

Can Oil Prices Forecast Exchange Rates?

Domenico Ferraro, Ken Rogoff and Barbara Rossi

Duke University Harvard University Duke University

March 29, 2011

Abstract

This paper investigates whether oil price shocks have a reliable and stable out-of-sample relationship with the Canadian/U.S Dollar nominal exchange rate. Despite state-of-the-art methodologies and clean data, we find paradoxically little systematic relation between oil prices and the exchange rate, especially if one takes the monthly and quarterly frequencies into account. In contrast, the very short term relationship between oil prices and exchange rates at the daily frequency is rather robust, and holds no matter whether we use contemporaneous (realized) or lagged oil price shocks in our regression. However, the short-term out-of-sample predictive ability is ephemeral, and it mostly appears after time variation in the forecasting ability of the models has been appropriately taken into account. We show that similar results hold for other currencies and commodity price shocks.

J.E.L. Codes: F31, F37, C22, C53.

Acknowledgments: We would like to thank K. Sheppard and D. Wang for providing data related to this project, as well as L. Kilian and J. Wright for comments.

1 Introduction

In this paper, we focus on a particular commodity price, namely oil prices, to predict the fluctuations in the U.S.-Canada's nominal exchange rates in a pseudo out-of-sample forecast experiment. Our results suggest that, despite incredibly refined and clean data, we find paradoxically little systematic relation between oil prices and the exchange rate, especially if one takes the monthly and quarterly frequencies into account. In contrast, the very short term relationship between oil and exchange rates is rather robust. The novelty of our approach is to consider data at daily frequencies, that capture the contemporaneous short-run movements in these variables, as well as to allow for time variation in the relative performance of the models. Our results indicate that the contemporaneous, realized oil price shocks do predict daily nominal exchange rates between Canada and the U.S., and their predictive ability is strongly significant. On the other hand, the predictive ability of the lagged realized oil price shocks is more ephemeral, and allowing for time variation in the relative performance is crucial to show that lagged commodity price shocks are statistically significant predictors of exchange rates out-of-sample. It is noteworthy that, although in-sample fit is stronger in monthly and quarterly data than in daily data, the out-of-sample predictive ability result breaks down for monthly or quarterly data, thus suggesting that not only the predictive ability is transitory, but also that the effects of shocks to oil prices on exchange rates are short-lived and that the frequency of the data is crucial to capture them.

Although the main focus is on the Canadian-U.S. dollar exchange rate and oil prices, due to the availability of data and its importance in the press,¹ we demonstrate that similar results hold for other commodity prices/exchange rates. In particular, for the Norwegian Krone-U.S. dollar exchange rate and oil prices, we find significant predictive ability of both contemporaneous and lagged oil prices. Similar results hold for the South African Rand-U.S. dollar exchange rate and gold prices. For the Australian-U.S. dollar and oil prices and the Chilean Peso-U.S. dollar exchange rate and copper prices, we find strong and significant predictive ability only with contemporaneous commodity prices as predictors.² Our result

¹For example, see the Wall Street Journal ("Canadian Dollar Slumps, Weighed Down By Softer CPI, Oil Prices", January 25, 2011, at <http://online.wsj.com/article/BT-CO-20110125-714898.html>) and "Canadian Dollar Foreign Exchange Pushes Higher on Oil Prices" at <http://www.foreignexchangeservice.co.uk/foreign-exchange-america/canada/01/2011/canadian-dollar-foreign-exchange-rate-pushes-higher-on-oil-prices.html>.

²Note, however, that the weight of oil on the Canadian commodity price index is between 20 and 25% (source: IMF), for Norway is about 20% (source: Statistics Norway), whereas for Australia it is only 4%

holds for in-sample daily data as well. We conjecture that the mechanism leading to this result is the fact that, for a small open economy exporting oil, the exchange rate should reflect fluctuations in oil prices (see Obstfeld and Rogoff, 1996).³ The effects of shocks to oil prices are immediately translated in changes in exchange rates and are very short-lived. This sheds light on why our out-of-sample forecasts are significant in daily data but not at monthly or quarterly frequencies.

To further study the link between oil prices and exchange rates, in addition to a simple regression of exchange rates on oil prices, we consider the asymmetric model by Kilian and Vigfusson (2009) as well as a threshold model where the oil price has asymmetric effects on the nominal exchange rate. Both the asymmetric and threshold model do not provide significantly better forecasts than the simple benchmark model. This result seems to suggest that, as in Kilian and Vigfusson (2009), asymmetries are not too relevant.

Our empirical results are noteworthy and provide clear evidence of a short term relationship between oil price shocks and exchange rate fluctuations, somewhat parallel to the very high frequency relationship people have found between unanticipated Federal Reserve interest rate / macroeconomic announcements and exchange rates (see Andersen et al., 2003, Faust et al., 2007 and Kilian and Vega, 2008). Our paper is clearly also related to the literature on using commodity prices (in particular, oil prices) to predict exchange rates. In particular, in a very recent paper Chen, Rogoff and Rossi (2010) find that exchange rates predict commodity prices both in-sample and out-of-sample; however, the out-of-sample predictive ability in the reverse direction (namely, the ability of the commodity price index to predict nominal exchange rates) is not strong at the quarterly frequency that they consider. Other papers have considered oil prices or more general commodity prices as exchange rate determinants, but mostly as an *in-sample* explanatory variables for *real* exchange rates. Amano and Van Norden (1998a,b) consider the in-sample relationship between real oil prices and the real exchange rate; Chen and Rogoff (2003) consider instead commodity price indices and find

(source: RBA statistics).

³In particular, it is likely that changes in oil prices depend on international factors outside the control of Canada, since oil is an internationally traded commodity and since Canada is a small country relative to other countries (such as the U.S. or Europe). Since the demand for oil is inelastic in the short run, an increase in oil prices causes the dollar value of the oil sold to rise. Canada is a major exporter of oil to the United States, and in principle, changes to the price of oil should have an impact on currency market; for example, since oil is usually quoted in U.S. dollars on the international markets, Canadian exporters will demand Canadian Dollars in exchange for U.S. dollars, thus appreciating the Canadian dollar relative to the U.S. Dollar. See also Backus and Crucini (2000).

in-sample empirical evidence in favor of their explanatory power for real exchange rates⁴ – see Alquist, Kilian and Vigfusson (2011) for a review of the literature on forecasting oil prices and Obstfeld (2002) for a discussion on correlation between nominal exchange rates and export price indices.

More generally, our paper is related to the large literature on predicting nominal exchange rates using macroeconomic fundamentals.⁵ In particular, empirical evidence in favor of predictive ability of macroeconomic fundamentals has been found mainly at longer horizons (see Mark, 1995; Chinn and Meese, 1995; and Engel, Mark and West, 2007), although inference procedures have been called into question (see Kilian, 1999; Berkowitz and Giorgianni, 2001; Rogoff, 2007; and Rossi, 2005, 2007). Cheung, Chinn and Pascual (2005) also concluded that none of the fundamentals outperform the random walk, although there is empirical evidence that models with Taylor rule fundamentals may have some predictive ability (Wang and Wu, 2008, Molodtsova and Papell, 2009; and Molodtsova, Nikolsko-Rzhevskyy and Papell, 2008). See also Faust, Rogers and Wright (2003), Kilian and Taylor (2003) and Engel, Mark and West (2007) for additional empirical evidence on predictive ability at longer horizons. Our paper focuses instead on short-horizon predictive ability, for which the empirical evidence in favor of the economic models has been more controversial, and shows that oil prices contain valuable information for predicting exchange rates *out-of-sample* in a country that is a significant oil exporter. Short horizon predictive ability has never been convincingly demonstrated in the literature, especially with the high statistical significance levels that we are able to find. Our result is rather the opposite of what is commonly found in the literature: we do find predictive ability using daily data, which disappears at longer horizons. Our paper is also related to Faust, Rogers and Wright (2003), who pointed out that predictive ability is easier to find in real-time data: our paper focuses only on real-time data, but uses an economic fundamental that is very different from the traditional fundamentals used in their paper (such as output, prices, money supply and the current account).

⁴Note that our paper significantly extends the scope of Chen and Rogoff (2003) by showing that oil prices have significant predictive ability in forecasting *nominal* exchange rates *out-of-sample*. Chen and Rogoff (2003) find a stronger *in-sample* correlation when using a non-energy price index, but this data is not available at daily frequencies.

⁵Since the seminal works by Meese and Rogoff (1983a,b, 1988), the literature has yet to find convincing empirical evidence that there exist standard macroeconomic fundamentals, such as interest rate differentials or income differentials, which are reliable predictors for exchange rate fluctuations. See for example Mark, Engel and West (2007), Rogoff (2007) and Rogoff and Stavrakeva (2008). Predictive ability, when it exists, is unstable over time (see Rossi, 2006, and Giacomini and Rossi, 2010).

The paper is organized as follows. Section 2 presents a description of the data. Section 3 shows our main empirical results for the contemporaneous oil price model, and Section 4 reports results for the lagged oil price model. Section 5 extends the analysis to other commodity prices and currencies, and Section 6 presents the empirical results for more general oil price models that allow for asymmetries and threshold effects. Section 7 concludes.

2 Data Description

Our study focuses on Canada for three reasons. The first is that crude oil represents the 21.4 percent of Canada's total exports over the period 1972Q1-2008Q1. The second is that Canada has a sufficiently long history of market-based floating exchange rate. Finally, Canada is a small-open economy whose size in the world oil market is relatively small to justify the assumption that it is a price-taker in that market. For the latter reason, crude oil price fluctuations serve as an observable and essentially exogenous terms-of-trade shock for the Canadian economy.

We use data on Canadian/U.S. dollar nominal exchange rates, oil prices, and Canadian and U.S. interest rates. The oil price series is the spot price of the West Texas Intermediate crude oil. West Texas Intermediate (WTI) is a type of crude oil used as a benchmark in oil pricing and the underlying commodity of New York Mercantile Exchange's oil futures contracts. The Canadian/U.S. dollar nominal exchange rate is from Barclays Bank International (BBI). Data at daily, monthly and quarterly frequency are end-of-sample. More precisely, we follow the end of sample data convention from Datastream: the monthly observation is the observation in the first day of the month, whereas the quarterly observation is the observation on the first day of the second month of the quarter. It is worth to recall that, while the previous literature focuses on monthly and quarterly frequencies, our study switches the focus to daily data, and provides a clean comparison of the results for the three frequencies. The data sample ranges from 12/14/1984 to 11/05/2010. The daily dataset contains 6756 observations, the monthly dataset 311, and the quarterly dataset 104. We acknowledge the availability of quarterly data for the Canadian/U.S. dollar nominal exchange rate since the early seventies, but we restrict our sample for the sake of comparison across frequencies.

To construct the daily Canada-U.S. interest rates differential data, we subtract the daily U.S. short-term interest rate from the daily Canadian short-term rate. The Canadian short-term interest rate is the daily overnight money market financing rate and the U.S. short-term rate is the daily Federal funds effective rate. The series of the daily Canadian overnight

money market financing rate is from Bank of Canada, whereas the series of the Federal funds rate is from the Board of Governors of the Federal Reserve System. From the daily data, we construct the monthly and quarterly series: the monthly observation is the observation of the first day of the month and the quarterly observation is the observation of the second month of the quarter.

We also extend our analysis to other currencies and commodities. The original series for the Norwegian krone/U.S., South African rand/U.S. dollar and Australian dollar/U.S. dollar nominal exchange rates are from Barclays Bank International (BBI). The series for the Chilean peso/U.S. dollar exchange rate is instead from WM Reuters (WMR). Beside the oil price series described above, we use prices for copper and gold. All commodity prices and exchange rates series are obtained from Datastream.

3 Can Oil Prices Forecast Exchange Rate Movements?

In this section, we analyze the relationship between oil prices and exchange rates by evaluating whether oil prices have predictive content for future exchange rates. We first show that oil prices have significant predictive content in out-of-sample forecasts in daily data. The predictive content however is much weaker at monthly frequencies, and completely disappears at quarterly frequencies.

The finding that oil prices do forecast nominal exchange rates overturns an important conventional result in the literature, namely the fact that nominal exchange rates are unpredictable. It is therefore crucial to understand the reasons why we find the predictability. We will show that: (i) predictability is very short-lived: it appears at daily frequencies, but is much weaker at monthly frequencies and nonexistent at quarterly frequencies; (ii) the predictability at daily frequencies is specific to oil prices and does not extend to other traditional fundamentals such as interest rates; (iii) predictability is extremely reliable, in the sense that it does not depend on the sample period; (iv) in addition, we verify that the predictability is present not only out-of-sample but also in-sample. While this section focuses on contemporaneous predictive content of oil prices, based on realized oil price shocks as predictors in the out-of-sample forecasting exercise, the next section verifies the robustness of the results to actual ex-ante predictive content by using lagged oil price shocks as predictors.

3.1 Out-of-Sample Forecast Analyses With Realized Fundamentals

We first assess the out-of-sample predictive ability of oil prices. We focus on the simplest oil price model:

$$\Delta s_t = \alpha + \beta \Delta p_t + u_t, \quad t = 1, \dots, T, \quad (1)$$

where Δs_t and Δp_t are the first difference of the logarithm of respectively the Canadian/U.S. dollar exchange rate⁶ and the oil price, T is the total sample size, and u_t is an unforecastable error term. Notice that the realized right-hand-side variable is used for prediction. In the forecasting literature such “ex-post” forecasts are made when one is not interested in ex-ante prediction but in the evaluation of predictive ability of a model given a path for some un-modelled set of variables – see West (1996). An important example of the use of such technique is Meese and Rogoff (1983). Meese and Rogoff (1983a,b, 1988) demonstrated that even using realized values of the regressors, traditional fundamentals such as interest rates, monetary or output differentials would have no predictive power for exchange rates. One of the objectives of this paper is to show that the use of a different fundamental, namely oil prices, can overturn this important finding at daily frequencies; we therefore use the same forecasting strategy.

We estimate the parameters of the model with rolling in-sample windows and produce a sequence of 1-step ahead pseudo out-of-sample forecasts conditional on the realized value of the commodity prices. Let Δs_{t+1}^f denote the one-step ahead pseudo out-of-sample forecast:

$$\Delta s_{t+1}^f = \hat{\alpha}_t + \hat{\beta}_t \Delta p_{t+1}, \quad t = R, R + 1, \dots, T - 1$$

where $\hat{\alpha}_t, \hat{\beta}_t$ are the parameter estimates obtained from a rolling sample of observations $\{t - R + 1, t - R + 2, \dots, t\}$, where R is the in-sample estimation window size. As discussed above, the pseudo out-of-sample forecast experiment that we consider utilizes the realized value of the change in the oil price as a predictor for the change in the exchange rate. The reason is that it is very difficult to obtain a model to forecast future changes in the oil price, since they depend on political decisions and unpredictable supply shocks. If we were to use past values of oil prices in our experiment, and the past values of oil prices were not good forecasts of future values of oil prices, we would end up rejecting the predictive ability of oil prices even though the reason for the lack of predictive ability is not the absence of a

⁶The value of the Canadian/U.S. exchange rate is expressed as the number of U.S. dollars per unit of Canadian dollars.

relationship between exchange rates and oil prices, but the poor forecasts that lagged price changes have for future price changes. To avoid this problem, we condition the forecast on the realized future changes in oil prices. It is important to note, however, that our exercise is not a simple in-sample fit exercise: we attempt to fit future exchange rates out-of-sample, which is a notably difficult enterprise.⁷

We compare the oil price-based forecasts with those of the random walk, which, to date, is the toughest benchmark to beat. We implement the Diebold and Mariano (1995)'s test of equal predictive ability, comparing the oil price model's forecasts with those of two benchmarks: the random walk with and without drift.⁸ We test the null hypothesis of equal predictive ability with daily, monthly and quarterly data. Figure 1(a) depicts the Diebold and Mariano (1995) test statistic for daily data computed with varying in-sample estimation window sizes. The size of the in-sample estimation window relative to the total sample size is reported on the x-axis. When the Diebold and Mariano (1995) statistic is less than -1.96, we conclude that the oil price model forecasts better than the random walk benchmark. The figure shows that, no matter the size of the in-sample window, the test strongly favors the model with oil prices.⁹ This result holds for both benchmarks: the random walk without drift (solid line with circles) and with drift (solid line with diamonds). Overall, we conclude that daily data show extremely robust results in favor of the predictive ability of the oil price model.¹⁰

INSERT FIGURE 1(a) HERE

⁷Our use of the terminology "out-of-sample fit" emphasizes that this section uses "realized" rather than "forecasted" out-of-sample predictors. The next section will focus on the more challenging exercise of using lagged values of the predictors, which will reflect actual out-of-sample forecasting ability.

⁸Even though our models are nested, we can use the Diebold and Mariano (2005) test for testing the null hypothesis of equal predictive ability at the estimated (rather than pseudo-true) parameter values – see Giacomini and Rossi (2009) for a discussion. Using the Clark and West's (2006) test, designed for nested models comparisons, would only strengthen our results in favor of the economic models.

⁹Differences in the value of the test statistics may come from either changing the size of the in-sample window or from considering different sub-samples of data. This issue is addressed later in the paper.

¹⁰Note that the MSFE ratio between the model and the random walk without drift is 0.94 for R=1/2, 0.93 for R=1/3 and 0.91 for R=1/5. Thus, the improvement in forecasting ability is non-negligible in economic terms. The MSFE of the random walk without drift is $3.2976 \cdot 10^{-5}$ for R=1/2, $2.6626 \cdot 10^{-5}$ for R=1/3 and $2.3396 \cdot 10^{-5}$ for R=1/5.

3.2 Why Are We Able to Find Predictive Ability?

Our empirical results greatly differ from the existing literature in two crucial aspects. First, we consider an economic fundamental for nominal exchange rates that is very different from those commonly considered in the literature, namely oil prices. Second, we focus on a different data frequency, daily rather than monthly or quarterly. Therefore, it is important to understand whether it is the frequency of the data or the nature of the fundamental that drives our results.

In a first experiment, we consider the model with oil prices but at the monthly and quarterly frequencies. Figure 1(b) shows Diebold-Mariano's (1995) test statistics for monthly and quarterly data, respectively. For quarterly data, we are never able to reject the null hypothesis of equal predictive ability. For monthly data, we find empirical evidence in favor of the model with oil prices, although the significance is much lower than that of daily data. Since previous research focused only on either monthly or quarterly data, this may explain why the existing literature never noticed the out-of-sample predictive ability in oil prices.

INSERT FIGURE 1(b) HERE

In a second experiment, we consider a model with traditional fundamentals. Traditional fundamentals include interest rate, output and money differentials (see Meese and Rogoff, 1983a,b, 1988, and Engel, Mark and West, 2007). Since output and money data are not available at the daily frequency, we focus on interest rate differentials. That is, we consider the interest rate model:

$$\Delta s_t = \alpha + \beta \Delta i_t + \mu_t \quad (2)$$

where Δi_t are the first difference of the interest rate differential between Canada and the U.S., and μ_t is an unforecastable error term.

Figure 2 reports the results. Panel A in Figure 2 shows that the interest rate model never forecasts better than the random walk benchmark; if anything, the random walk without drift benchmark is almost significantly better. Panels B and C show that similar results hold at the monthly and quarterly frequency.

We conclude that the most likely reason why we are able to find predictive ability is the new fundamental that we consider (the oil price) rather than the frequency of the data.

INSERT FIGURE 2 HERE

3.3 Instabilities in Forecast Performance

The existing literature on the effects of oil price shocks on the economy points to the existence of instabilities over time – see Mork (1989), and Hooker (1996). In particular, Mork (1989) found that the behavior of GNP growth is unstable and indeed correlated with the state of the oil market. Hooker (1996) provided sub-sample analyses and also found empirical evidence of structural instability. To evaluate whether potential instabilities may affect the forecast performance of the oil price model, we report results of the Fluctuation test proposed by Giacomini and Rossi (2010). The latter suggests to report rolling averages of (standardized) Mean Squared Forecast Error differences over time to assess whether the predictive ability changes over time. The in-sample estimation window is one-half of the total sample size and the out-of-sample period equals five hundred days. Panel A in Figure 3 shows the Fluctuation Test for daily data. The figure plots the relative performance (measured by the Diebold and Mariano's (1995) statistics) for the oil price model (eq. 1) against the random walk without drift (solid line with circles) and with drift (solid line with diamonds), together with the 5% critical values (solid lines). Since the values of the statistic are below the (negative) critical value, we reject the null hypothesis of equal predictive ability at each point in time, and conclude that the oil price model forecasts better in some periods. Visual inspection of the graph suggests that the oil price model performs significantly better than the random walk after 2005. Panels B and C in Figure 3 show the results of the Fluctuation test for monthly and quarterly data. For monthly and quarterly data the in-sample window size is the same as in daily data and equals one-half of the total sample, whereas the out-of-sample window is chosen to be the same across frequencies. At the monthly and quarterly frequencies we do not detect significant predictive ability improvements of the oil price model over the random walk.

INSERT FIGURE 3 HERE

3.4 In-sample Fit Analysis

To assess whether the out-of-sample predictive ability is related to the in-sample fit of the models, we estimate the oil price model, eq. (1), over the entire sample period with daily, monthly and quarterly data. Table 1 shows empirical results on the fit of the model. The constant α , is never statistically significant. The coefficient on the growth rate of the oil price β , instead, is statistically significant at any standard level of significance, and for all frequencies. The in-sample fit of the model (measured by the R^2) improves when considering

quarterly data relative to monthly and, especially, daily data. Comparing these results with those in the previous section, interestingly, it is clear that the superior in-sample fit at monthly, and especially quarterly, frequencies does not translate into superior out-of-sample forecasting performance.

INSERT TABLE 1 HERE

In order to assess whether there is time variation in the in-sample measures of fit, we next estimate the oil price model over rolling sub-samples, fixing the rolling-sample size at one-half of the total sample. This produces a sequence of R^2 's and t-statistics for the slope coefficient on the oil price growth rate. Panel A in Figure 4(a) shows the results for daily data. There is clear evidence of time variation in the in-sample predictive ability of oil prices. The slope coefficient on oil prices is statistically significant at the 5% significance level in most of the sample, but becomes increasingly significant towards the end of the sample. This clearly shows that oil prices have become increasingly more important for the Canadian/U.S. dollar exchange rate over the sample period. The same pattern holds for the R^2 measure of fit, reported in Panel A in Figure 4(b). The R^2 remains roughly constant for most of the sample period, but starts to increase over the years 2004/2005 and reaches a value of approximately 0.14 at the end of the sample.

INSERT FIGURES 4(a,b) HERE

Panel B in Figure 4(a) shows that similar results hold for monthly data. The R^2 of the rolling regressions vary over time and tends to increase towards the end of the sample. Note that the R^2 at the end of the sample period climb to values around 0.20. Regarding the slope coefficient on oil price growth, it is not statistically different from zero for most of the early sample period, but it becomes significant towards the end of the sample. This is mirrored by an increasing R^2 over the same time period. It is also worth to note that is the high significance at the end of the sample that could drive the overall significance shown in Table 1. The results are similar for quarterly data, reported in Panel C of Figure 4(a).

The main conclusion that we can draw from the in-sample analysis is that the frequency of the data does not matter for in-sample analysis, at least when we evaluate the oil price model over the full sample. When we look at regressions over rolling sub-samples, we find mixed evidence. The common pattern across frequencies is the increasing importance of oil prices in explaining in-sample variation of the Canadian/U.S. dollar exchange rate. However,

while the oil price is statistically different from zero over most of the sample period for daily data, for monthly and quarterly data the significance is more sporadic.¹¹

INSERT TABLE 1(b) HERE

4 Can Lagged Oil Prices Forecast Exchange Rates?

The previous section focused on regressions where the realized value of the oil price shock is used to predict exchange rates contemporaneously. In reality, forecasters would not have access to realized values of oil price shocks when predicting future exchange rates. So, while the results in the previous section are important to establish the existence of a stronger link between oil price shocks and exchange rates in daily data (relative to monthly and quarterly data), they would not be useful for practical forecasting purposes. In this section, we consider a stricter test by studying whether lagged (rather than contemporaneous) oil price shocks have predictive content for future exchange rates. We first show that the predictive ability now depends on the estimation window size, being more favorable to the model with lagged oil prices only for large in-sample estimation window sizes. We also find that the predictive ability is now more ephemeral, pointing to strong empirical evidence of time variation in the relative performance of the model with lagged oil prices relative to the random walk benchmark only for short periods of time. Only once that time variation is taken into account, we can claim that the model with lagged oil prices significantly forecasts better than the random walk benchmark at the daily frequency. On the other hand, the same model at the monthly and quarterly frequencies never forecasts significantly better than the random walk. Also, using lagged interest rates never improves the forecasting ability relative to the random walk (with or without drift). The empirical evidence in favor of the model with lagged daily oil prices clearly demonstrates that it is important not only to consider daily frequencies but also to allow for the possibility that the relative forecasting performance of the models is time varying, as the predictive ability is very transitory.

4.1 Out-of-Sample Forecast Analyses With Lagged Fundamentals

We focus on the following model with lagged oil prices:

$$\Delta s_t = \alpha + \beta \Delta p_{t-1} + u_t, \quad t = 1, \dots, T, \quad (3)$$

¹¹Table 1(b) reports in-sample estimates of the interest rate model, eq. (2). The coefficient on the interest rate is never significant at any of the frequencies.

where Δs_t and Δp_t are the first difference of the logarithm of respectively the Canadian/U.S. dollar exchange rate and the oil price, T is the total sample size, and u_t is an unforecastable error term. Notice that the lagged value of the right-hand-side variable is used for prediction in eq. (3), whereas the realized value of the explanatory variable was used in eq. (1).

We estimate the parameters of the model with rolling in-sample windows and produce a sequence of 1-step ahead pseudo out-of-sample forecasts conditional on the lagged value of commodity prices. Let Δs_{t+1}^f denote the one-step ahead pseudo out-of-sample forecast:

$$\Delta s_{t+1}^f = \hat{\alpha}_t + \hat{\beta}_t \Delta p_t, \quad t = R, R+1, \dots, T-1$$

where $\hat{\alpha}_t, \hat{\beta}_t$ are the parameter estimates obtained from a rolling sample of observations $\{t-R+1, t-R+2, \dots, t\}$, where R is the in-sample estimation window size. As before, we compare the oil price-based forecasts with those of the random walk by using the Diebold and Mariano's (1995) test. Figure 5(a) reports the Diebold and Mariano's (1995) test statistic for daily data computed with varying in-sample estimation windows. The size of the in-sample estimation window relative to the total sample size is reported on the x-axis. Clearly, predictability depends on the estimation window size. The Diebold and Mariano's (1995) statistic is negative for large in-sample window sizes, for which model (3) forecasts better than both the random walk with and without drift; however, the opposite happens for small in-sample window sizes. Since the Diebold and Mariano (1995) statistic is never less than -1.96, we conclude that the oil price model never forecasts significantly better than the random walk benchmark *on average* over the out-of-sample forecast period.¹²

INSERT FIGURE 5(a) HERE

Figure 5(b) reports forecast comparisons for the same model, eq. (3) at the monthly and quarterly frequencies. The model estimated at monthly and quarterly frequencies forecasts worse than the one estimated in daily data. Again, the model with monthly data does show some predictive ability for the largest window sizes, although it is not statistically significant, whereas the quarterly data model never beats the random walk.

INSERT FIGURE 5(b) HERE

However, Figure 5(c) demonstrates that, once we allow the relative performance of the models to be time-varying, the most interesting empirical results appear. Panel A in Figure

¹²Note that the MSFE ratio between the model and the random walk without drift is 0.99 for most window size. The MSFE of the random walk without drift is the same as in footnote 10.

5(b) reports the Fluctuation test in daily data. It is clear that there is strong significant evidence in favor of the model with lagged prices, especially around 2007, both against the random walk with and without drift. Panels B and C show, instead, that there was *never* statistically significant empirical evidence in favor of the model for monthly and quarterly data (in particular, against the toughest benchmark, the driftless random walk).

INSERT FIGURE 5(c) HERE

Note that the predictive ability again disappears if we use other economic fundamentals, such as interest rates differentials. Figure 5(d) reports the same analysis for the model with lagged interest rate differentials:

$$\Delta s_t = \alpha + \beta \Delta i_{t-1} + \varepsilon_t. \quad (4)$$

Clearly, in this case, the model's forecasts never beat the random walk's forecasts, no matter what the estimation window size is.

INSERT FIGURE 5(d) HERE

5 Other Commodity Prices and Exchange Rates

In this section, we show that our results are not confined to the case of the Canadian-U.S. dollar exchange rate and oil prices. We consider the predictive ability of exchange rates of other exporting countries vis-a-vis the U.S. dollar for a few additional commodity prices. In particular, we consider: (a) the price of copper (in U.S. dollars) and the Chilean Peso/U.S. dollar exchange rate; (b) the gold price (in U.S. dollars) and the South African Rand/U.S. dollar exchange rate; (c) the oil price and the Norwegian Krone/U.S. dollar exchange rate. The sample we consider is from 1/3/1994 to 9/16/2010 and the data are from Datastream. We will show that in the Norwegian Krone and the South African Rand case, oil prices and gold prices, respectively, statistically improve forecasts of exchange rates no matter whether the oil price is a contemporaneous regressor or a lagged regressor, when we allow for time variation in the relative forecasting performance of the models. The predictive ability is present only for the contemporaneous regression model for the other countries/commodity prices.

Figures 6(a)-(f) show the empirical results for forecasting the Norwegian Krone/U.S. dollar exchange rate using oil prices. In this case, the data show a clear forecasting improvement

over a random walk both in the model with contemporaneous regressors (eq. 1) at daily frequencies (see Figure 6(a)) as well as in monthly data (see Figure 6(b)), no matter which window size is used for estimation. The forecasting improvement is statistically significant in both cases, although the predictive ability again becomes statistically insignificant at quarterly frequencies. Figure 6(c) shows that the predictive ability disappears in the model with lagged fundamentals (eq. 3) under the assumption that the relative performance of the models is constant over the entire out-of-sample span of the data. However, when allowing the models' forecasting performance to change over time (Figure 6(f), Panel A), the model with lagged regressors does forecast significantly better than the random walk benchmark. Note that the performance of the lagged regressor model in monthly and quarterly frequencies is never significantly better than the random walk benchmark even if we allow the forecasting performance to change over time (see Figure 6(f), Panels B and C).

INSERT FIGURE 6(a-f) HERE

Figures 7(a)-(f) show that similar results hold when considering the South African Rand exchange rate and gold prices. Figure 7(a) shows that the predictive ability of contemporaneous gold prices is statistically significant, no matter whether the benchmark model is a random walk with or without drift, and no matter which in-sample window size the researcher chooses. In monthly and quarterly data, instead, Figure 7(b) demonstrates that fluctuations in gold prices may improve the predictive ability over a random walk model, especially for relatively large in-sample window sizes, and the forecasting improvement is never significant. Only when we allow for time variation, the model beats both the random walk with drift and without drift at daily data (Figure 7(e), Panel A). At monthly frequencies, the model beats the random walk with drift but not the driftless random walk, and at quarterly frequencies the model is always worse than both (Figure 7(e), Panels B,C). Interestingly, again, the model with lagged data never performs better than the random walk when we do not allow for time variation, no matter the frequency of the data (Figures 7(c,d)). However, when we allow for time variation (Figure 7(f)), it is clear that the model beats the driftless random walk (although it does not beat the random walk with drift) in daily data (Panel A); there is some evidence that the model beats the driftless random walk also at the quarterly frequency, but not at the monthly frequency (Panels B,C).

INSERT FIGURE 7(a-f) HERE

Figure 8(a) shows that the price of copper has a clear advantage for predicting the Chilean Peso/U.S. dollar exchange rate in the model with contemporaneous regressors at daily frequencies relative to the random walk model (with or without drift), and it is strongly statistically significant. Figure 8(b) demonstrates that such predictive ability becomes statistically insignificant when considering end-of-sample monthly and quarterly data. However, the forecasting performance disappears in the lagged regressor model (Figures 7c,d), even if we allow for time variation in the forecasting performance (Figure 8(f)).

INSERT FIGURE 8(a-f) HERE

Finally, Figure 9 considers predicting the Australian/U.S. dollar exchange rate using oil prices. Again, the predictive ability is extremely significant in daily data in the contemporaneous model (Figure 9(a)), significant at 5% in monthly data and not significant in quarterly data (Figure 9(b)). The predictive ability is also not significant in the model with lagged fundamentals at any frequency (Figures 9(c,d,f)).

INSERT FIGURE 9(a-f) HERE

6 Non-linear Models with Asymmetries and Thresholds

The recent debate on whether oil price shocks have asymmetric effects on the economy motivates us to consider such models in our forecasting experiment. Hamilton (2003) found significant asymmetries of oil price shocks on output, and Kilian and Vigfusson (2009) contend the finding claiming that the asymmetries are insignificant. Herrera, Lagalo and Wada (2010) suggest that non-linearities are really important in sectoral-based data. In this section, we evaluate whether it is possible to improve upon the simple oil price model by using non-linear models that account for the asymmetric effects of oil prices.

The model with asymmetries follows Kilian and Vigfusson (2009). We consider a model where the exchange rate response is asymmetric in oil price increases and decreases:

$$\Delta s_t = \alpha_+ + \beta_+ \Delta p_t + \gamma \Delta p_t^+ + u_t \quad (5)$$

where $\Delta p_t^+ = \begin{cases} \Delta p_t & \text{if } \Delta p_t > 0 \\ 0 & \text{otherwise.} \end{cases}$ ¹³ Our goal is to compare the forecasting ability of the model with asymmetries (5) with the linear model in eq. (1).

¹³See also Kilian (2008a,b) for analyses of the effects of oil price shocks on typical macroeconomic aggre-

In addition, we also consider a threshold model in which “large” changes in oil prices have additional predictive power for the nominal exchange rate:

$$\Delta s_t = \alpha_q + \beta_q \Delta p_t + \gamma_q \Delta p_t^q + u_t \quad (6)$$

where

$$\Delta p_t^q = \begin{cases} \Delta p_t & \text{if } \Delta p_t > 80\text{th quantile of } \Delta p_t \text{ and } < 20\text{th quantile of } \Delta p_t \\ 0 & \text{otherwise.} \end{cases}$$

and the quantiles of Δp_t are calculated over the full sample.¹⁴

We focus again on the representative case of the Canadian-U.S. dollar exchange rate and oil prices. We report results on both out-of-sample predictive ability as well as in-sample fit. To preview our findings, the empirical evidence shows that, although both the model with asymmetries and the model with threshold effects are not rejected in-sample, their forecasting ability is worse than that of the linear model, eq. (1). The same results hold especially when using lagged non-linear explanatory variables.

6.1 Out-of-Sample Forecasting Analysis with Contemporaneous Regressors

Figure 10(a) reports result for the asymmetric model and the threshold model for daily data. Both figures show the test statistic for testing the difference in the Mean Squared Forecast Errors (MSFE) of either models (5) or model (6) versus the MSFE of the linear model, eq. (1). The figures reports the test statistics calculated using a variety of sizes for the in-sample estimation window, whose size relative to the total sample size is reported on the x-axis. Negative values in the plot indicate that the linear model, eq. (1), is better than the competitors. Figure 10(b) reports results for monthly and quarterly data.

INSERT FIGURES 10(a,b) HERE

In general the simple oil price model outperforms the asymmetric model. Regarding the threshold model, the evidence is not clear cut. The threshold model is statistically better

gates, such as GDP, and Bernanke, Gertler and Watson (1997), Herrera (2008) and Herrera and Pesavento (2009) on the relationship between oil prices, inventories and monetary policy.

¹⁴We calculate the thresholds over the full sample to improve their estimates. While this gives an unfair advantage to the threshold models at beating the simple model, we still find that, even with the best estimate of the threshold, the model does not beat the simple linear model, eq. (1).

than the simple oil price model when the in-sample window size is large, whereas the result is the opposite when it is small. Figure 10(b) shows that for monthly and quarterly data the non-linear models are never statistically better than the simple linear model, and the linear model is significantly better than the non-linear models for some window sizes.

To evaluate whether forecast instabilities are important, we also implemented Fluctuation tests. Figures A.1 and A.2 in the Appendix report the results of the Fluctuation Test for both the asymmetric and threshold models at all frequencies. The figures show that the asymmetric and threshold models are never statistically better than the linear oil price model at any point in time.

6.2 Out-of-Sample Forecasting Analysis with Lagged Regressors

Let's consider asymmetric and threshold models with lagged fundamentals. The asymmetric model with lagged fundamentals is:

$$\Delta s_t = \alpha_+ + \beta_+ \Delta p_{t-1} + \gamma \Delta p_{t-1}^+ + u_t \quad (7)$$

whereas the threshold model with lagged fundamentals is:

$$\Delta s_t = \alpha_q + \beta_q \Delta p_{t-1} + \gamma_q \Delta p_{t-1}^q + u_t. \quad (8)$$

Figures 10(c) reports the forecast comparison test of models (7) and (8). Negative values indicate the linear model (3) performs better than the asymmetric model. The figure shows that the simple, linear model forecasts significantly better than both the asymmetric and threshold model for most in-sample window sizes. Figure 10(d) shows results for monthly and quarterly data. At these lower frequencies, the out-of-sample forecast performance of the simple oil price model and that of non-linear models are not statistically different.

INSERT FIGURES 10(c,d) HERE

6.3 In-sample Analysis

We estimate the asymmetric and threshold models over the entire sample period with daily, monthly and quarterly data. We focus on eqs. (5) and (6) in the in-sample analysis, since information on contemporaneous regressors is clearly available in-sample. Table 2 reports the results for the asymmetric model. As for the linear oil price model, the constant, α_+ , is never statistically significant. The slope coefficient on the rate of growth of the oil price,

β_+ , is statistically significant at the 5% percent level at all data frequencies. The coefficient γ_+ , designed to capture potential asymmetric effects, is instead never significantly different from zero.

INSERT TABLES 2 AND 3 HERE

Table 3 reports results for the threshold model. The coefficient γ_q , which quantifies the effects of “large” oil price changes, is never statistically different from zero at the 5% significance level. The coefficient on the growth rate of the oil price, β_q , instead, is statistically significant at the 5% significance level for daily and quarterly data.¹⁵

The fact that the coefficients on Δp_t^+ and Δp_t^q are never statistically different from zero suggests that non-linear oil price effects do not have in-sample predictive ability.

Panel A in Figures 11(a,b) report the rolling regressions results for the asymmetric model with daily data; Panel A in Figure 11(a) reports t-statistics on β_+ (line with circles) and γ (line with diamonds) in rolling regressions for the threshold model with daily data. As in section (3.4), we run regressions over rolling sub-samples of fixed size, and estimate a sequence of R^2 ’s and t-statistics for the slope coefficient. The series of t-statistics confirms that oil price shocks have important and significant effects for exchange rates at daily frequencies, while, at the same time, non-linearities are not important.¹⁶ Note that the asymmetries seem to play some role in fitting data at monthly frequencies, although Figure 10(d) showed that the in-sample predictive ability never translates into out-of-sample predictive power. Similar results hold for the threshold model (see Figures 12(a,b)).

INSERT FIGURES 11(a,b) AND 12(a,b)

This last observation, together with the full sample evidence, leads to the conclusion that it is the growth rate of the oil price the explanatory variable that contains in-sample predictive ability: asymmetries or thresholds effects do not provide any additional predictive power. We conclude that non-linearities in the effects of oil prices are not empirically important, neither in-sample nor out-of-sample, for explaining or predicting exchange rate fluctuations.

¹⁵Similarly to the linear oil price and asymmetric models, the R^2 of the threshold model for quarterly data are much larger than those for daily data.

¹⁶The R^2 ’s sequences of both models display the same pattern of the simple oil price model. The in-sample fit increases toward the end of the sample period.

7 Conclusions

Our empirical results suggest that oil prices can predict the Canadian/U.S. dollar nominal exchange rate at daily frequency, in the sense of having a stable out-of-sample relationship. However, the predictive ability is not evident at quarterly and monthly frequencies. When using contemporaneous realized daily oil price shocks to predict exchange rates, the predictive power of oil prices is robust to the choice of the in-sample window size and it does not depend on the sample period under consideration. When using the lagged oil price shocks to predict exchange rates, the predictive ability is more ephemeral and only shows up in daily data after allowing the relative forecasting performance of the oil price model and the random walk to be time-varying. Both the out-of-sample and in-sample analyses suggest that the frequency of the data is important to detect the predictive ability of oil prices, as the out-of-sample predictive ability breaks down when considering monthly and quarterly data. Following Kilian and Vigfusson (2009), we also consider two models aimed at modeling potentially important non-linearities in the oil price-exchange rate relationship. We find that non-linearities do not significantly improve upon the simple linear oil price model.

Our results suggest that the most likely explanations for why the existing literature has been unable to find evidence of predictive power in oil prices is that they focused on low frequencies where the short-lived effects of oil price shocks wash away and that the superior predictive ability of the oil price shocks is very transitory.

At the same time, our results also raise interesting questions. For example, does the Canadian/U.S. dollar exchange rate respond to demand or supply shocks to oil prices? It would be interesting to investigate this question by following the approach in Kilian (2009). It would also be interesting to consider predictive ability at various horizons by adjusting the current exchange rate for recent changes in the price of oil over a longer period (e.g. a week). We leave these issues for future research.

References

- Alquist, R., L. Kilian and R.J. Vigfusson (2011), “Forecasting the Price of Oil”, in: G. Elliott and A. Timmermann, *Handbook of Economic Forecasting* Vol. 2, Elsevier-North Holland.
- Amano, R. and S. Van Norden (1998a), “Exchange Rates and Oil Prices”, *Review of International Economics*.
- Amano, R. and S. Van Norden (1998b), “Oil Prices and the Rise and Fall of the U.S. Real Exchange Rate”, *Journal of International Money and Finance*.
- Andersen, T., Bollerslev, T., Diebold, F.X. and Vega, C. (2003), “Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange,” *American Economic Review* 93, 38-62.
- Backus, D.K. and M.J. Crucini (2000), “Oil Prices and the Terms of Trade”, *Journal of International Economics* 50(1), 185-213.
- Berkowitz, J. and L. Giorgianni (2001), “Long-horizon exchange rate predictability?”, *Review of Economics and Statistics* 83, 81–91.
- Bernanke, B.S., M. Gertler, and M. Watson (1997), “Systematic Monetary Policy and the Effects of Oil Price Shocks,” *Brookings Papers on Economic Activity* 1, 91-142.
- Chen, Y. and K.S. Rogoff (2003), “Commodity Currencies”, *Journal of International Economics* 60, 133-169.
- Chen, Y., K.S. Rogoff and B. Rossi (2010), Can Exchange Rates Forecast Commodity Prices?, *Quarterly Journal of Economics*, forthcoming.
- Chinn, M.D., and R.A. Meese, (1995), Banking on currency forecasts: how predictable is change in money? *Journal of International Economics* 38, 161–178.
- Cheung, Y.W., M.D. Chinn, and A.G. Pascual (2005), “Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?”, *Journal of International Money and Finance* 24, 1150-1175.
- Clark, T., and K. D. West (2006), “Using Out-of-sample Mean Squared Prediction Errors to Test the Martingale Difference Hypothesis”, *Journal of Econometrics* 135, 155-186.
- Diebold, F.X., and R. Mariano (1995), “Comparing Predictive Accuracy”, *Journal of Business and Economic Statistics* 13, 253-263.
- Engel, C., N. Mark and K.D. West (2007), “Exchange Rate Models Are Not as Bad as You Think,” in: *NBER Macroeconomics Annual*, D. Acemoglu, K.S. Rogoff and M. Woodford, eds., Cambridge, MA: MIT Press.

Faust, J., J.H. Rogers and J.H. Wright (2003), “Exchange rate forecasting: the errors we’ve really made”, *Journal of International Economics* 60(1), 35-59.

Faust, J., J.H. Rogers, S-Y.B. Wang and J.H. Wright (2007), “The High-Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements”, *Journal of Monetary Economics* 54, 1051-1068.

Giacomini, R. and B. Rossi (2010), Forecast Comparisons in Unstable Environments, *Journal of Applied Econometrics* 25(4), 595-620.

Hamilton, J.D. (2003), “What is an Oil Shock?” *Journal of Econometrics* 113, 363-398.

Herrera, A.M. (2008), “Oil Price Shocks, Inventories, and Macroeconomic Dynamics,” *mimeo*, Department of Economics, Michigan State University.

Herrera, A.M., L.G. Lagalo and T. Wada (2010), “Oil Price Shocks and Industrial Production: Is the Relationship Linear?”, *mimeo*, Wayne State University.

Herrera, A.M. and E. Pesavento (2007), “Oil Price Shocks, Systematic Monetary Policy and the ‘Great Moderation’”, *Macroeconomic Dynamics*.

Hooker, M.A. (1996), “What Happened to the Oil Price-Macroeconomy Relationship?” *Journal of Monetary Economics* 38, 195-213.

Kilian, L. (1999), “Exchange rates and monetary fundamentals: what do we learn from long-horizon regressions?”, *Journal of Applied Econometrics* 14, 491–510.

Kilian, L. (2008a), “The Economic Effects of Energy Price Shocks”, *Journal of Economic Literature* 46, 871-909.

Kilian, L. (2008b), “Exogenous Oil Supply Shocks: How Big Are They and How Much Do they Matter for the U.S. Economy?” *Review of Economics and Statistics* 90, 216-240.

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), 1053-1069.

Kilian, L. and C. Vega (2008), “Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices,” *mimeo*, University of Michigan.

Kilian, L. and M.P. Taylor (2003), “Why is it so difficult to beat the random walk forecast of exchange rates?”, *Journal of International Economics*, 60(1), 85-107.

Kilian, L. and R. Vigfusson (2009), “Pitfalls in Estimating the Asymmetric Effects of Oil Price Shocks”, *mimeo*, U. of Michigan.

Kilian, L. and R. Vigfusson (2009), “Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?”, *mimeo*, U. of Michigan.

Mark, N.C. (1995), “Exchange rates and fundamentals: evidence on long-horizon prediction”, *American Economic Review* 85, 201–218.

Meese, R. and K.S. Rogoff (1983a), "Exchange Rate Models of the Seventies. Do They Fit Out of Sample?", *The Journal of International Economics* 14, 3-24.

Meese, R. and K.S. Rogoff (1983b), "The Out of Sample Failure of Empirical Exchange Rate Models," in: *Exchange Rates and International Macroeconomics*, J. Frankel, ed. (Chicago: University of Chicago Press for NBER).

Meese, R. and K.S. Rogoff (1988), "Was it Real? The Exchange Rate-Interest Differential Relation Over the Modern Floating Rate Period", *The Journal of Finance* 43, 923-948.

Molodtsova, T. and D. H. Papell (2009), "Out-of-sample exchange rate predictability with Taylor rule fundamentals", *Journal of International Economics* 77(2) 167-180.

Molodtsova, T., A. Nikolsko-Rzhevskyy and D. H. Papell (2008), "Taylor rules with real-time data: A tale of two countries and one exchange rate", *Journal of Monetary Economics* 55, S63-S79.

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, 740-744.

Obstfeld, M. (2002), "Exchange Rates and Adjustment: Perspectives from the New Open-Economy Macroeconomics", *Monetary and Economic Studies* Special Edition, 23-46.

Obstfeld, M., and K.S. Rogoff (1996), *Foundations of International Macroeconomics*, Cambridge, MA: MIT Press.

Rogoff, K.S. (2007), "Comment to: Exchange Rate Models Are Not as Bad as You Think", in: *NBER Macroeconomics Annual*, D. Acemoglu, K.S. Rogoff and M. Woodford, eds., Cambridge, MA: MIT Press.

Rogoff, K.S. and V. Stavrakeva (2008), "The Continuing Puzzle of Short Horizon Exchange Rate Forecasting," *NBER Working paper* No. 14071.

Rossi, B. (2005), "Testing Long-Horizon Predictive Ability with High Persistence, and the Meese-Rogoff Puzzle", *International Economic Review* 46(1), 61-92.

Rossi, B. (2006), "Are Exchange Rates Really Random Walks", *Macroeconomic Dynamics* 10(1), 20-38.

Rossi, B. (2007), "Comment to: Exchange Rate Models Are Not as Bad as You Think", in: *NBER Macroeconomics Annual*, D. Acemoglu, K.S. Rogoff and M. Woodford, eds., Cambridge, MA: MIT Press.

Wang, J. and J. Wu (2008), "The Taylor Rule and Forecast Intervals for Exchange Rates", *Federal Reserve Bank of Dallas Working Paper* No. 22.

West, K.D. (1996), "Asymptotic Inference about Predictive Ability", *Econometrica* 64, 1067-1084.

Figures and Tables

Table 1. Estimates of the Basic Linear Model with Oil Prices

	Daily	Monthly	Quarterly
R^2	0.03	0.09	0.21
α	-0.000 (-0.67)	-0.000 (-0.57)	-0.002 (-0.552)
β	-0.03 (-14.36)	-0.059 (-5.43)	-0.085 (-5.18)

Notes to the Table. The model is eq. (1); t-statistics reported in parentheses.

Table 1b. Estimates of the Model with Interest Rates

	Daily	Monthly	Quarterly
R^2	0.00001	0.0014	0.0008
α	-0.00001 (-0.25)	-0.0007 (-0.39)	-0.0007 (-0.12)
β	0.00002 (0.12)	0.0004 (0.42)	-0.0004 (-0.12)

Notes to the Table. The model is eq. (2); t-statistics reported in parentheses.

Table 2. Estimates of the Asymmetric Model

	Daily	Monthly	Quarterly
R^2	0.029	0.08	0.20
α	- 0.000 (-0.53)	0.000 (0.28)	-0.002 (-0.35)
β	-0.030 (-9.15)	-0.046 (-2.42)	-0.08 (-3.13)
γ	0.000 (0.04)	-0.029 (-0.88)	-0.001 (-0.02)

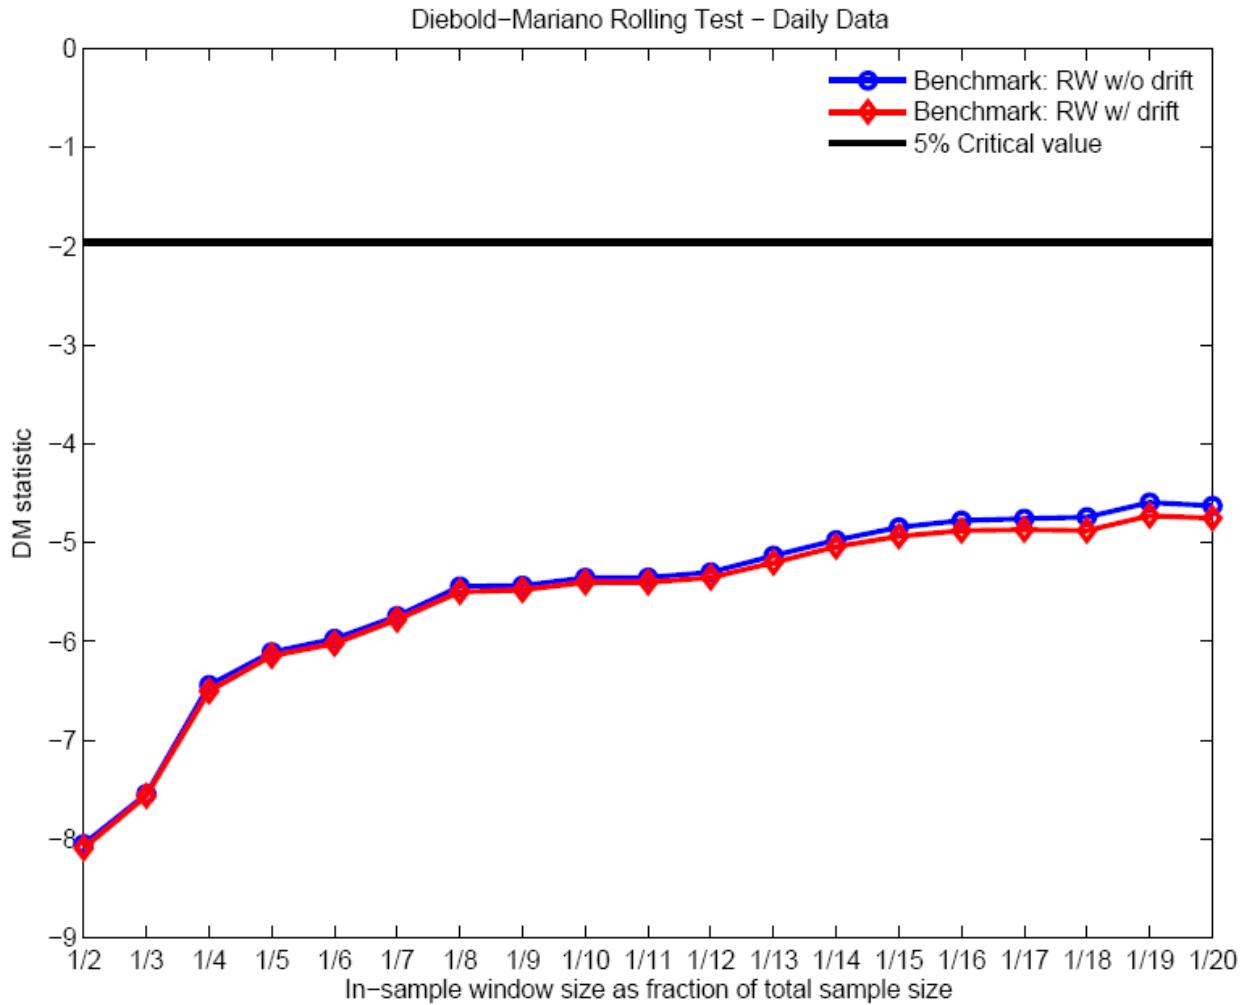
Notes to the Table. The model is eq. (5); t-statistics reported in parentheses.

Table 3. Estimates of the Threshold Model

	Daily	Monthly	Quarterly
R^2	0.03	0.08	0.20
α	-0.000 (-0.62)	-0.000 (-0.62)	-0.002 (-0.54)
β	-0.04 (-4.15)	-0.05 (-1.41)	-0.08 (-1.49)
γ	0.008 (0.87)	-0.014 (-0.39)	-0.002 (-0.027)

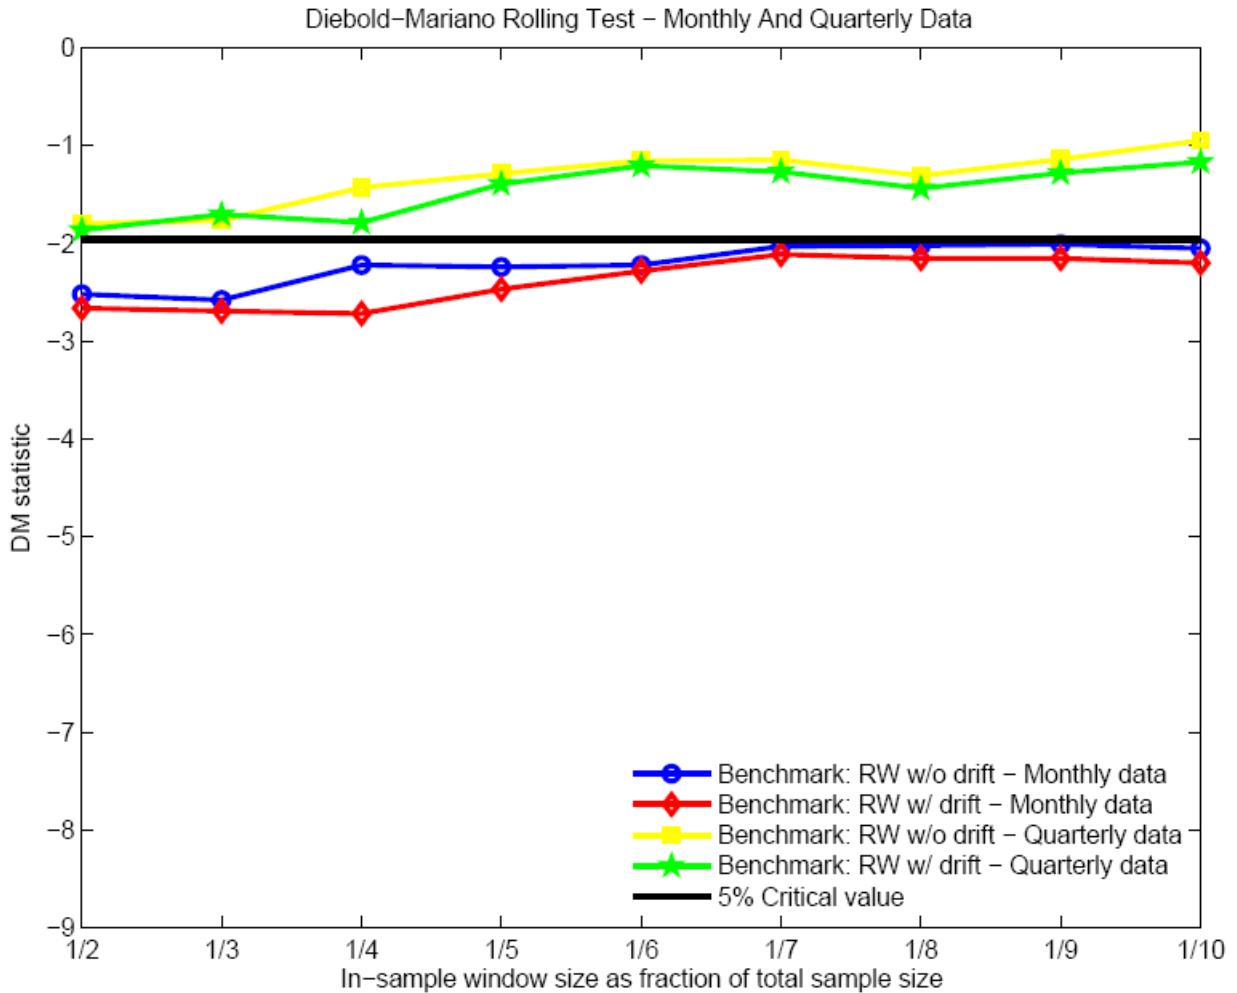
Notes to the Table. The model is eq. (6); t-statistics reported in parentheses.

Figure 1(a). Oil Price Model. Forecasting Ability in Daily Data



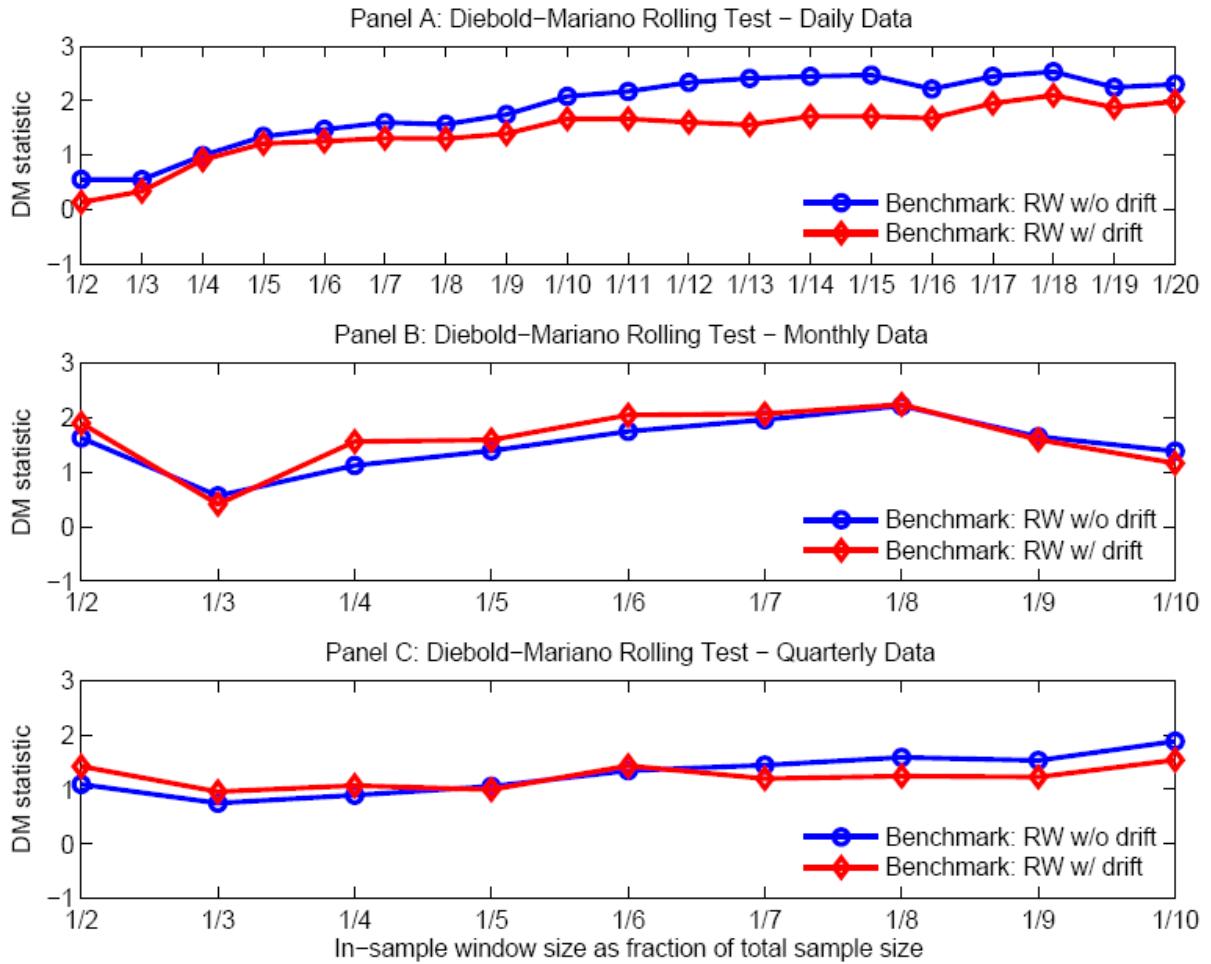
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for daily data and several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark.

Figure 1(b). Oil Price Model. Forecasting Ability in Monthly and Quarterly Data



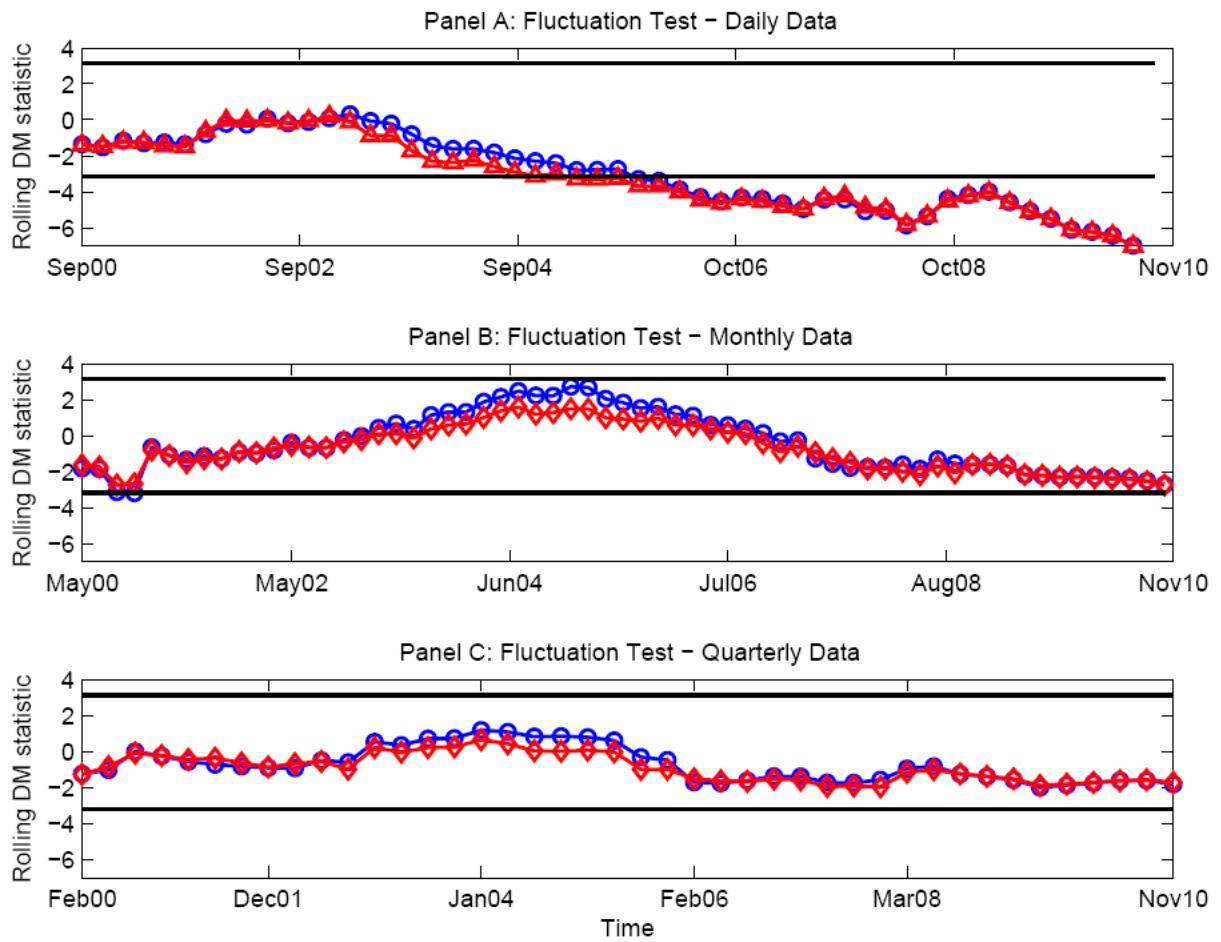
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles for monthly data and squares for quarterly data) as well as relative to the Random Walk with drift benchmark (line with diamonds for monthly data and stars for quarterly data) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark.

Figure 2. The Interest Rate Model.



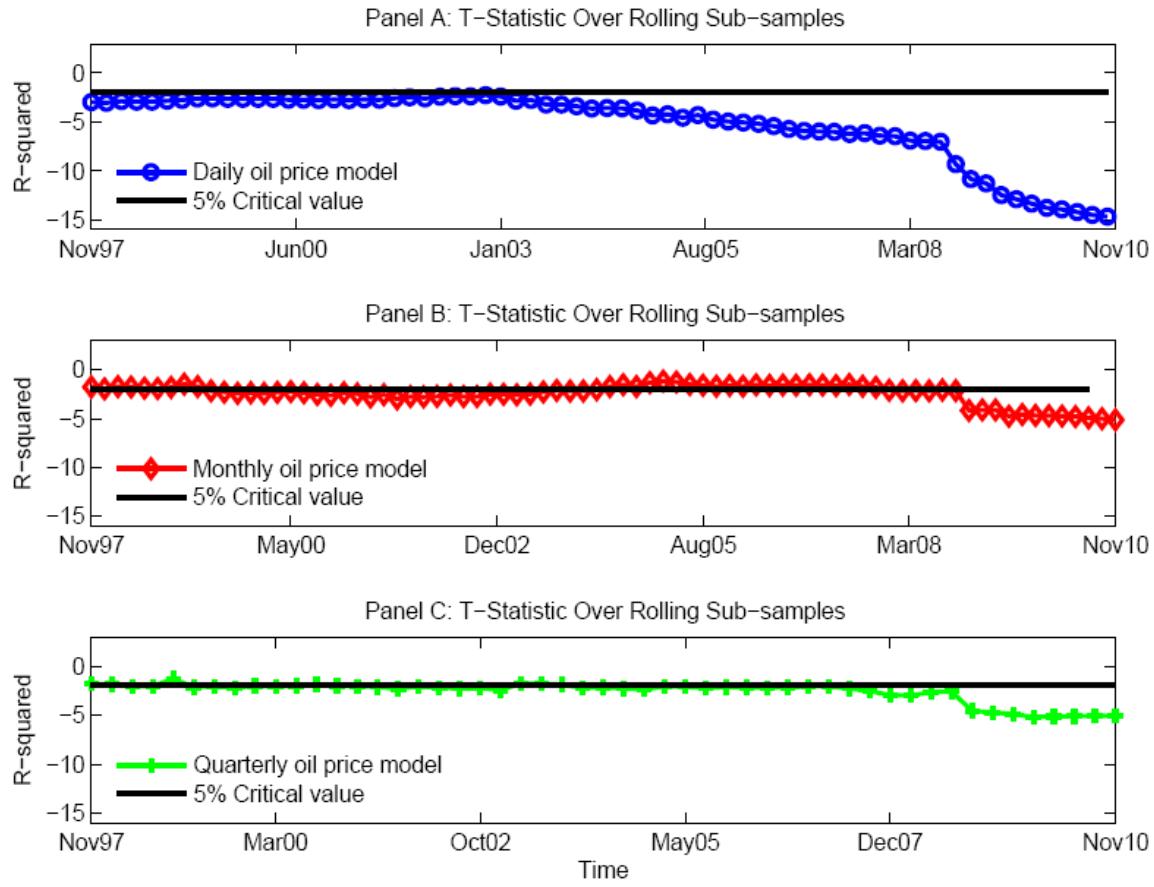
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (2) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (2) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (2) forecasts significantly better than its benchmark.

Figure 3. Fluctuation Test For the Oil Price Model



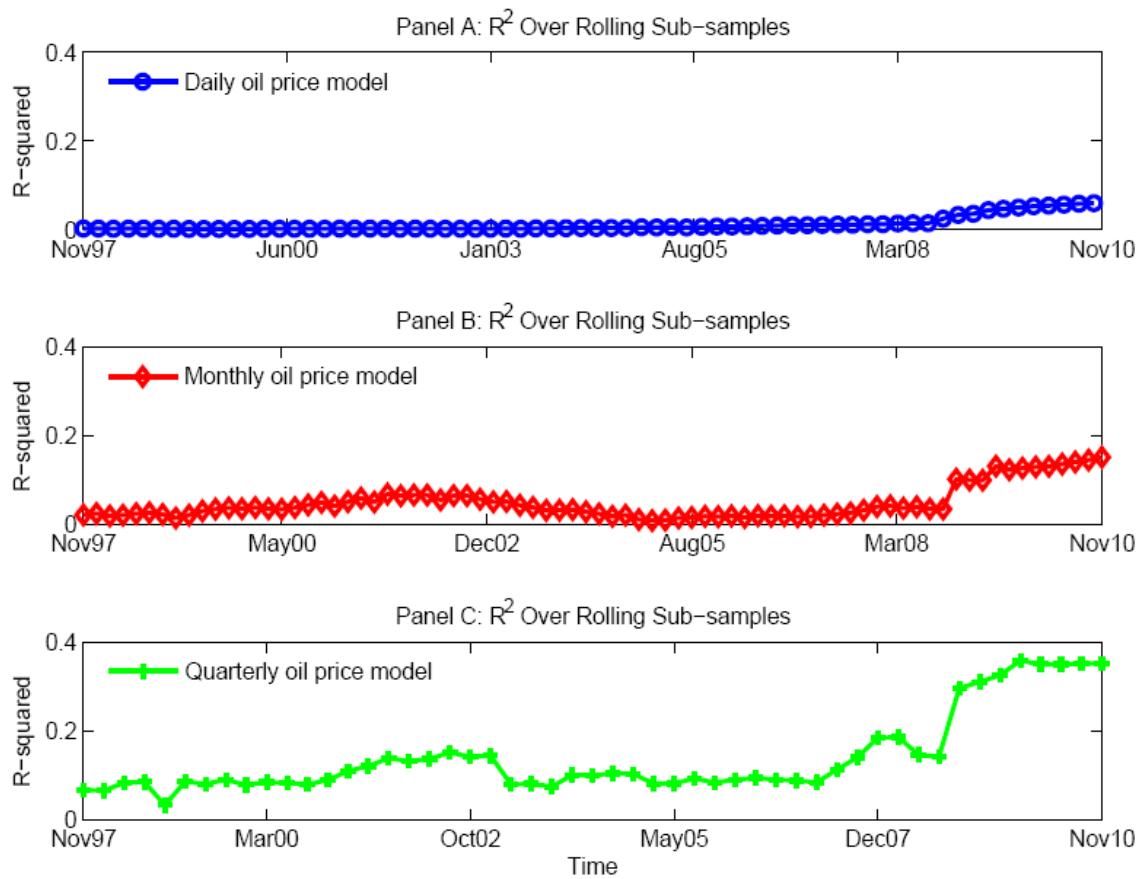
Notes to the Figure. The figure reports the Giacomini and Rossi's (2009) Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with triangles). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Fluctuation test statistic: if the estimated test statistics is below this line, Model (1) forecasts significantly better than its benchmark.

Figure 4(a). In-sample Fit of Oil Price Model – T-statistics Over Time



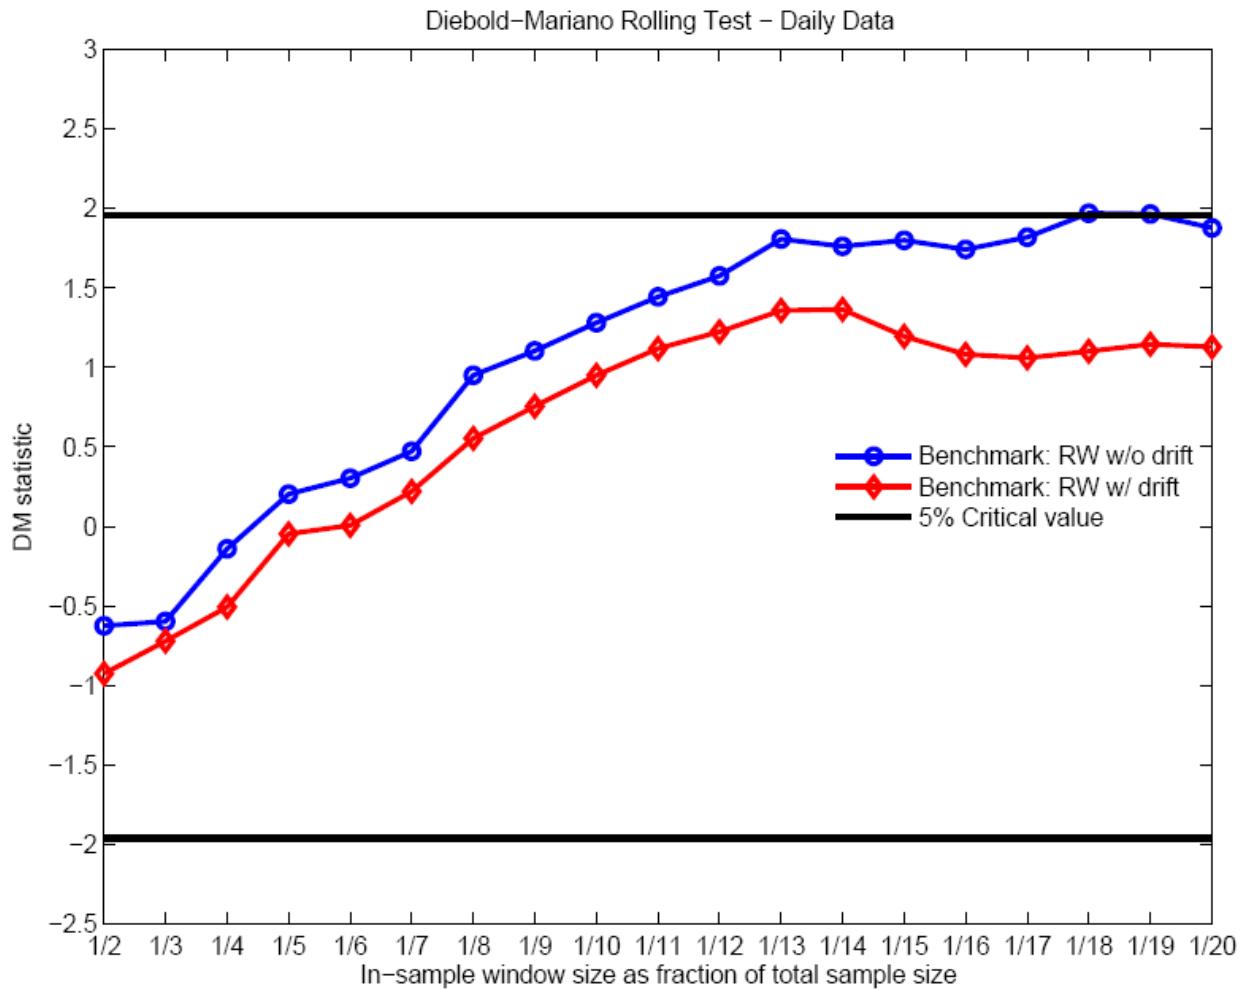
Notes to the Figure. The figure reports in-sample t-statistics for comparing forecasts of Model (1) calculated over rolling samples (dates reported on the x-axis). The continuous line indicates the critical value of the t-statistic: if the estimated test statistics is below this line, the coefficient on the oil price in Model (1) is statistically significantly negative. Top panel is for daily data, middle panel for monthly and bottom panel for quarterly data.

Figure 4(b). In-sample Fit of Oil Price Model – R^2 statistics Over Time



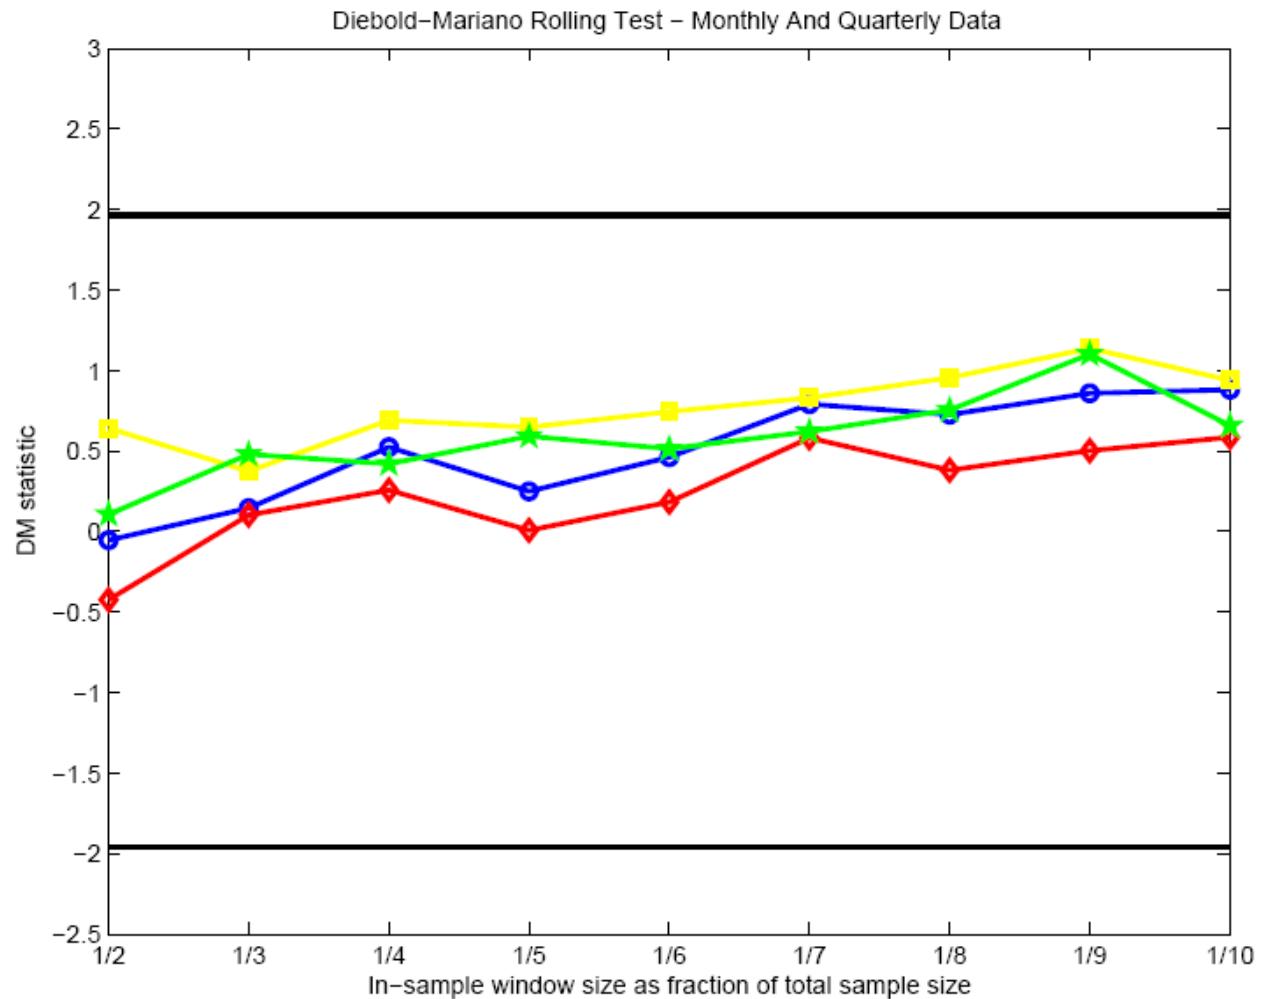
Notes to the Figure. The figure reports in-sample R^2 statistics for comparing forecasts of Model (1) calculated over rolling samples (dates reported on the x-axis).

Figure 5(a). Oil Price Model. Forecasting Ability in Daily Data



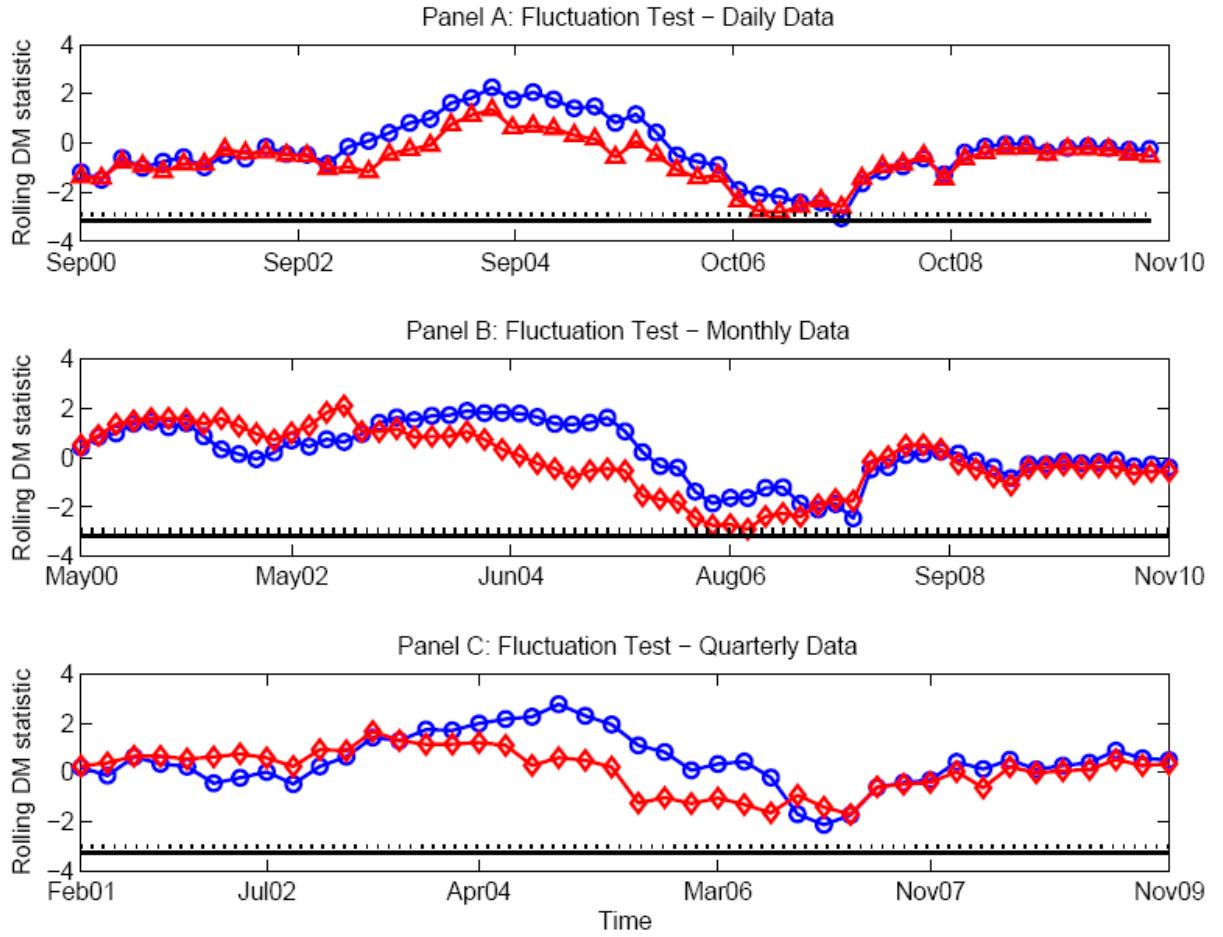
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (3) relative to a Random Walk without drift benchmark (line with circles for monthly data and squares for quarterly data) as well as relative to the Random Walk with drift benchmark (line with diamonds for monthly data and stars for quarterly data) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (3) forecasts significantly better than its benchmark.

Figure 5(b). Oil Price Model. Forecasting Ability in Monthly and Quarterly Data



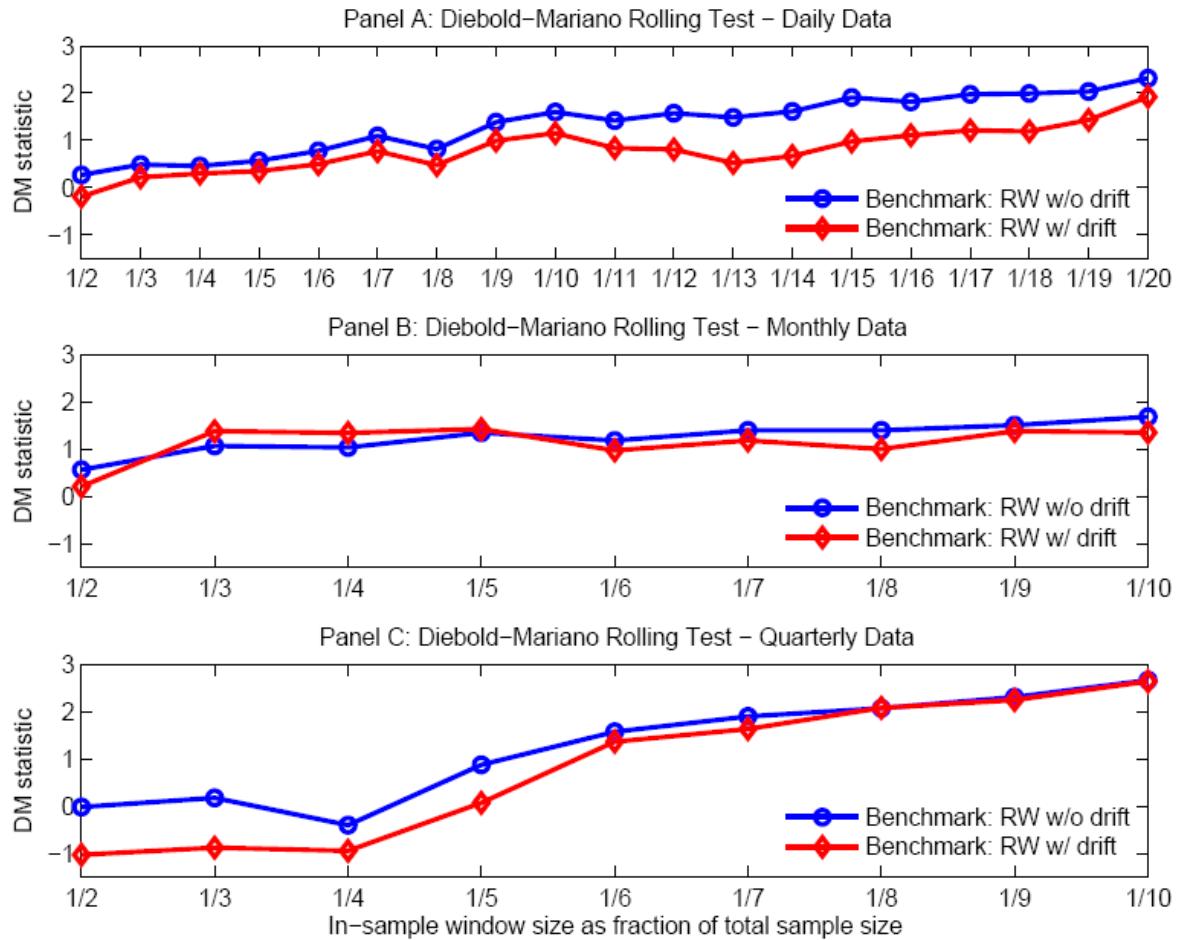
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (3) relative to a Random Walk without drift benchmark (line with circles for monthly data and squares for quarterly data) as well as relative to the Random Walk with drift benchmark (line with diamonds for monthly data and stars for quarterly data) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (3) forecasts significantly better than its benchmark.

Figure 5(c). Fluctuation Test For the Oil Price Model



Notes to the Figure. The figure reports the Giacomini and Rossi's (2009) Fluctuation test statistic for comparing forecasts of Model (3) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with triangles). Negative values indicate that Model (3) forecasts better. The continuous line indicates the critical value of the Fluctuation test statistic: if the estimated test statistics is below this line, Model (3) forecasts significantly better than its benchmark. The continuous and dashed lines denote, respectively, the two-sided 5% and 10%-level critical values.

Figure 5(d). The Interest Rate Model. Forecasting Ability in Daily, Monthly and Quarterly Data



Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (4) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for daily, monthly and quarterly data and several in-sample window sizes (x-axis). Negative values indicate that Model (4) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (4) forecasts significantly better than its benchmark.

Figure 6(a). Norw. Krone and Oil.
Daily Data, Contemp. Model

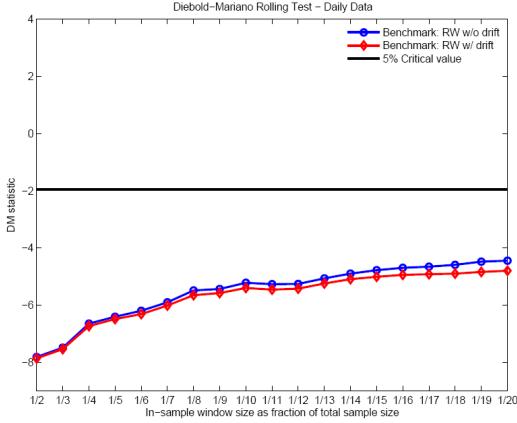


Figure 6(c). Norw. Krone and Oil.
Daily Data, Lagged Model

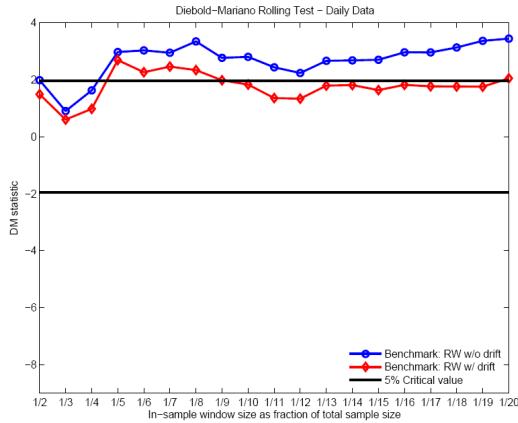


Figure 6(b). Norw. Krone and Oil.
Monthly and Quarterly Contemp. Model

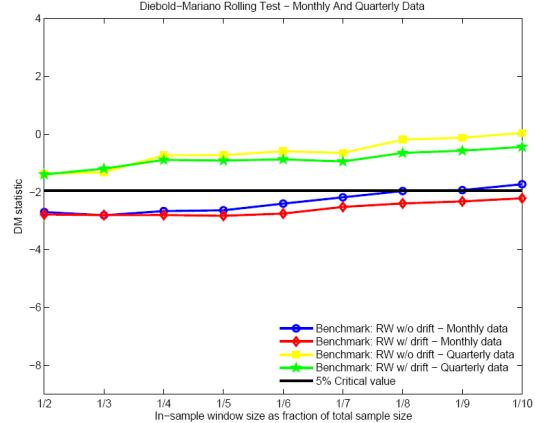
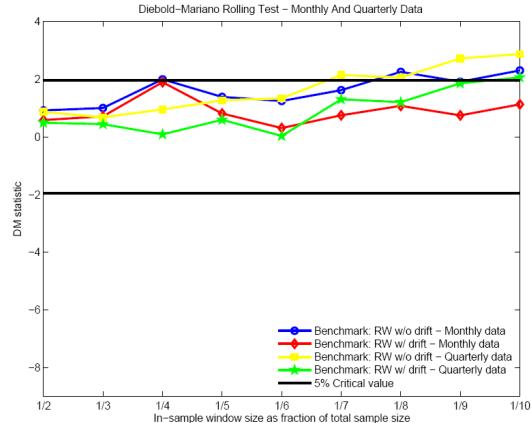


Figure 6(d). Norw. Krone and Oil.
Monthly and Quarterly Lagged Model



Notes to the Figure. Panels (a,b) report Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for daily (panel (a)), monthly and quarterly data (panel (b)), and several in-sample window sizes (x-axis). Similarly, Panels (c,d) report the same analysis for Model (3). Negative values indicate that Model (1) or (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) or (3) forecast significantly better than its benchmark.

Figure 6(e). Norw. Krone and Oil. Fluctuation Test, Contemp. Model

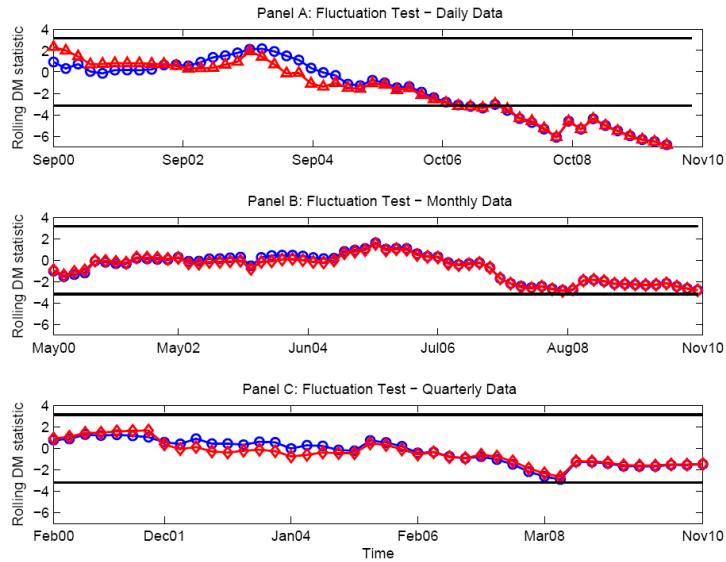
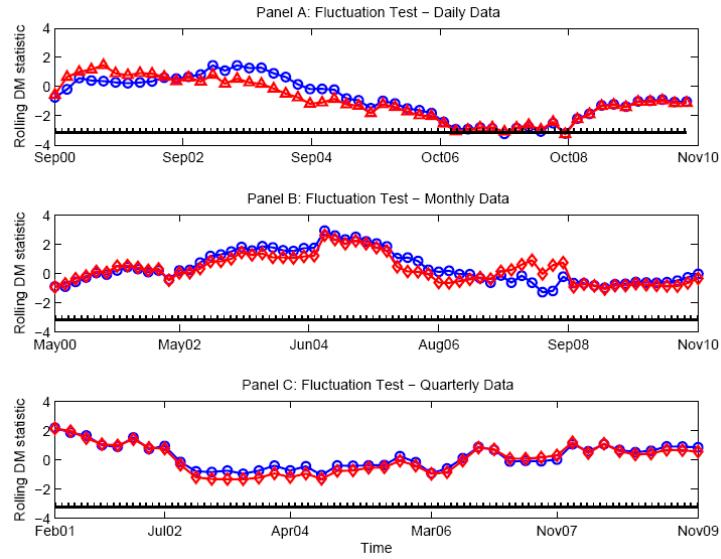
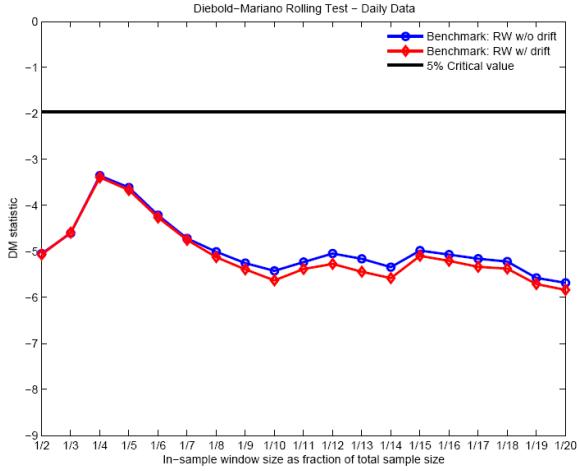


Figure 6(f). Norw. Krone and Oil. Fluctuation Test, Lagged Model

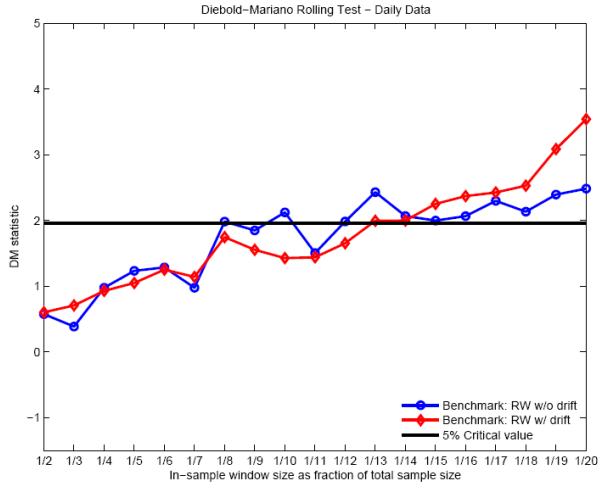


Notes to the Figure. Panel (a) reports the Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated at daily, monthly and quarterly frequencies, and several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark. Panel (b) reports the same analysis for Model (3).

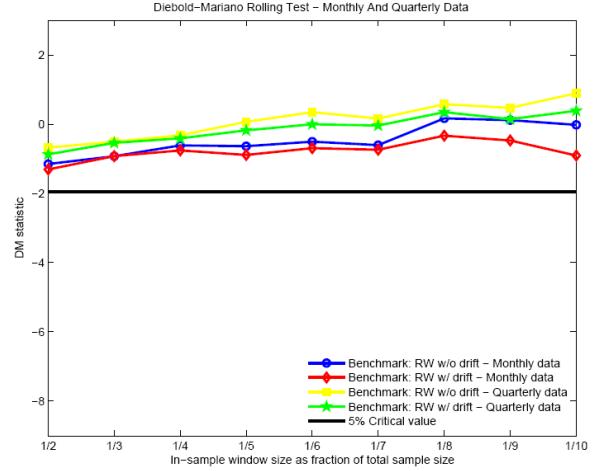
**Figure 7(a). S.A. Rand and Gold.
Daily Data, Contemp. Model**



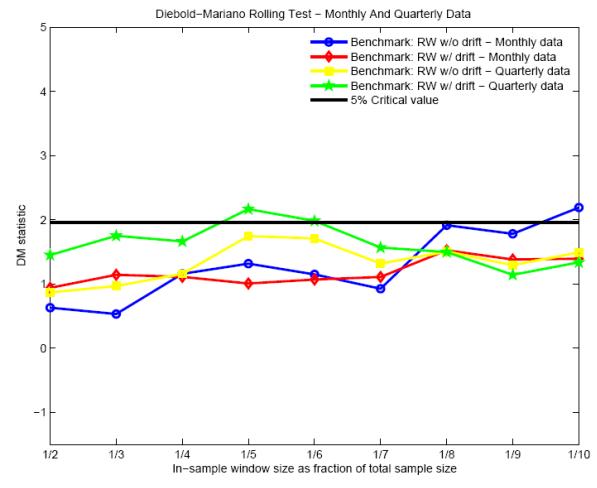
**Figure 7(c). S.A. Rand and Gold.
Daily Data, Lagged Model**



**Figure 7(b). S.A. Rand and Gold.
Monthly and Quarterly Contemp. Model**



**Figure 7(d). S.A. Rand and Gold.
Monthly and Quarterly Lagged Model**



Notes to the Figure. Panels (a,b) report Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for daily (panel (a)), monthly and quarterly data (panel (b)), and several in-sample window sizes (x-axis). Similarly, Panels (c,d) report the same analysis for Model (3). Negative values indicate that Model (1) or (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) or (3) forecast significantly better than its benchmark.

Figure 7(e). S.A. Rand and Gold. Fluctuation Test, Contemp. Model

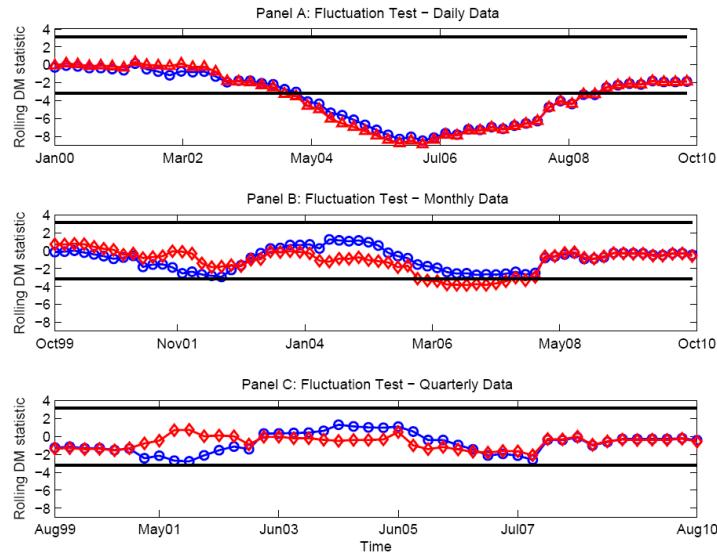
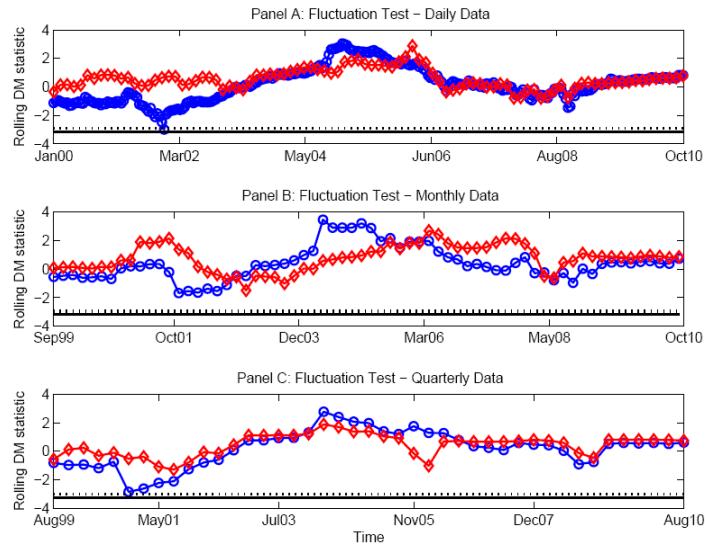
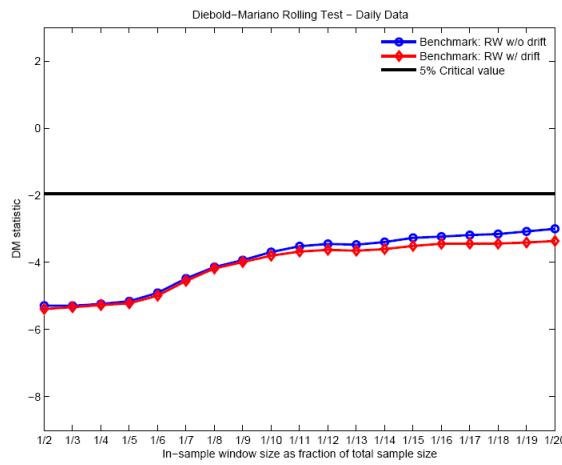


Figure 7(f). S.A. Rand and Gold. Fluctuation Test, Lagged Model

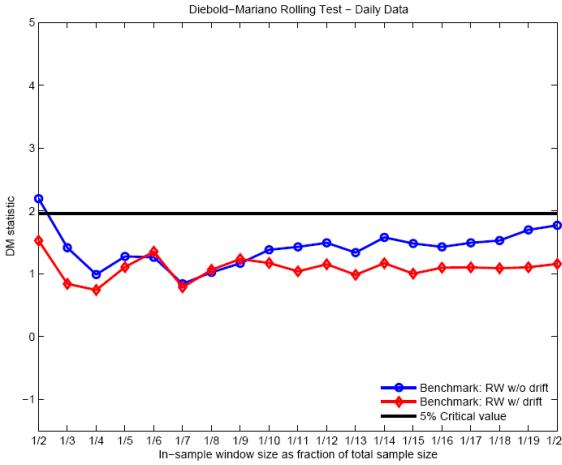


Notes to the Figure. Panel (a) reports the Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated at daily, monthly and quarterly frequencies, and several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark. Panel (b) reports the same analysis for Model (3).

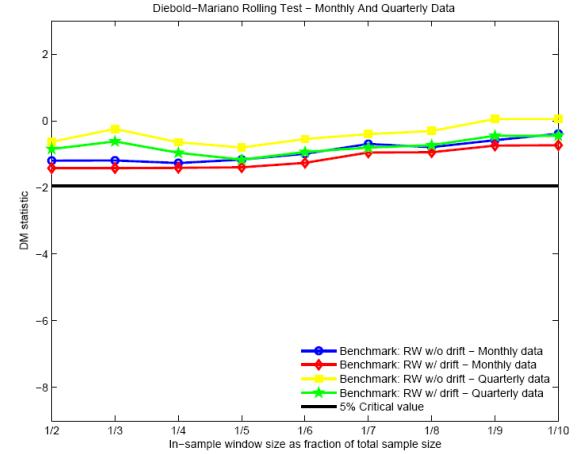
**Figure 8(a). Chilean Peso and Copper.
Daily Data, Contemp. Model**



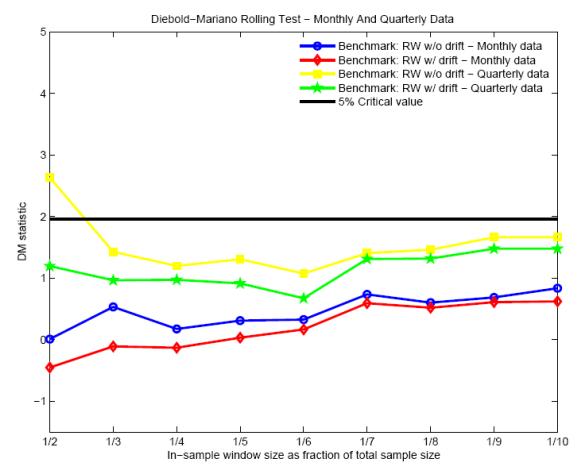
**Figure 8(c). Chilean Peso and Copper.
Daily Data, Lagged Model**



**Figure 8(b). Chilean Peso and Copper.
Monthly and Quarterly Contemp. Model**



**Figure 8(d). Chilean Peso and Copper.
Monthly and Quarterly Lagged Model**



Notes to the Figure. Panels (a,b) report Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for daily (panel (a)), monthly and quarterly data (panel (b)), and several in-sample window sizes (x-axis). Similarly, Panels (c,d) report the same analysis for Model (3). Negative values indicate that Model (1) or (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) or (3) forecast significantly better than its benchmark.

Figure 8(e). Chilean Peso and Copper. Fluctuation Test, Contemp. Model

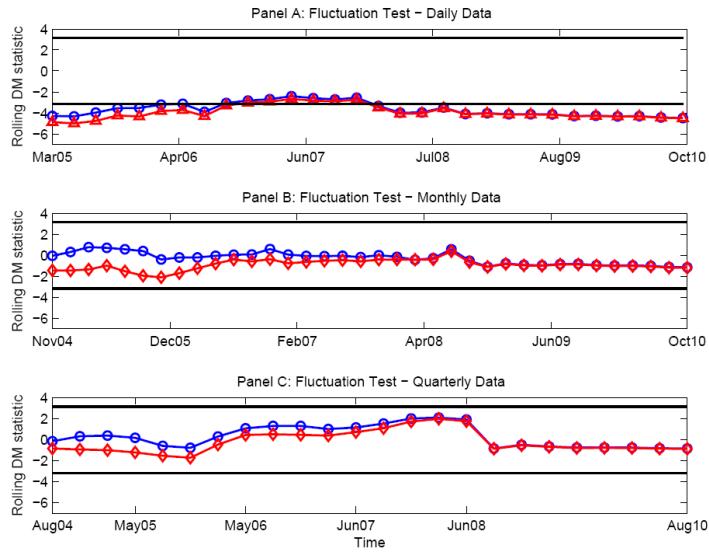
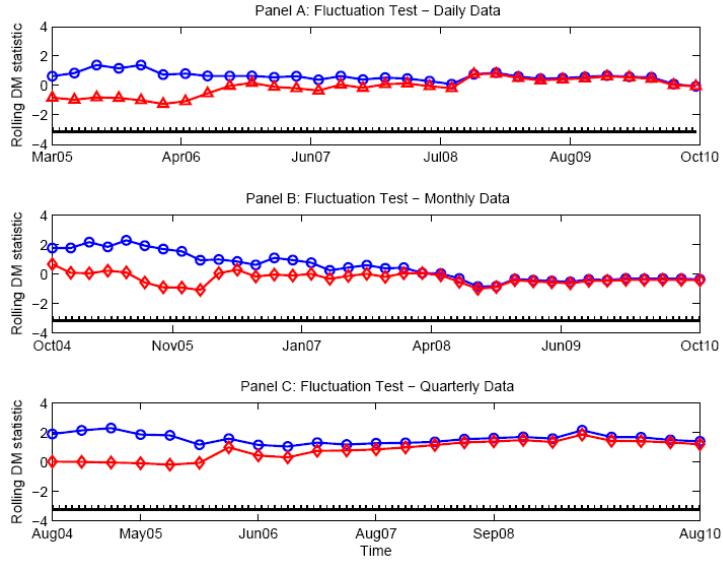
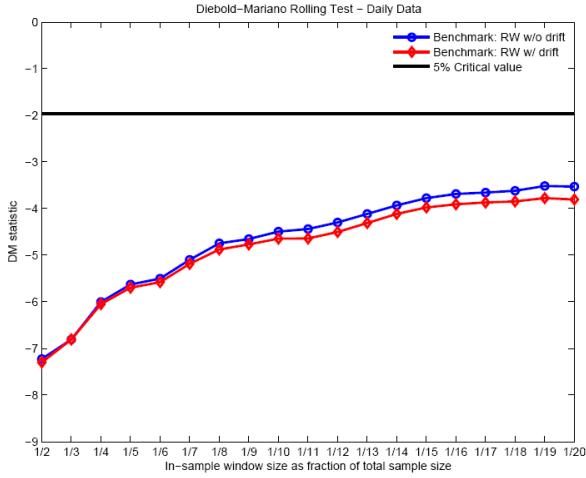


Figure 8(f). Chilean Peso and Copper. Fluctuation Test, Lagged Model

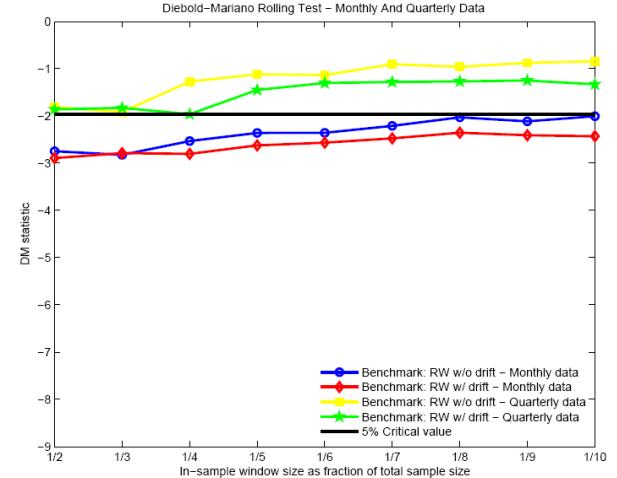


Notes to the Figure. Panel (a) reports the Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated at daily, monthly and quarterly frequencies, and several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark. Panel (b) reports the same analysis for Model (3).

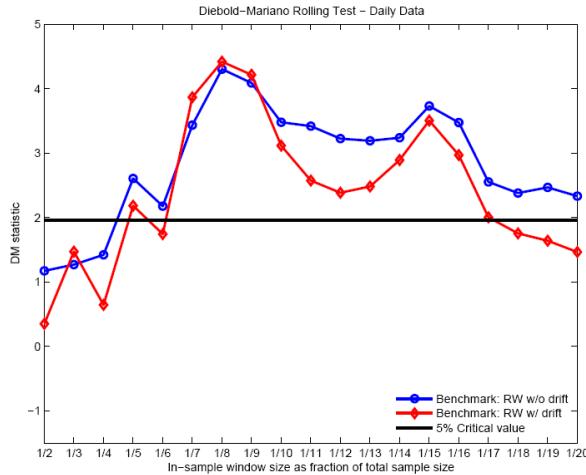
**Figure 9(a). Austr. \\$ and Oil.
Daily Data, Contemp. Model**



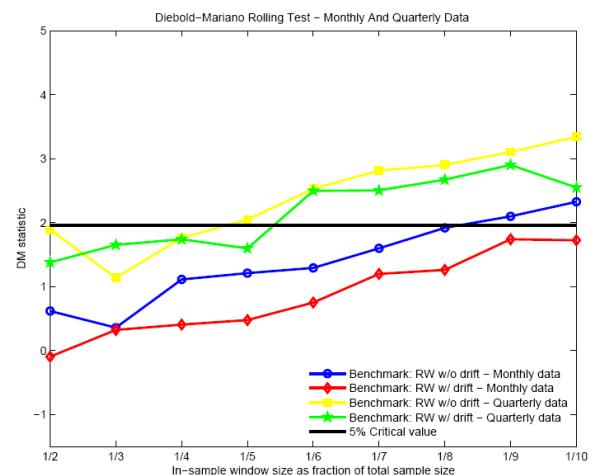
**Figure 9(b). Austr. \\$ and Oil.
Monthly and Quarterly Contemp. Model**



**Figure 9(c). Austr. \\$ and Oil.
Daily Data, Lagged Model**



**Figure 9(d). Austr. \\$ and Oil.
Monthly and Quarterly Lagged Model**



Notes to the Figure. Panels (a,b) report Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated for daily (panel (a)), monthly and quarterly data (panel (b)), and several in-sample window sizes (x-axis). Similarly, Panels (c,d) report the same analysis for Model (3). Negative values indicate that Model (1) or (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) or (3) forecast significantly better than its benchmark.

Figure 9(e). Australian \$ and Oil. Fluctuation Test, Contemp. Model

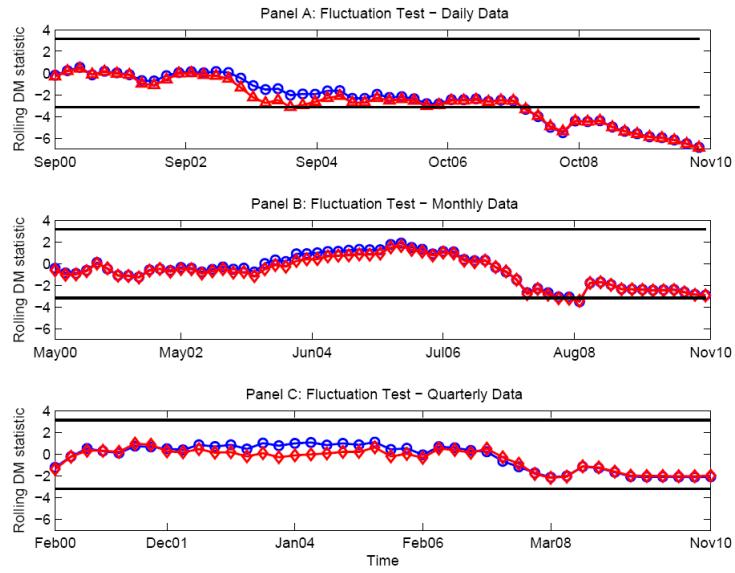
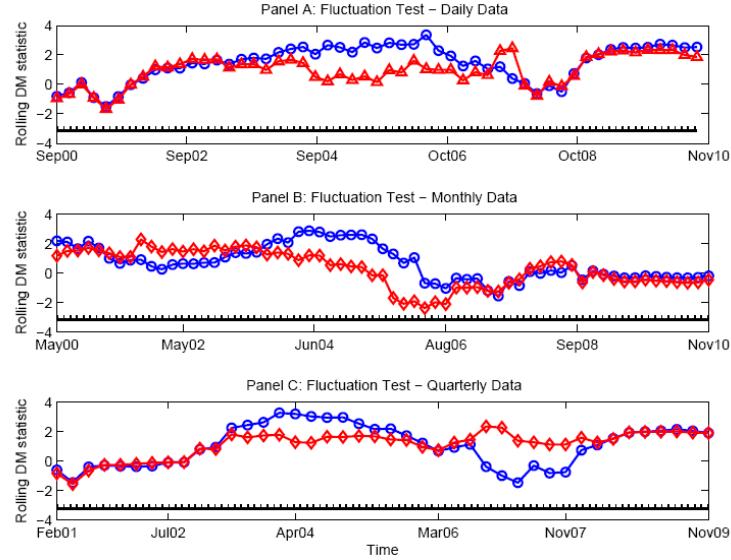
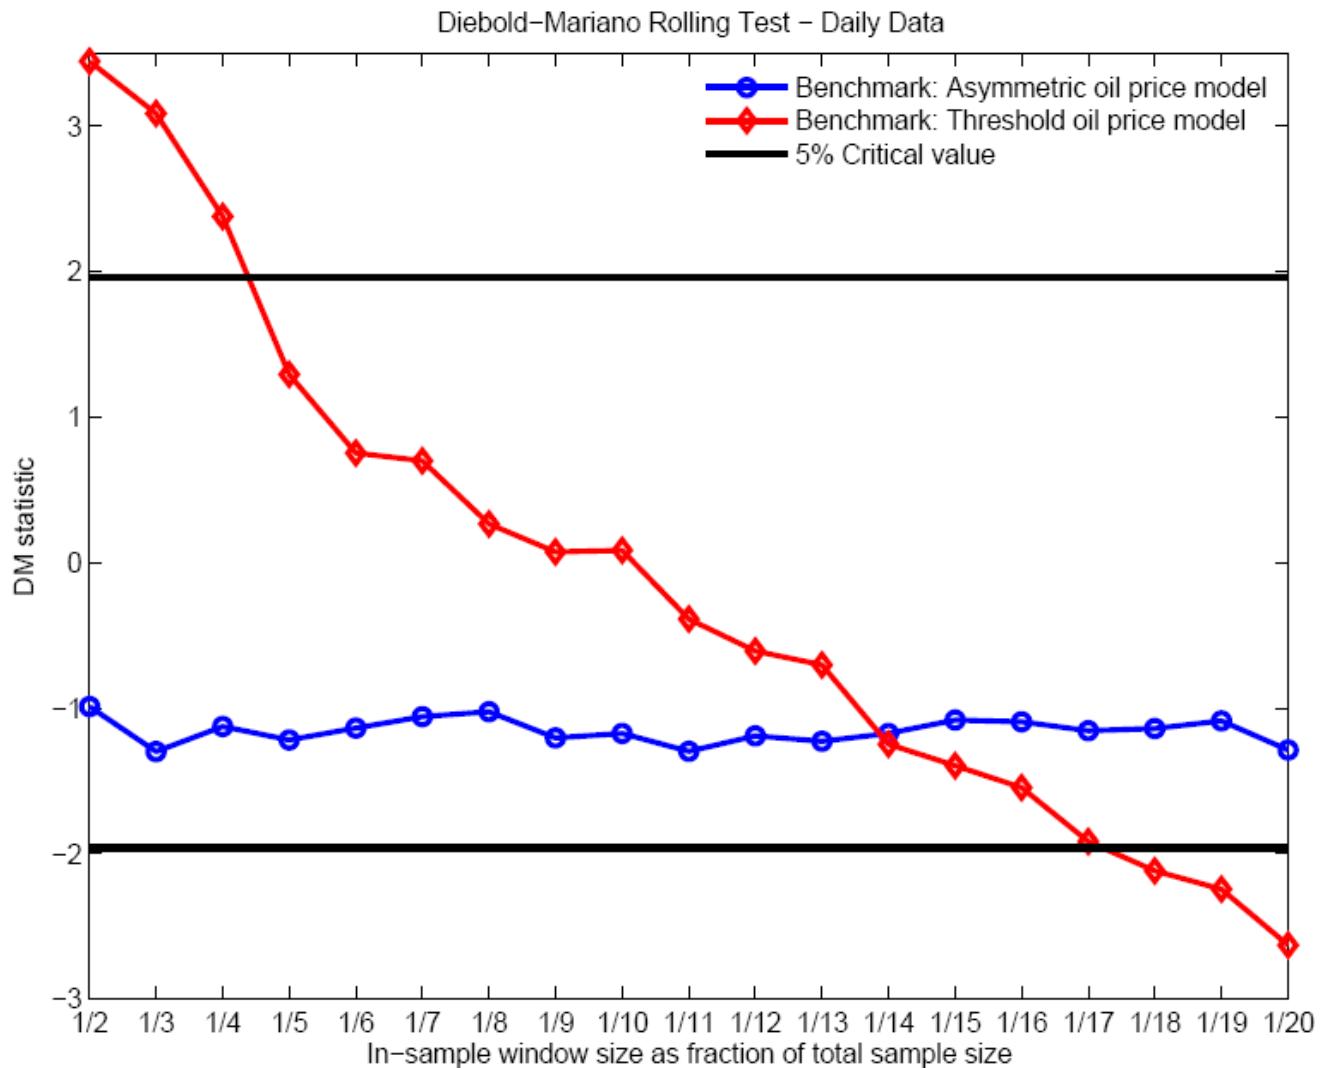


Figure 9(f). Australian \$ and Oil. Fluctuation Test, Lagged Model



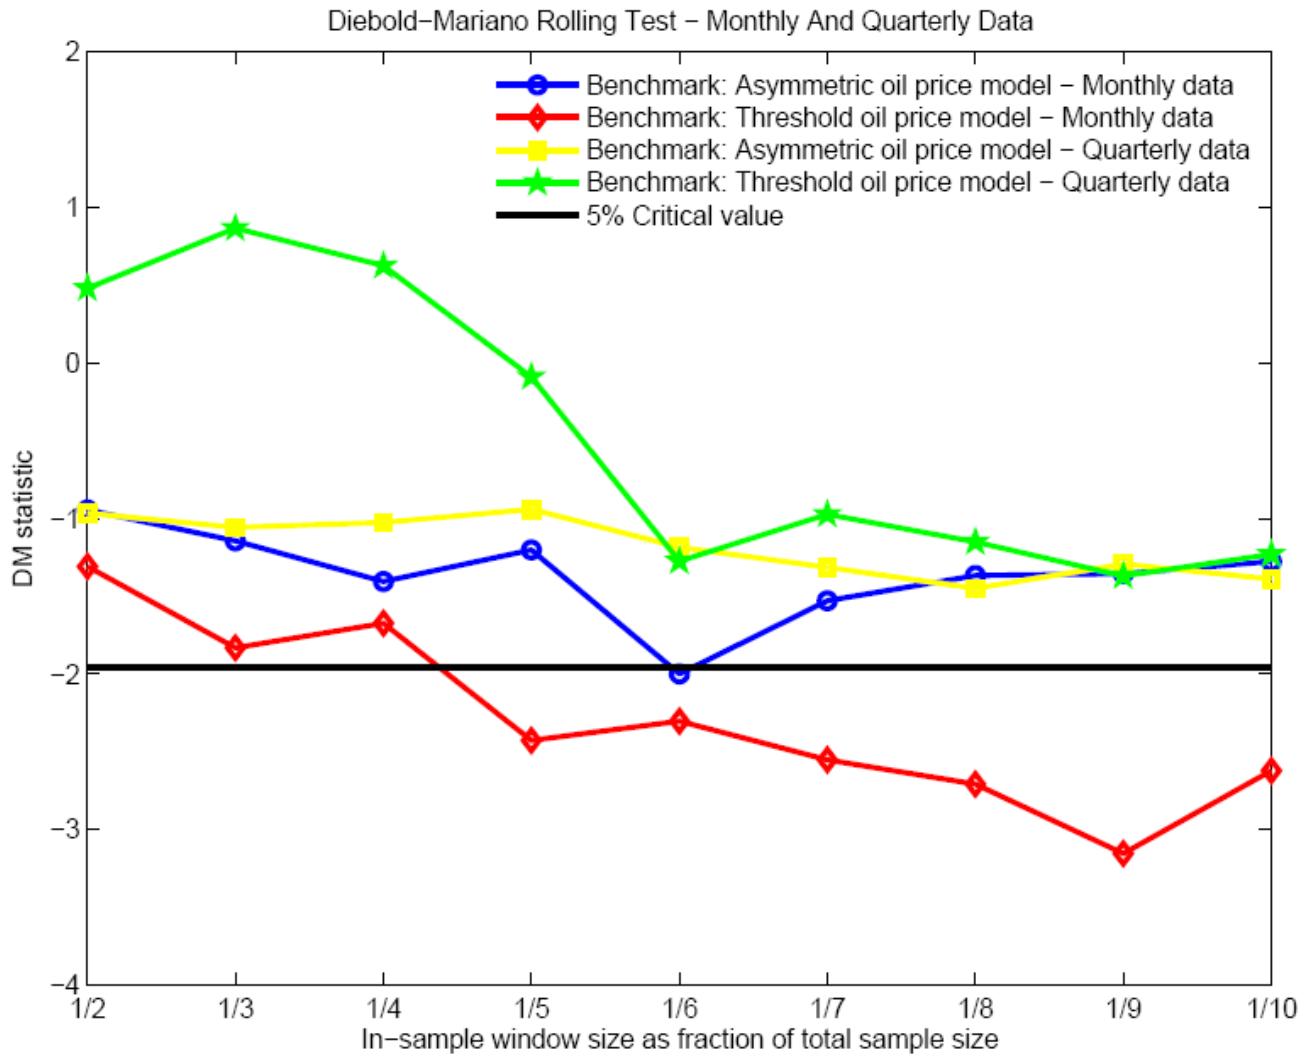
Notes to the Figure. Panel (a) reports the Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) calculated at daily, monthly and quarterly frequencies, and several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark. Panel (b) reports the same analysis for Model (3).

Figure 10(a). Asymmetric and Threshold Models. Forecasting Ability in Daily Data



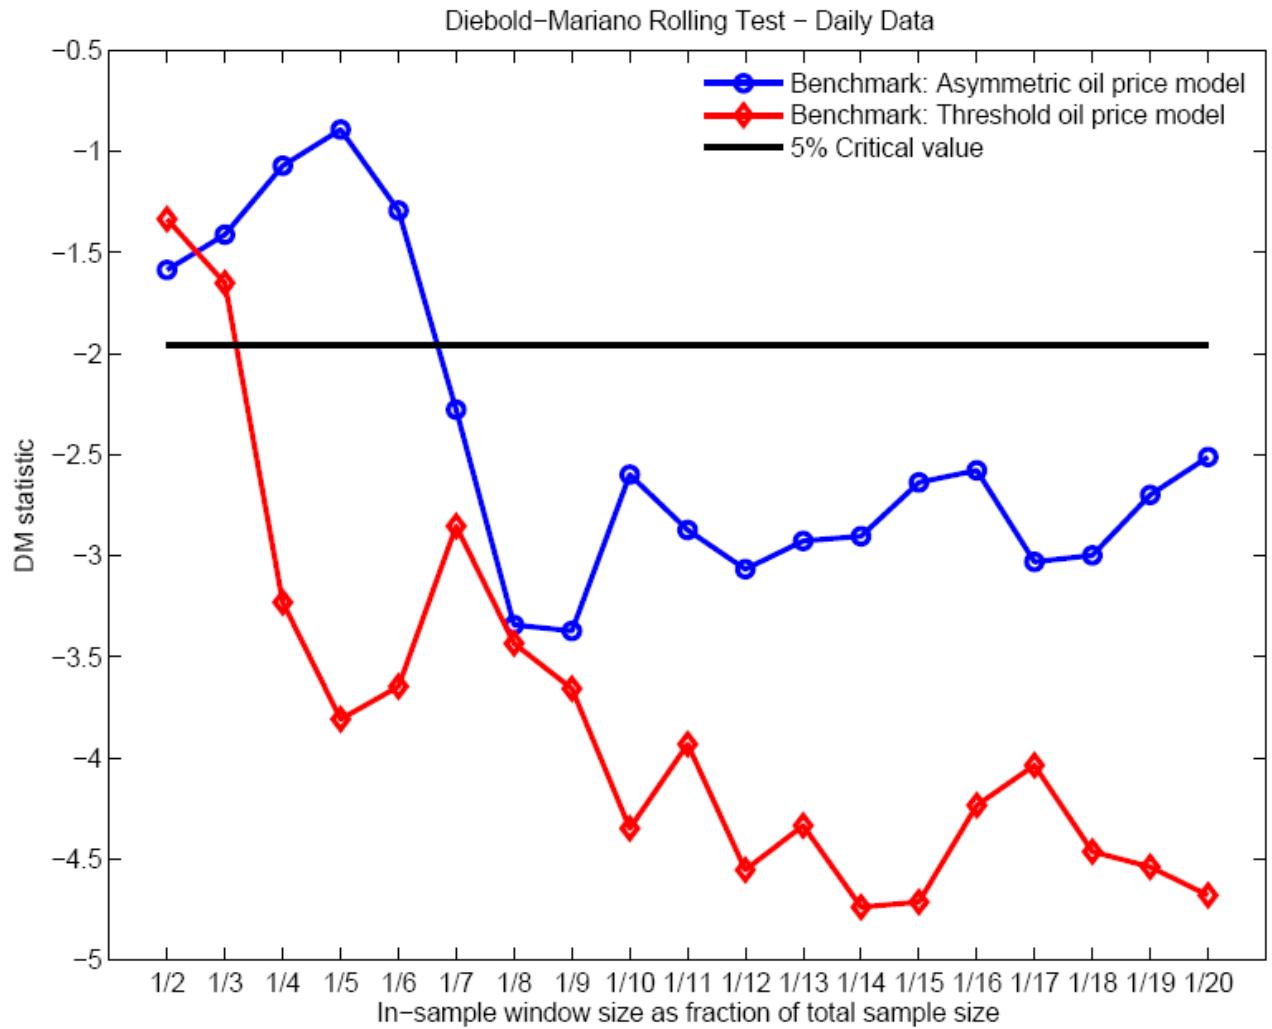
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to Model (5) (line with circles) as well as the forecasts of Model (1) relative to Model (6) (line with diamonds) calculated for daily data and several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark, and when it is above this line Model (1) forecasts worse.

Figure 10(b). Asymmetric and Threshold Models. Forecasting Ability in Monthly and Quarterly Data



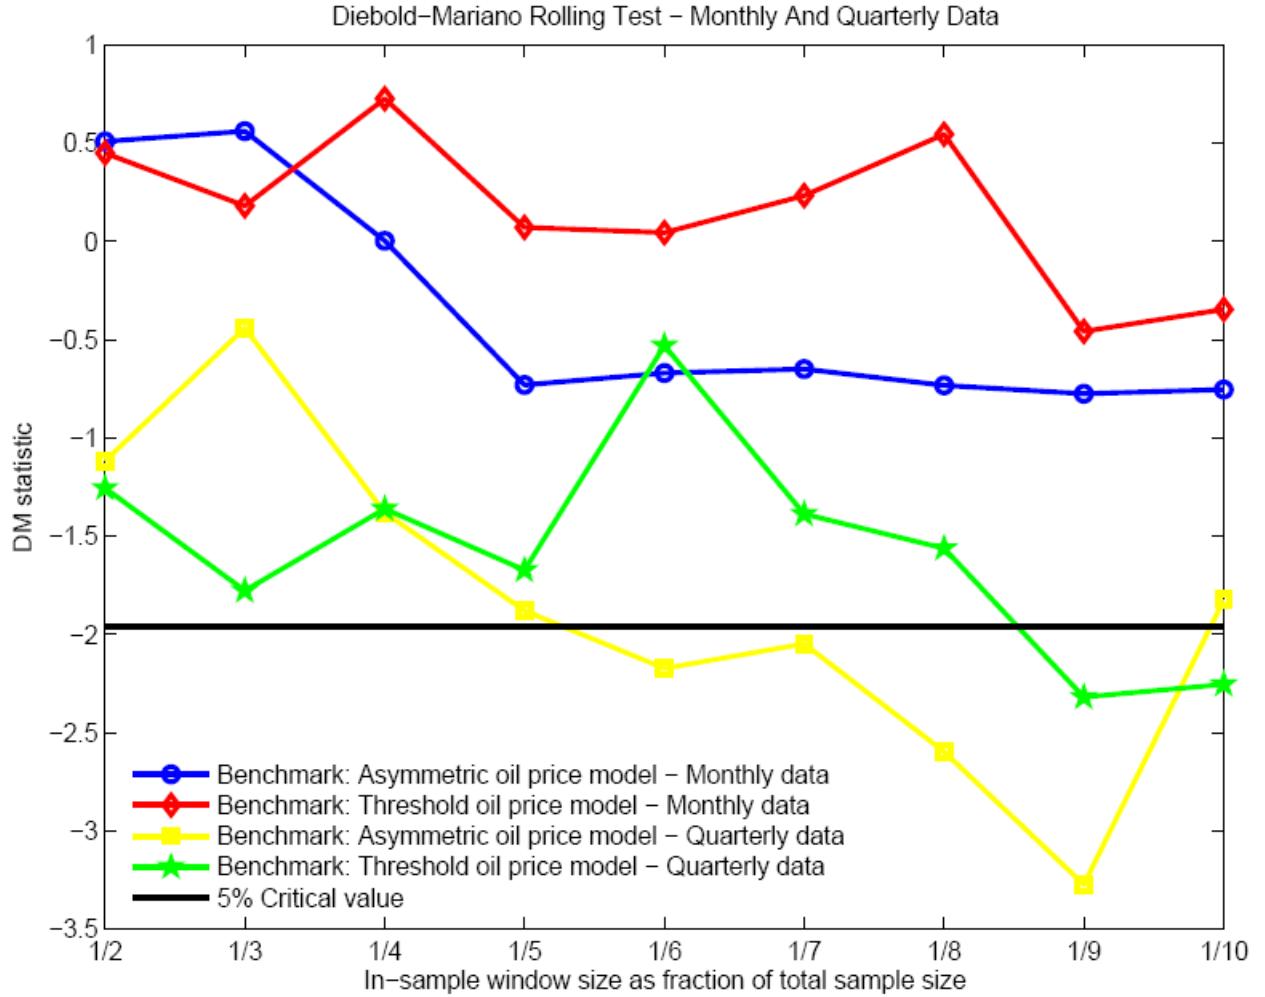
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (1) relative to Model (5) (line with circles for monthly data and line with squares for quarterly data) as well as the forecasts of Model (1) relative to Model (6) (line with diamonds for monthly data and line with stars for quarterly data) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (1) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (1) forecasts significantly better than its benchmark.

Figure 10(c). Asymmetric and Threshold Models. Forecasting Ability in Daily Data, Lagged Model



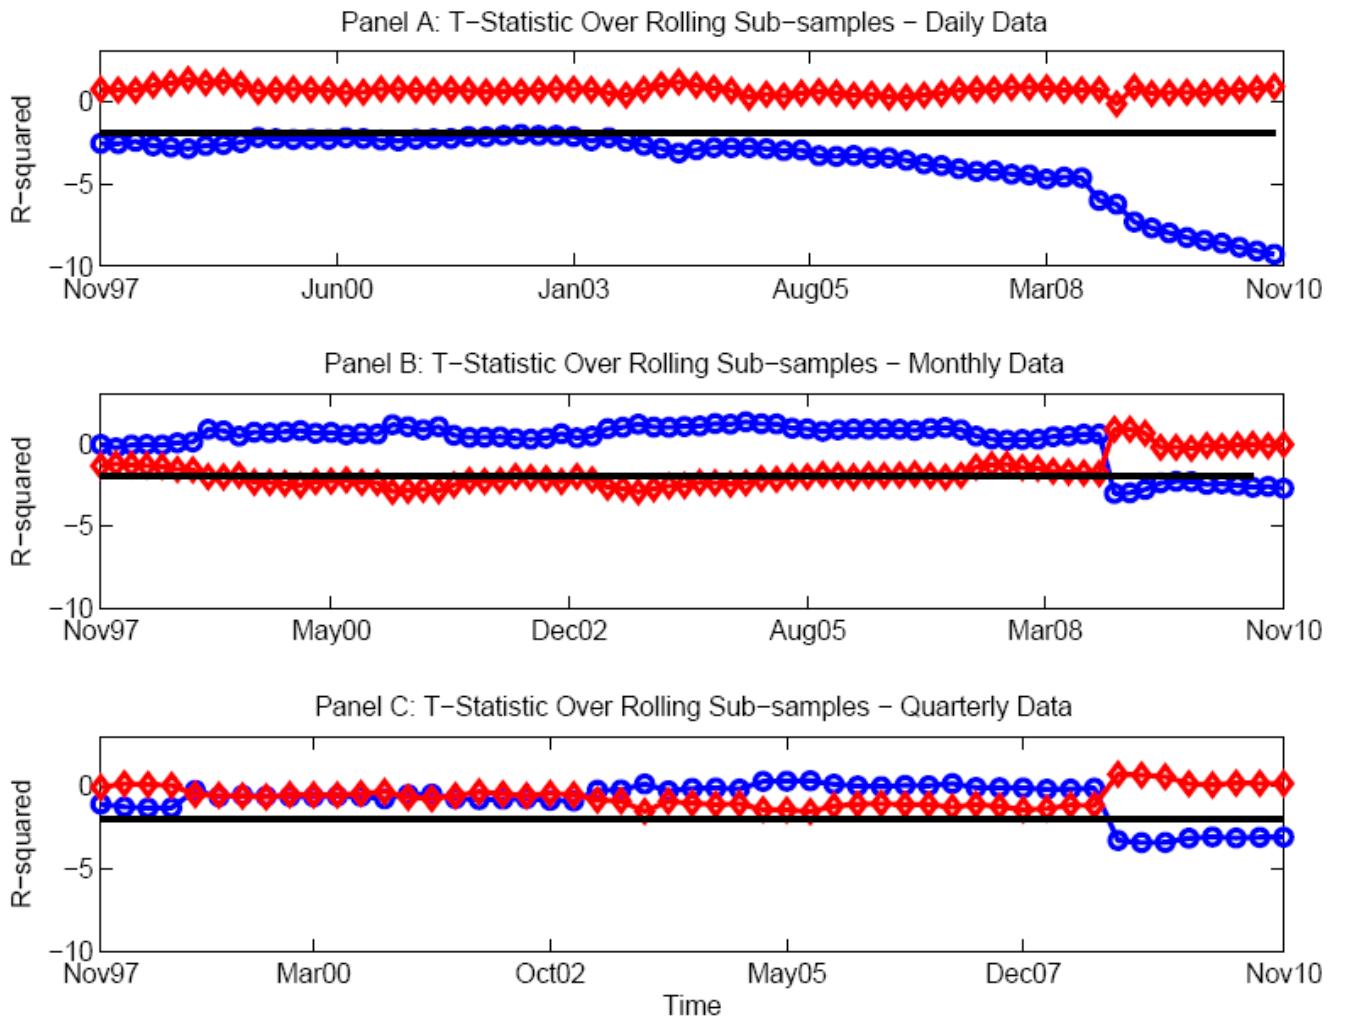
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (3) relative to Model (5) (line with circles for monthly data and line with squares for quarterly data) as well as the forecasts of Model (3) relative to Model (6) (line with diamonds for monthly data and line with stars for quarterly data) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (3) forecasts significantly better than its benchmark.

Figure 10(d). Asymmetric and Threshold Models. Forecasting Ability in Monthly and Quarterly Data, Lagged Model



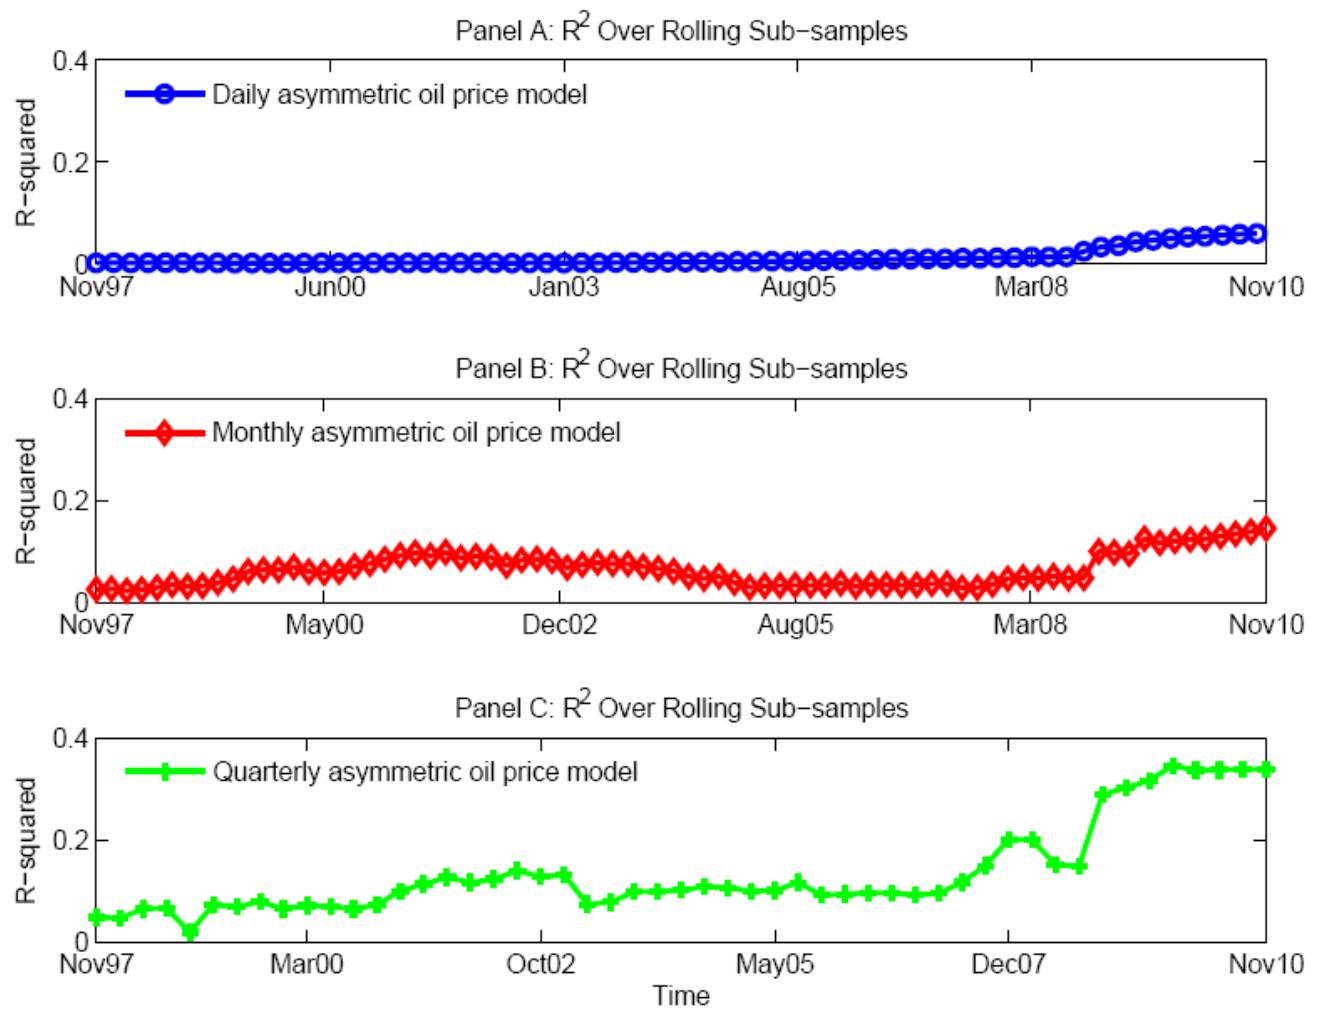
Notes to the Figure. The figure reports Diebold and Mariano's (1995) test statistic for comparing forecasts of Model (3) relative to Model (5) (line with circles for monthly data and line with squares for quarterly data) as well as the forecasts of Model (3) relative to Model (6) (line with diamonds for monthly data and line with stars for quarterly data) calculated for several in-sample window sizes (x-axis). Negative values indicate that Model (3) forecasts better. The continuous line indicates the critical value of the Diebold and Mariano's (1995) test statistic: when the estimated test statistics are below this line, Model (3) forecasts significantly better than its benchmark.

Figure 11(a). In-sample Fit of the Asymmetric Model – T-statistics Over Time



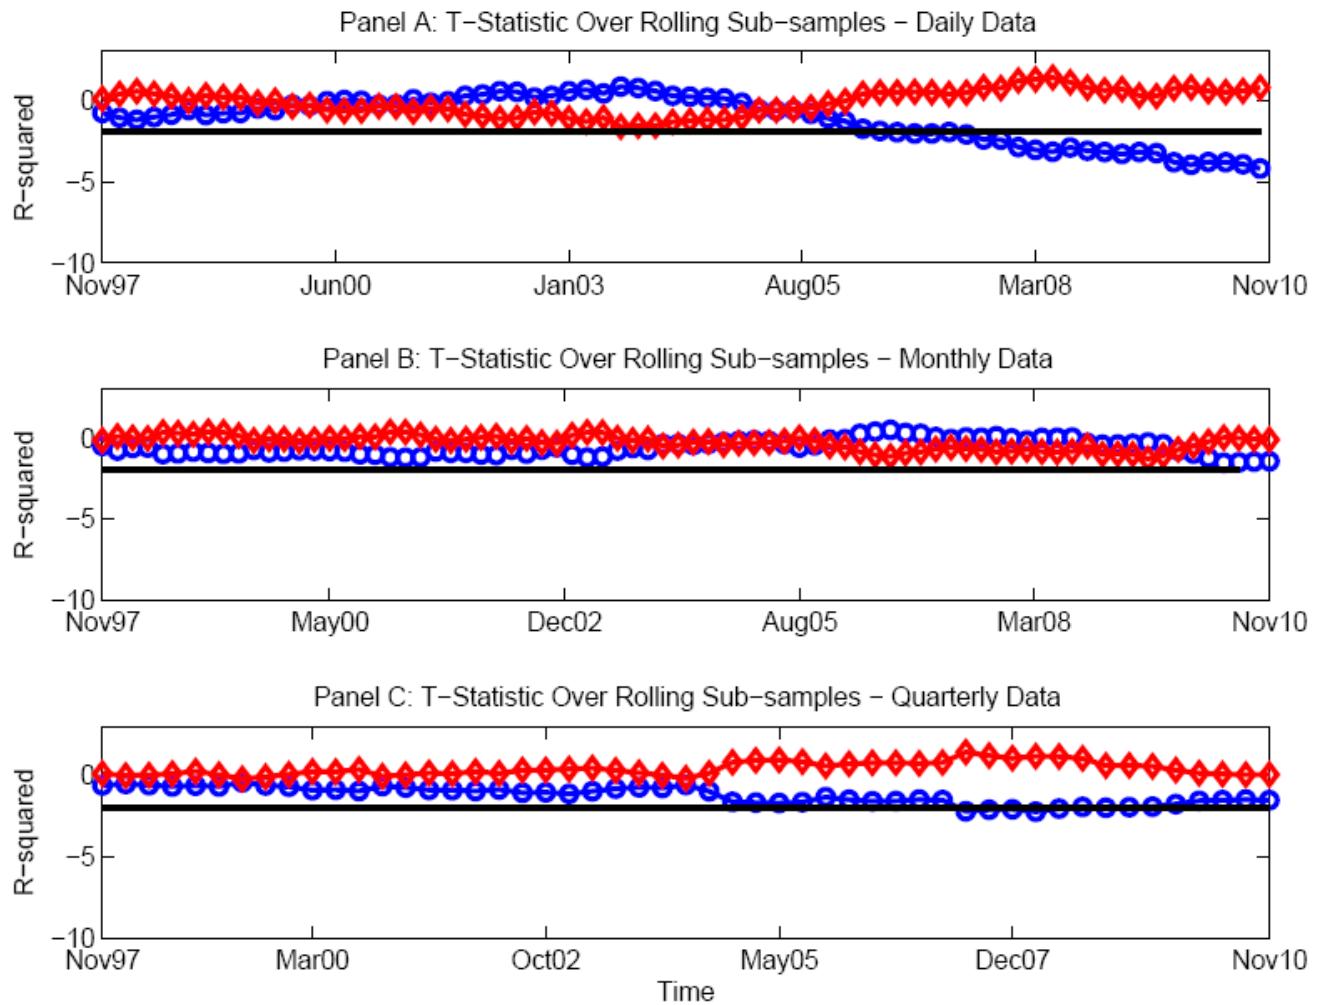
Notes to the Figure. The figure reports in-sample t-statistics for comparing forecasts of Model (1) calculated over rolling samples of size equal to 1/2 of the total sample size (dates reported on the x-axis). The line with circles is the t-statistic on the coefficient on the oil price growth rate and the line with diamonds is the t-statistic on the coefficient of the non-linear variable, calculated over rolling samples of data. The continuous line indicates the critical value of the t-statistic: if the estimated test statistics is below this line, the relevant coefficient in Model (1) is statistically significantly negative.

Figure 11(b). In-sample Fit of the Asymmetric Model – R^2 statistics Over Time



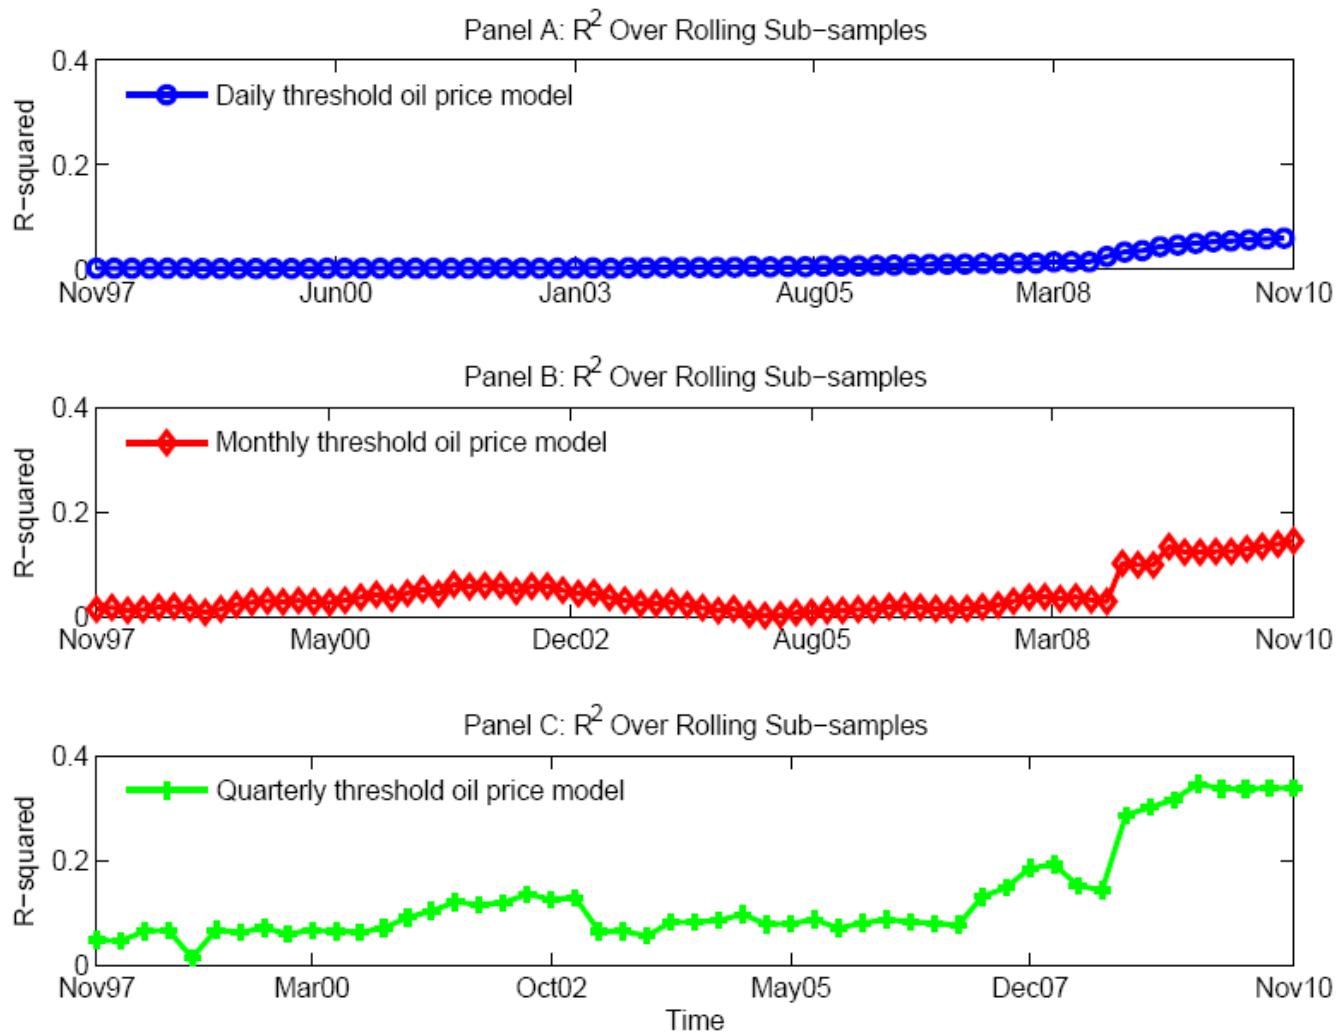
Notes to the Figure. The figure reports in-sample R^2 statistics for comparing forecasts of Model (1) calculated over rolling samples of 1/2 of the total sample size (dates reported on the x-axis).

Figure 12(a). In-sample Fit of the Threshold Model – T-statistics Over Time



Notes to the Figure. The figure reports in-sample t-statistics for comparing forecasts of Model (1) calculated over rolling samples of size equal to $1/2$ of the total sample size (dates reported on the x-axis). The line with circles is the t-statistic on the coefficient on the oil price growth rate and the line with diamonds is the t-statistic on the coefficient of the non-linear variable, calculated over rolling samples of data. The continuous line indicates the critical value of the t-statistic: if the estimated test statistics is below this line, the coefficient on the oil price in Model (1) is statistically significantly negative.

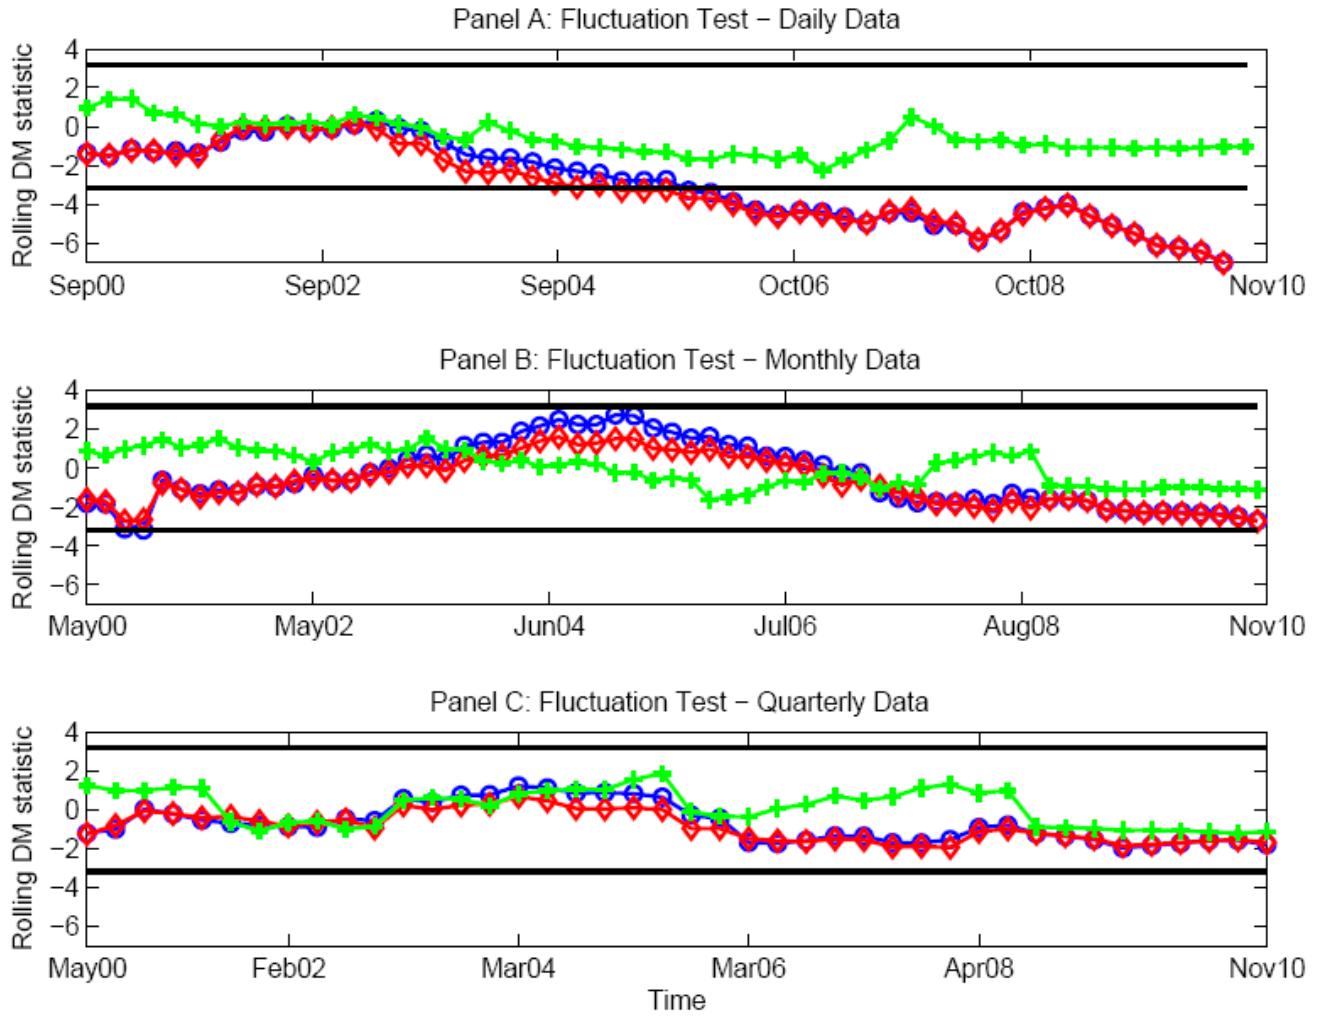
Figure 12(b). In-sample Fit of the Threshold Model – R^2 statistics Over Time



Notes to the Figure. The figure reports in-sample R^2 statistics for comparing forecasts of Model (1) calculated over rolling samples (dates reported on the x-axis).

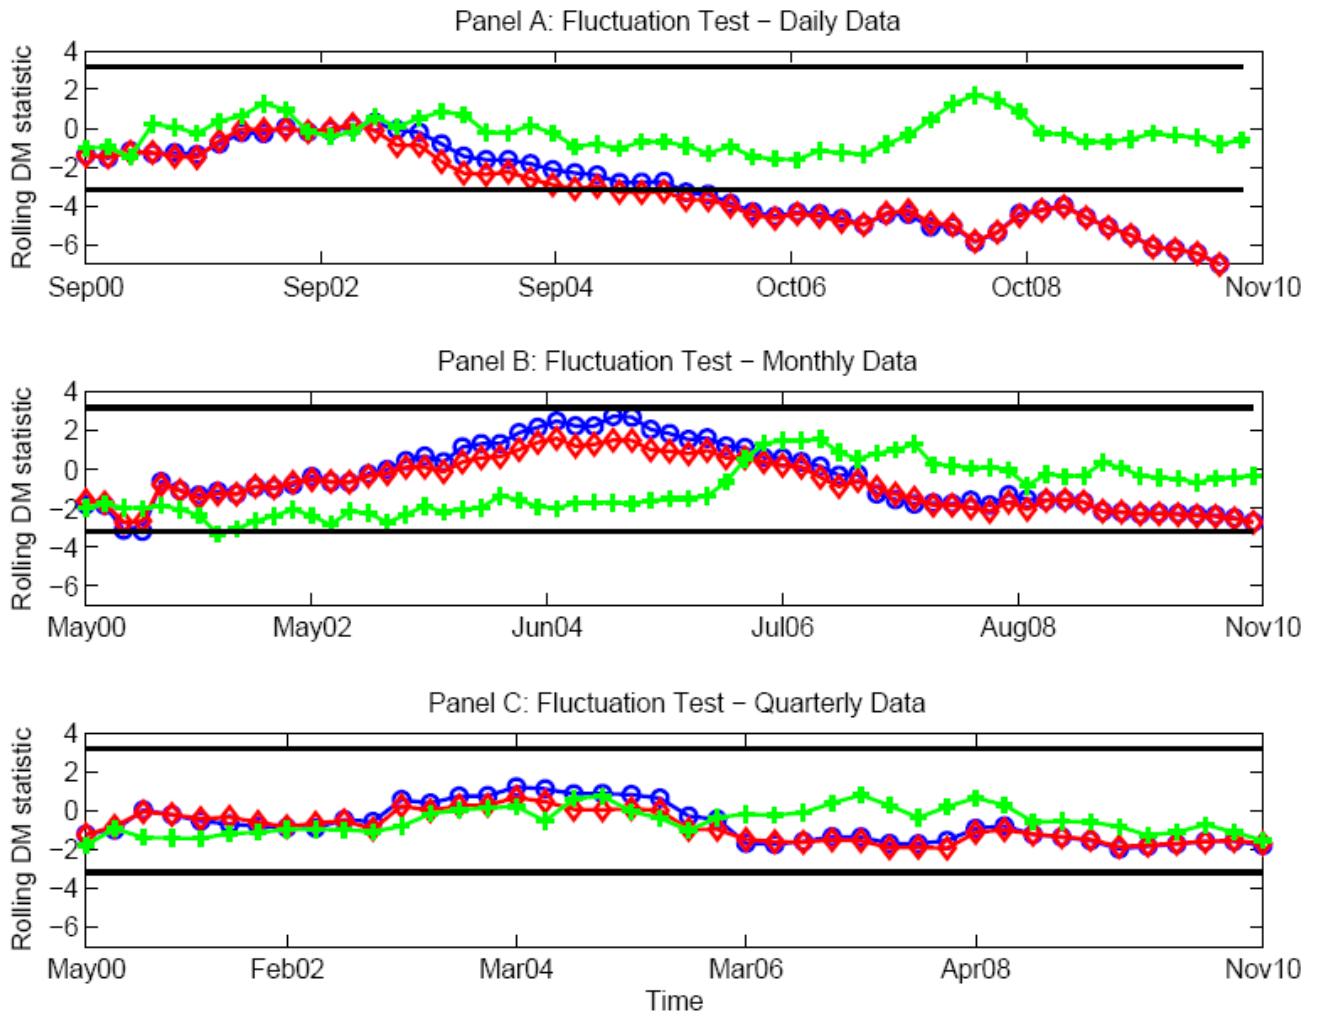
Appendix

Figure A.1



Notes to the Figure. The figure reports the Giacomini and Rossi's (2010) Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) and the Asymmetric Model (line with pluses). Negative values indicate that Model (5) forecasts better. The continuous line indicates the critical value of the Fluctuation test statistic: if the estimated test statistics is below this line, Model (5) forecasts significantly better than its benchmark.

Figure A.2



Notes to the Figure. The figure reports the Giacomini and Rossi's (2010) Fluctuation test statistic for comparing forecasts of Model (1) relative to a Random Walk without drift benchmark (line with circles) as well as relative to the Random Walk with drift benchmark (line with diamonds) and the Threshold Model (line with pluses). Negative values indicate that Model (6) forecasts better. The continuous line indicates the critical value of the Fluctuation test statistic: if the estimated test statistics is below this line, Model (6) forecasts significantly better than its benchmark.