C., & Kilian, L. (2014a). What central bankers need to know about forecasting oil prices. *International Economic Review*, *55*(3), 869–889.
Baumeister, C., & Kilian, L. (2014b). Forecasting the real price of oil in a changing world: a forecast combination approach. *Journal of Business and Economic Statistics* (in press).
Baumeister, C., Kilian, L., & Zhou, X. (2013). *Are product spreads useful for forecasting oil prices? An empirical evaluation of the Verleger hypothesis*. Mimeo, University of Michigan.
Bernard, J.-T., Khalaf, L., Kichian, M., & Yelou, C. (2013). *On the long-term dynamics of oil prices: learning from combination forecasts*. Mimeo, Carleton University.
Chen, S. S. (2014). Forecasting crude oil price movements with oil-sensitive stocks. *Economic Inquiry*, *52*, 830–844.
Chiu, C. W., Eraker, B., Foerster, A. T., Kim, T. B., & Seoane, H. D. (2014). Estimating VARs sampled at mixed or irregularly spaced frequencies: a Bayesian approach. *Journal of Financial Econometrics* (in press).
Clark, T. E., & McCracken, M. W. (2009). Tests of equal predictive ability with real-time data. *Journal of Business and Economic Statistics*, *27*, 441–454.
Clark, T. E., & McCracken, M. W. (2012). *Nested forecast model comparisons: a new approach to testing equal accuracy*. Mimeo, Federal Reserve Bank of St. Louis.
Clements, M. P., & Galvao, A. B. (2008). Macroeconomic forecasting with mixed-frequency data: forecasting US output growth. *Journal of Business and Economic Statistics*, *26*, 546–554.
Clements, M. P., & Galvao, A. B. (2009). Forecasting US output growth using leading indicators: an appraisal using MIDAS models. *Journal of Applied Econometrics*, *24*, 1187–1206.
Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, *13*, 253–263.
Fattouh, B., Kilian, L., & Mahadeva, L. (2013). The role of speculation in oil markets: what have we learned so far? *Energy Journal*, *34*, 7–33.
Foroni, C., Ghysels, E., & Marcellino, M. (2013). Mixed-frequency vector autoregressive models. *Advances in Econometrics*, *32*, 247–272.
Foroni, C., Marcellino, M., & Schumacher, C. (2014). U-MIDAS: MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society, Series A* (in press).
Frankel, J. (2008). The effect of monetary policy on real commodity prices. In J. Campbell (Ed.), *Asset prices and monetary policy* (pp. 291–327). University of Chicago Press.
Ghysels, E. (2012). *Macroeconomics and the reality of mixed frequency data*. Mimeo, University of North Carolina.
Ghysels, E., & Wright, J. H. (2009). Forecasting professional forecasters. *Journal of Business and Economic Statistics*, *27*, 504–516.
Hamilton, J. D. (2008). Daily monetary policy shocks and new home sales. *Journal of Monetary Economics*, *55*, 1171–1190.
Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity*, *1*, 215–261.
Hamilton, J. D., & Wu, J. C. (2014). Risk premia in crude oil futures prices. *Journal of International Money and Finance*, *42*, 9–37.
Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy*, *39*, 137–175.
Inoue, A., & Kilian, L. (2004). In-sample or out-of-sample tests of predictability: which one should we use? *Econometric Reviews*, *23*, 371–402.
Kilian, L. (2009). Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *American Economic Review*, *99*, 1053–1069.
Kilian, L., & Lee, T. K. (2014). Quantifying the speculative component in the real price of oil: the role of global oil inventories. *Journal of International Money and Finance*, *42*, 71–87.
Kilian, L., & 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.
Knittel, C. R., & Pindyck, R. S. (2013). *The simple economics of commodity price speculation*. Mimeo, MIT.
Pesaran, M. H., & Timmermann, A. (2009). Testing dependence among serially correlated multicategory variables. *Journal of the American Statistical Association*, *104*, 325–337.
Schorfheide, F., & Song, D. (2014). Real-time forecasting with a mixed frequency VAR. *Journal of Business and Economic Statistics* (in press).
Strumpf, D. (2013). Goldman cuts the near-term Brent crude forecast to $100 a barrel. *Wall Street Journal*, April 23.
Ye, M., Zyren, J., & Shore, J. (2005). A monthly crude oil spot price forecasting model using relative inventories. *International Journal of Forecasting*, *21*, 491–501.

Christiane Baumeister is a Principal Researcher in the International Economic Analysis Department at the Bank of Canada. She received her Ph.D. in Economics from Ghent University in 2010.

Pierre Guérin is a Senior Analyst in the International Economic Analysis Department at the Bank of Canada. He received his Ph.D. in Economics from the European University Institute in 2011.

Lutz Kilian is Professor of Economics at the University of Michigan. He received his Ph.D. in Economics from the University of Pennsylvania in 1996 and his M.A. in Development Banking from The American University in 1988. Prior to his Ph.D., he worked for the Inter-American Development Bank in Washington, DC. During 2001–03 he served as an adviser to the European Central Bank in Frankfurt/M., Germany.