Oil Price Shocks and Stock Return Predictability *

Lars Qvigstad Sørensen †
Norwegian School of Economics and Business Administration

First Draft: December 2008
This Draft: January 2009

Abstract

Recent research has documented that oil price changes lead the aggregate market in most industrialized countries, and has argued that it represents an anomaly - an underreaction to information that investors can profit from. I identify oil price changes that are caused by exogenous events and show that it is only these oil price changes that predict stock returns. The exogenous events usually correspond to periods of extreme turmoil - either military conflicts in the Middle East or OPEC collapses. Given the source of the predictability, I question its usefulness as a trading strategy.

*Thanks to Jonas Andersson, Thore Johnsen, Per Östberg, and seminar participants at the NHH Finance Seminar Series, and the FIBE 2009 Conference at NHH for helpful conversations and comments.
†E-mail: Lars.Sorensen@nhh.no

1 Introduction

The predictability of asset returns is one of the more controversial topics in financial economics. According to the Efficient Markets Hypothesis (EMH), all information in investors' information set should be incorporated into asset prices, thus leaving asset returns unpredictable. However, researchers and practitioners have identified a number of variables that seem to predict returns, at least in-sample. Some examples of variables that have been shown to predict aggregate market returns are inflation (Fama and Schwert, 1977), the default spread and the term premium (Fama and French, 1989) and the dividend-price ratio (e.g. Campbell and Shiller, 1988; Fama and French, 1988). Other financial ratios, such as the book-to-market ratio (e.g. Kothari and Shanken, 1997) and the earnings ratio (e.g. Lamont, 1998) have also been shown to predict returns. Recently, other variables that are not price-based have also been found to predict the aggregate stock market, such as the output gap (Cooper and Priestley, 2008). Traditionally, these variables have been thought to capture time-varying risk, a concept which is not in conflict with a modern interpretation of informationally efficient markets. In addition there is a growing literature that abandons the idea of informational efficiency; this literature argues that behavioral biases induce stock return predictability. In one recent example Baker and Wurgler (2007) argue that investor sentiment predicts returns. Another recent example is Hong, Torous and Valkanov (2007) who find that certain industries lead the aggregate market; they argue that this is due to underreaction to information.

In this paper I examine how changes in the oil price affect stock markets. Oil is by far the most important commodity and as of September 2008 energy accounted for around 75% of the Goldman Sachs Commodity Index (S&P GSCI). Crude oil accounted for 40% of the entire index and Brent crude oil for another 15%, thus leaving oil products with a weight of around 55% of the entire commodity index. The individual commodities in the index are weighted by their respective world production quantities. Given its large weight, it is not surprising that changes in oil prices may be an important factor for fluctuations in the economy. It is also plausible that oil price changes are important for understanding changes in stock prices. However, no consensus seems to have been reached regarding the impact of oil price changes on stock returns. Two recent papers, Driesprong, Jacobsen and Maat (2008) and Pollet (2004), document that oil price changes predict stock returns. Driesprong et al. (2008) show that the predictability is strongest for developed markets and less pronounced for emerging markets.

Goyal and Welch (2008) show that most predictors, such as the dividend-price ratio, have performed poorly over the last 30 years. They examine a number of variables which have been suggested as predictors of the equity premium and show that the statistical significance of these variables comes from the years up to and including the oil shock of 1973-1975. This period is exactly when oil prices start fluctuating and hence my sample using oil as a predictor starts in this period. I show that oil, unlike the other known predictors, seemingly predicts well both in-sample (IS) and out-of-sample (OOS). Goyal and Welch (2008) write that the signal from a predictive regression could give an investor confidence in the signal if it showed (i) good IS and reasonable OOS performance over the sample period, (ii) an upward drift in both IS and OOS plots, (iii) an upward drift not only in special periods (such as the outbreak of a war etc.).[1] I am particularly interested in the third point above; the oil price reacts sharply to e.g. military conflicts in the Middle East and OPEC crises. According to the advice from Goyal and Welch, an investor should not feel confident in the signal from the predictor if these military conflicts etc. are the only periods with good predictability. However, this is exactly what I find to be the case, thus casting doubt on the oil strategy's virtue as an equity investment strategy.

A large literature has examined the macroeconomic effects of oil price shocks in the U.S. with a particular emphasis on the effects on growth and inflation, see e.g. Barsky and Kilian (2004) and Hamilton (2008) for a review. In particular, some of the largest fluctuations in the U.S. economy have been preceded by political events in the Middle East that have reduced the oil supply from this region. For instance from 1973 to 1975 the oil price doubled, and oil importing countries experienced both inflation and rising unemployment. The later oil price increases during the Iranian revolution and leading up to the Iran-Iraq war also caused stagflation. Turmoil in the oil market has also had positive effects on the economy; following the OPEC collapse in 1985 and 1986 the aggregate supply curve for oil was shifted to the right, oil prices fell dramatically, output grew, unemployment fell and inflation was low. However, the positive effects of the oil price decrease were not as large as the adverse effects following the oil price increases in the 1970s. Based on the above-mentioned events, it has been suggested that U.S. macroeconomic variables respond in a nonlinear fashion to oil shocks. Mork (1989) showed that oil price increases affect the economy whereas decreases do not. Another transformation

[1]The plots are explained in Section 4.

has been suggested by Lee, Ni and Ratti (1995) who argue that *"an oil shock is likely to have greater impact in an environment where oil prices have been stable than in an environment where oil price movement has been frequent and erratic"*. They control for this by scaling the oil price by the conditional volatility as measured by a GARCH model. The effect of the Lee et al. (1995) transformation is that a small shock that occurs in a calm period will be scaled up whereas a large shock in a volatile period will be scaled down. Elaborating on these ideas, Hamilton (2003) further examines the functional form for oil and the macroeconomy.

I examine whether specific episodes, asymmetries and nonlinearities are important considerations for the predictability of equity returns. In doing so I build on the ideas in Hamilton (2003) and try to isolate an exogenous component of oil price changes that is attributable to supply shocks or fear of supply disruptions. Exogeneity here refers to oil price changes that are not caused by global or U.S. macroeconomic conditions but rather are caused by events that are exogenous with respect to the macroeconomy, such as political events in the Middle East. I start by constructing a counterfactual scenario where I remove historic events that led to large oil price changes. This significantly weakens the predictability, particularly OOS as demonstrated in Section 4.1. In subsequent sections I try to formalize this heuristic evidence by considering nonlinear transformations of the oil price changes; I construct measures that separate the oil price changes that arise because of exogenous events from those that do not; these are henceforth labeled exogenous and endogenous oil price changes. The exogenous changes pick out the events that were caused by wars and OPEC crises. I demonstrate that it is only the exogenous component that forecasts returns. Furthermore, the effect is asymmetric; large decreases in the oil price predict higher returns whereas large increases predict low returns only for a few countries. This asymmetric effect disappears once the oil price change is scaled by its time-varying volatility.

The paper proceeds as follows. Section 2 describes the data; Section 3 performs the initial linear regressions; Section 4 analyzes when oil price changes perform well as a linear predictor and when they do not; Section 5 performs regressions using nonlinear transformations of the oil price changes as the predictor; Section 6 discusses whether the predictability is truly an anomaly or whether it can be explained by time-varying discount rates; Section 7 concludes.

2 Data and descriptive statistics

2.1 The oil price

Throughout I use the West Texas Intermediate (WTI) price of crude oil as my measure of the oil price since it is often used as a benchmark in oil pricing and is also used as the underlying commodity price for the New York Mercantile Exchange's (NYMEX) oil futures contracts. Many studies that exclusively focus on the U.S. use the producer price index (PPI) for crude oil. However, since I consider both U.S. and international stock returns it seems less appropriate and furthermore, as noted by Mork (1989), it may be misleading during the price controls of the 1970s. The real price of crude oil is deflated with the consumer price index (CPI). Fig. 1 shows the nominal and the real WTI price of crude oil. The oil price was approximately constant from 1960 to 1973 when it rose sharply as a consequence of the Yom Kippur War and the oil embargo. I therefore start the sample in 1973:01 as this seems to be the first point in time where investors could reasonably expect there to be a relation between oil prices, the real economy and the stock markets. The grey bars in the graph mark important events in the oil market. As is evident from the graph, many of these events triggered huge changes in the oil price. The important oil events fall into two categories; they are either military conflicts in the Middle East or events concerning the Organization of the Petroleum Exporting Countries (OPEC). I shall focus on some of these episodes later in the paper and therefore explain briefly below the most important of the episodes that are marked with a grey bar in Fig. 1.

2.1.1 Oil supply shocks

In the subsequent analysis I treat an oil supply shock as exogenous with respect to the world economy, as has been common in the previous literature; for instance, the exogeneity of the oil price collapse in 1985-1986 was an important element in the analysis in Lamont (1997). A key question regarding each of these oil event dates is whether or not the event was unexpected. Arguably, they all contained elements of surprise, except perhaps the second Gulf war. The month of the events listed below and also in Fig. 1 is the month with the largest crude oil production disruption for the nation or group of nations most directly affected by the event.

The Yom Kippur war (October 1973) was fought from October 6 to October 26, 1973 between Israel and a coalition of Arab states. In response to the U.S. decision to support Israel during the war, the Organization of Arab Petroleum Exporting Countries (OAPEC) announced an oil embargo on October 15. As seen in Fig. 1, these events in 1973 resulted in large upward jumps in the oil price. After the military conflicts in 1973 the volatility of the oil price increased permanently from its pre 1973 low level. The Iranian revolution (November 1978) started early in 1978 and culminated when the Shah fled Iran in January 1979 and the monarchy was replaced by an Islamic republic under Ayatollah Khomeini. Real crude oil prices nearly doubled from the outbreak of the Iranian revolution to the Iran-Iraq war. The Iran-Iraq war (October 1980) lasted from September 1980 to August 1988 and October 1980 saw the highest real price of oil over the sample, matched only by the year-end 2007 price. The OPEC collapse (January 1986) marked the first large decrease in oil prices over the sample period. Saudi Arabia which had the swing-producer role in OPEC and had bore most of the production cuts abandoned its swing-producer role in late 1985 and aggressively increased its market share. In response, other OPEC members followed and the market was flooded with oil. The result was a sharp fall in oil prices in the end of 1985 and beginning of 1986. The Persian Gulf war (August 1990) started in August 1990 and ended in February 1991. The outbreak of the war resulted in a sharp upward spike in oil prices and an equally sharp decline shortly thereafter. The last three events in the graph did not cause such dramatic movements in the oil price. The OPEC meeting in 1999 marked a low point for the oil price before it again started rising; after the terrorist attacks in 2001, oil prices first declined but then rose sharply; the outbreak of the second gulf war started an upward trend in crude oil prices that persisted until year-end 2007.

2.1.2 Oil demand shocks

Oil demand shocks mainly describe increases or decreases in global demand for crude oil. The exceptional growth in China and India has resulted in large recent demand increases for crude oil. It is believed that demand from these two economies is responsible for much of the recent increase in oil prices. Oil price increases stemming from demand shocks have two effects. On the one hand, increased global aggregate demand represents a stimulus to the economy. On the other hand, the increased oil price may dampen that effect in oil-importing economies due to e.g. the increased cost of energy. It is not clear which of these two effects will dominate.

2.2 Equity data

I use mainly the MSCI equity indexes from 1972:12 to 2007:12 for the analyses across countries; all equity indexes are monthly and in the local currency. There are both total return and price indexes available, thus enabling computation of returns with and without dividend payments included. This also enables calculation of dividend payments and dividend-price ratios. I do not have industry level data available for the MSCI indexes, so for the industry level breakdown, I use the Datastream industry indexes. The countries I consider are the G7 countries, Norway and the World market. Norway has been included because it is the only major oil exporter for which equity data is available over the entire sample period. It is thus likely that the Norwegian market will respond in a different way to oil price changes than the major oil importing countries. I find that this is indeed the case. The first price quoted for the Datastream indexes is in 1973:01, so the first return available is in 1973:02. The exception is Norway for which Datastream index returns are available from 1980:02.

2.3 Interest rates and CPI data

Interest rates and CPI data are downloaded from Global Financial Data, except for Japan where the CPI data is downloaded from Datastream. I use the 3-month T-bill as the risk-free rate when computing the equity premium in all countries. For Norway, I have joined the 3-month bill with a 2-year government bond based on linear regressions because of a lack of data for the 3-month bill.

2.4 Descriptive statistics

The upper panel of Table 1 gives descriptive statistics for both real log oil price changes and real MSCI equity returns. Over the sample the real oil price is in fact the most volatile of the time series with an annual volatility of 28.61%. Norway has the highest mean real return over the sample, which may be caused by the large discoveries of oil on the Norwegian continental shelf. The first oil was extracted from the Norwegian continental shelf in 1971. Norway thus discovered the oil just in time to partake in the tumultous years following the OPEC price increases in 1973-1974 and again later in 1979-1980. The mean returns, mean returns in excess of the 90-day T-bill, the standard deviation of returns and the median are reported in annual terms. The min and max values are the min and max of a particular month and have not been annualized. The rightmost column reports the Jarque-Bera statistic. The null hypothesis of normality is rejected at all conventional significance levels for all the series. The critical value at the 5% level is 5.83.

3 Predictive regressions assuming a linear relation

I estimate the univariate predictive regressions for all G7 countries, Norway and the World market. In this section I consider regressions of the form

$$r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}, \tag{1}$$

where r is a log return and the right-hand side includes lagged values of oil price changes. Much of the literature on predictability has focused on the case where the explanatory variable follows a persistent AR(1) process which has close to a unit root. This causes problems with respect to inference and the possibility of spurious results as discussed in e.g. Stambaugh (1999). However, in our case, the change in the log oil price is used as the independent variable and from Table 1 the first order autocorrelation of changes in the log real oil price is 0.17, i.e. far from a unit root. Therefore, we need not be concerned about unit root induced biases or non-stationarity of the regressor.

I report the t-statistics both under strict OLS assumptions and correcting for heteroskedasticity and autocorrelation (HAC) using the Newey and West (1987) procedure.[2] Following Newey and West I include autocovariances up to five lags in the the monthly sample running from 1973:01 to 2007:12.[3] The results are found in Table 2 which re-

[2]The White heteroskedasticity consistent estimator gives similar results for the asymptotic covariance matrix of the coefficients.

[3]The lag length q is set to $q = floor\left(4\left(\frac{T}{100}\right)^{2/9}\right)$, where the floor function maps its argument to the nearest integer that is less than or equal to the argument.

ports the results for the univariate regressions in four cases. The upper panel considers nominal and real log returns. The bottom panel shows the results for excess rather than actual returns. Table 2 shows that changes in the oil price has strong predictive power for both nominal and real returns and for both actual and excess returns. Nominal oil price changes forecast nominal actual and excess returns in six out of the nine markets considered. Real oil price changes forecast real actual returns in seven and excess returns in six markets. Oil price changes thus seem to be a strong predictor of returns. However, one has to bear in mind that the nine markets considered are not independent. The bottom panel of Table 1 shows the correlations between the real actual returns below the diagonal. Covariances of the real actual returns are reported on and above the diagonal. Since the same regressor is used across all markets, we would expect returns in two highly correlated countries to be either both forecasted or both not forecasted by oil price changes. The coefficient on oil is negative in all markets and significant at the 5% level in all markets except Canada, Japan and Norway. It is interesting to note that oil price has no predictive power in Norway, which is the only country in the sample which is a major net oil exporter with a full history of equity returns over the sample. It is also of interest to note that the predictability is weak in Canada as well, a major oil producer albeit not a net oil exporter, and the country that exports the largest amount of crude oil to the U.S. Table 3 examines the predictability evidence further by using industry level data. Since I have no industry level data available for the MSCI indexes, I use the Datastream ICB industry indexes instead.[4] Not all returns series are available in all markets; series that are not available have been marked with a dash in the table. The two rightmost columns report the number of significant t-statistics by industry across countries and the total number of countries for which the time series were available. The two bottom rows report the number of significant t-statistics by country and again the total number of time series available per country. In the top panel, oil price changes are lagged one period; in the bottom panel, the oil price changes are contemporaneous. An interesting observation in the top panel is that to a large degree the same industries are forecasted by oil price changes across countries. I confirm the finding in Driesprong et al. (2008) that the one industry that is least predicted by oil is the oil and gas industry. Driesprong et al. (2008) interpret this as an underreaction to information or gradual diffusion of information as in Hong and Stein (1999). According to this theory, the oil and gas industry is not predicted by oil price changes because oil has a first order effect in this industry. Hence, information

[4] The ICB industries are Oil and Gas, Basic Materials, Industrials, Consumer Goods, Health Care, Consumer Services, Telecommunications, Utilities, Financials and Technology.

pertaining to the oil price is incorporated into prices without a lag. In other industries, where oil has an important, yet more second-order effect, it may take time for the information to be incorporated fully into prices. However, a challenge to this theory seems to be that the ICB supersector Automobiles is strongly forecasted in all markets where it is available.[5] This would contradict the claim that information about oil is quickly incorporated into industries where energy prices have a first-order effect. The automobiles industry is one of the industries that is most sensitive to petroleum prices.

Stock returns can be decomposed into the sum of risk-free rates and excess returns. The ability of the oil price to forecast returns is not due to its ability to forecast interest rates. First, as shown in Table 2 lagged oil price changes forecast both actual and excess returns. Second, unreported results show that when using real returns, lagged oil price changes forecasts interest rate changes only in Italy and marginally in Norway. To address causality further, Table 4 reports p-values from F-tests of Granger causality for various lags using excess returns and real oil price changes. Stock returns do not Granger cause oil price changes in any of the countries. Oil price changes Granger cause stock returns in the same countries as in Table 2 at lags 1 and 6. At 12 lags France and the World market become insignificant at the 5% level. In short, oil price changes lead stock returns, but there is no feedback from the stock markets to the oil the oil price.

The economic significance of oil price shocks as measured in the predictive regressions is large. The annualized standard deviation for real oil price changes is 8.26% per month (28.61% annualized). The beta coefficent on lagged oil for real actual returns is -0.11. This implies that a one standard deviation increase in the oil price leads to an annualized drop in expected returns in the U.S. of -0.91% per month or -10.90% annually.

4 A first stab at explaining the oil predictability

The in-sample predictability in the previous section seems impressive. In this section I try to scrutinize the predictability results to find out when and why oil prices predict returns. I examine both in-sample and out-of-sample predictability using a newly devel-

[5] Unreported evidence shows that the t-statistic on lagged oil price changes is significant in France, Germany, Italy, Japan, UK, U.S. and World. The Automobiles industry is not available for the full sample in Canada and Norway.

oped graphical tool suggested by Goyal and Welch (2003, 2008). The graphs represent a diagnostic tool to assess when the prediction from the univariate OLS model outperforms a random walk model. Hence, two models are considered: under the null hypothesis, the prediction for the equity premium at any point in time is just the prevailing mean, i.e.

$$E_{t-1}^{(pm)}\left[r_t^e\right] = \frac{1}{t-1}\sum_{i=1}^{t-1} r_{M,i}^e. \tag{2}$$

Under the alternative hypothesis, the equity premium is predicted by a regressor $x_t = \Delta o_t$, which is the change in the log oil price from t-1 to t. The regression equation is

$$r_{t+1}^e = \alpha + \beta x_t + \varepsilon_{t+1}. \tag{3}$$

To emphasize the recursive procedure used for the OOS forecasts, I use time subscripts on the estimated coefficients to denote the information included in their estimation. The prediction for the equity premium in period t+1 using only information available at time t is then

$$E_t^{(x)}\left[r_{t+1}^e\right] = \hat{\alpha}_t + \hat{\beta}_t x_t. \tag{4}$$

Note that only information up to time t is included in the recursive estimates of the parameters. Therefore, in the OOS models we get updated parameter estimates at each time point that is included in the estimation period. I use 20 time periods before I start estimating the model. Figure 2 shows the time-varying regression coefficient on lagged oil for the U.S. stock market. In the top panel I have plotted the important oil dates as grey bars; in the bottom panel I have plotted the NBER recession dates as grey bars. The width of the bar corresponds to the period when the economy went from peak to trough. The beta coefficient for the U.S. settles fairly quickly to a stable level after the turmoil in 1973. Figure 3 shows the beta coefficients on lagged oil for the remaining countries in the study. The evidence is now more mixed and in some countries the coefficient on lagged oil seems unstable. In particular it seems the OPEC collapse in 1985-1986 had an important, albeit differential, impact on the beta coefficient. In Italy the coefficient went from positive to negative over the course of the OPEC collapse.

The squared forecast errors in each period for the two models are

$$SE^{(pm)}(t) = \left(r_t^e - E_{t-1}^{(pm)}\left[r_t^e\right]\right)^2 \tag{5}$$

$$SE^{(x)}(t) = \left(r_t^e - E_{t-1}^{(x)}\left[r_t^e\right]\right)^2 . \tag{6}$$

Subtracting Eq. (6) from Eq. (5) gives a measure of how much better the OLS model predicted than the prevailing mean model in a given period. At each point in time τ, the cumulative OOS sum of squared errors is measured as

$$\text{Net-SSE}(\tau) = \sum_{t=1}^{\tau} SE^{(pm)}(t) - SE^{(x)}(t), \tag{7}$$

where $SE^{(pm)}$ is the squared OOS prediction error in year t when the prevailing mean is used as the forecast of the next return and $SE^{(x)}$ is the squared OOS prediction error when the oil price change is used as the predictor of the next return. If the graph is upward sloping, the equity premium is better predicted by lagged oil price changes than by a random walk. Whenever the graph decreases, the OLS model cannot predict better than the prevailing mean.

The in-sample coefficient estimates correspond to using the entire sample period in the estimate. This is equal to the last OOS coefficient estimates since we have used the entire history in the sample to compute the coefficients. The forecast errors in sample are obtained using the entire sample under consideration.

Figure 4 shows the graph for the U.S. Again I have split the figure into two panels. The top panel plots the important oil dates as grey bars; the bottom panel plots the NBER recession dates as grey bars. Note that a minimum number of data points is needed before we can start predicting the equity premium out of sample; I have chosen 20 initial data points. Therefore, we can predict IS before OOS and hence the red dotted line starts before the blue solid line. I have constructed the graphs such that they are both zero when the OOS forecast period starts. This means that if the OLS model has predicted better than the random walk for the initial 20 data points, the red dotted line is constructed to start below zero. Figure 4 reveals some noteworthy facts about the ability of lagged oil price changes to explain asset returns. It appears that the predictability comes from a few episodes only. Even more interesting is the fact that these episodes correspond almost exactly to what I labeled oil supply shocks above. It is in the periods identified as either military conflicts or OPEC crises that the predictability emerges. Three periods stand out in particular: the Yom Kippur war, The OPEC collapse and the first Gulf war. In the U.S. the predictability is strongest around July 1990 and the subsequent months. If we look at the second panel of Fig. 4 we see that this period corresponds to the peak of the business cycle as defined by NBER. It therefore appears that oil price changes predicted the recession. However, it may be just as plausible that the U.S. equity market plummeted for reasons other than oil price changes in this period.[6]

Figure 5 shows the graphs for the remaining countries in the study. Consistent with the results in Table 2, the predictability is weak in Canada and Japan and non-existent in Norway. Interestingly, the difference between the IS and OOS results are largest for the countries where oil price changes do not predict the equity premium. This indicates a more unstable relation between oil and asset prices in those markets. In the remaining countries, oil strongly forecasts returns. From the graphs in Fig. 5 we see that they share a common feature with the U.S. results: the exogenous oil shocks result in periods of remarkable predictability, in particular the episodes in 1973, 1985-1986 and 1990.

4.1 Removing "oil dates"

Based on the results in Figs. 4 and 5, I consider what happens when I remove dates that represent what I have identified as exogenous shocks. Consider a matrix of observations $\mathbf{Z}$ where each row consists of a vector

$$\mathbf{z}'_{t+1} = \left[r_{t+1}, \Delta o_t\right].$$

When removing certain dates from the sample, I remove a row $\mathbf{z}'_t$. This means that when I remove an exogenous oil event that has caused a dramatic change in the oil price, I also remove the subsequent reaction in the stock market. I remove the oil dates 1973:10-1973:11, 1978:10-1979:01, 1980:10, 1986:01-1986:08, 1990:07-1991:02 and 2003:02-2003:04, i.e. a total of 26 out of 419 $\mathbf{z}'_{t+1}$ entries in the $\mathbf{Z}$ matrix have been removed. The result of removing the oil dates is a significant drop in the predictability. Only three out of

[6] Mankiw (2003) writes that this recession was caused by *"several contractionary shocks to aggregate demand: tight monetary policy, the savings-and-loan crisis, and a fall in consumer confidence coinciding with the Gulf War"*.

nine equity premia, Germany, Italy and the United Kingdom, remain significantly predictable from oil price changes after removing these dates, thus illustrating that much of the predictability comes from these extreme events. Figure 6 shows the IS and OOS graphs for the countries where oil price changes were significant predictors prior to the removal of dates. The graphs show an increased discrepancy between the IS and OOS predictability after having removed the oil dates. The weaker OOS results mean that under the counterfactual scenario that the military and political events discussed in Section 2 did not take place, an equity investor following the oil signal would significantly underperform compared to a random walk strategy in France, the U.S. and World. In fact, the investor would never have been ahead of the random walk benchmark in these three markets, except for a brief period in France. In the remaining three markets with previous strong predictability results, the investor would still be slightly ahead of the random walk benchmark as of year-end 2007. However, had the sample ended earlier, the random walk strategy would have outperformed the OLS model in these three cases as well.

5 Functional form in the predictive regressions

There is no consensus regarding the effects oil price changes have on the stock markets. The analysis in Section 3 was conducted with a linear specification, as was done in Driesprong et al. (2008). Such a specification predicts a symmetric response from oil price increases and decreases. In this section I show that a linear specification seems inappropriate. Instead, I examine nonlinear specifications, some of which have been suggested in the macroeconomics literature (see e.g. Hamilton (2003)). If the true functional form is nonlinear, then inference based on a linear regression can be biased. In the remainder I work with the entire sample and not the sample where I had removed the important oil dates. For brevity I focus on the log equity premium and the real change in log oil prices. I assume that the predictive regression is of the form

$$r_{t+1} = f\left(\Delta o_t\right) + \varepsilon_{t+1}, \tag{8}$$

where the function f may represent a nonlinear relation. Even though I consider nonlinear transformations of the oil price, the transformations are still linear in the parameters.

Hence they can be described as linear regressions of the form

$$r_{t+1} = \alpha + \boldsymbol{\beta}'\tilde{\mathbf{x}}_t + \varepsilon_{t+1}, \tag{9}$$

for a corresponding specification of the regressor $\tilde{\mathbf{x}}_t$.

5.1 Positive and negative oil price changes

I start by examining a nonlinear transformation that was first suggested by Mork (1989). In particular, I examine the separate effects of oil price increases and oil price decreases as many papers in the macroeconomics literature have suggested that oil price increases affect the economy adversely, whereas oil price decreases do not affect the economy positively, at least not to a similar extent. The measure proposed by Mork (1989) to correct the relation between oil prices and GDP is

$$\Delta o_t^+ = \begin{cases} \Delta o_t & , \quad \Delta o_t > 0 \\ 0 & , \quad \Delta o_t \leq 0 \end{cases} . \tag{10}$$

Hamilton (2003) addresses the question of functional form in the case of the oil price and GDP and finds that the estimated effects of oil price increases are larger than those implied by the linear relation. However, it is not clear that this applies to the equity markets. First, stock market data are available at much higher frequencies than macro data, which are typically quarterly. Second, in an informationally efficient market, all available information should be incorporated into prices. A priori, it may therefore be just as likely that the markets react positively to an oil price decrease. I therefore also consider the negative equivalent

$$\Delta o_t^- = \begin{cases} \Delta o_t & , \quad \Delta o_t < 0 \\ 0 & , \quad \Delta o_t \geq 0 \end{cases} . \tag{11}$$

Table 5 shows the results which are interesting and perhaps surprising. The results for stock return predictability are quite the opposite of what has been found in the macroeconomics literature for the oil and GDP relation. It is the oil price decreases, not the increases, that predict stock returns more strongly. This finding is consistent with investment theory; risk averse investors with concave utility require a positive risk premium to invest in risky assets. From Eq. (9) the forecasted equity premium from the predictive regression is a constant plus the estimated coefficient multiplied with the measure of the oil price change. If a positive oil price change is large enough, it forecasts a negative risk premium according to this specification, and investors will sell stocks until a positive expected risk premium is restored. Campbell and Thompson (2008) impose the restriction that investors rule out a negative equity premium in their forecasts and truncate the forecast to zero whenever it is negative. They find that this improves the performance of the predictive regressions, similar to what I find above.

5.2 Extreme changes in the oil price over a horizon

Hamilton (2003) proposes another measure to gauge the effects of oil price shocks on GDP. He compares today's oil price against the maximum oil price attained over a prespecified horizon (12 months or 36 months). The measure of the oil shock is the amount by which the log oil price exceeds its maximum value over a previous period, e.g. one year or three years. If oil prices are lower than they have been at some point over this previous period, no oil shock is said to have occured. The idea is that it is the large oil price changes that affect the economy adversely and especially if they come unexpectedly. Hamilton's measure is

$$\Delta o_t^{\dagger,+} = \begin{cases} \ln\left(o_t / o_{t-1}^{\max}\right) & , \quad o_t > o_{t-1}^{\max} \\ 0 & , \quad o_t \leq o_{t-1}^{\max} \end{cases} , \tag{12}$$

where the max oil price is defined as

$$o_{t-1}^{\max} = \max\left\{o_{t-1}, \ldots, o_{t-k}\right\} \tag{13}$$

and k is the prespecified horizon. As in the above case, there is no reason to consider only the effects of positive oil price changes on stock prices. I therefore construct a corresponding measure that is non-zero if it is lower than the minimum oil price over the same prespecified horizon. The corresponding negative measure is

$$\Delta o_t^{\dagger,-} = \begin{cases} \ln\left(o_t / o_{t-1}^{\min}\right) & , \quad o_t < o_{t-1}^{\min} \\ 0 & , \quad o_t \geq o_{t-1}^{\min} \end{cases} , \tag{14}$$

where $o^{\min}_{t-1}$ is defined as in Eq. (13) with the min operator instead of the max. Figure 7 plots the oil price measures described in Eqs. (12) and (14). The figure shows that these transformations pick out exactly the exogenous events identified previously. I also define a third measure that considers extreme oil price changes, whether they be positive or negative

$$\Delta o^{\dagger,\pm}_t = \begin{cases} \ln\left(o_t/o^{\max}_{t-1}\right) & , \quad o_t > o^{\max}_{t-1} \\ 0 & , \quad o^{\min}_{t-1} \leq o_t \leq o^{\max}_{t-1} \\ \ln\left(o_t/o^{\min}_{t-1}\right) & , \quad o_t < o^{\min}_{t-1} \end{cases} . \tag{15}$$

The results from regressions using the three measures above as the explanatory variables are reported in Table 6. The top panel considers minimum and maximum returns over a period of 12 months whereas the lower panel sets the horizon to 36 months. Compared to the results in Table 5, the predictability from positive oil price changes are almost non-existent. Only two out of nine countries are predicted by $\Delta o^{\dagger,+}_t$ when $k = 12$ and three markets when $k = 36$. Surprisingly, Norwegian equity returns which have hitherto showed a remarkable resistance against any predictability, has a significant GMM-corrected t-statistic when $k = 36$. Again, the predictability is strongest for the negative oil price changes; Norwegian equity returns are the only ones that are not predicted by $\Delta o^{\dagger,-}_t$. The forecasting ability of $\Delta o^{\dagger,\pm}_t$ is also strong; returns in six of the countries have significant t-statistics when $k = 12$ and seven when $k = 36$. This is in sharp contrast to what Hamilton (2003) finds for the oil and GDP relation; this relation only becomes significant when considering the maximum measure.

We can think of the minimum and maximum transformations as picking out the oil price changes that are caused by exogenous events and setting the oil price changes not caused by these events to zero. I define the endogenous measures as the log change in the real oil price when the exogenous measures $\Delta o^{\dagger,+}_t$, $\Delta o^{\dagger,-}_t$ and $\Delta o^{\dagger,\pm}_t$ above are zero. The one corresponding to the maximum is thus

$$\Delta \tilde{o}^{\dagger,+}_t = \begin{cases} 0 & , \quad o_t > o^{\max}_{t-1} \\ \Delta o_t & , \quad o_t \leq o^{\max}_{t-1} \end{cases} , \tag{16}$$

the one corresponding to the minimum is

$$\Delta \tilde{o}^{\dagger,-}_t = \begin{cases} 0 & , \quad o_t < o^{\min}_{t-1} \\ \Delta o_t & , \quad o_t \geq o^{\min}_{t-1} \end{cases} , \tag{17}$$

and the one corresponding to both extremes is

$$\Delta\tilde{o}_t^{\dagger,\pm} = \begin{cases} 0 & , \quad o_t > o_{t-1}^{\max} \\ \Delta o_t & , \quad o_{t-1}^{\min} \leq o_t \leq o_{t-1}^{\max} \\ 0 & , \quad o_t < o_{t-1}^{\min} \end{cases} . \tag{18}$$

The results from estimating the predictive regressions using this measure are reported in Table 7. Arguably, the measures $\Delta o_t^{\dagger,+}$, $\Delta o_t^{\dagger,-}$ and $\Delta o_t^{\dagger,\pm}$ extract the exogenous component of the oil price changes, whereas the measures $\Delta\tilde{o}_t^{\dagger,+}$, $\Delta\tilde{o}_t^{\dagger,-}$ and $\Delta\tilde{o}_t^{\dagger,\pm}$ contain the pure endogenous change in oil prices. Together, Tables 6 and 7 then establish a very interesting fact: the asset return predictability comes from the exogenous changes in oil prices. The endogenous part has no power to forecast returns. This again illustrates that an oil-based equity trading strategy entails no riskless profits. An investor would have to be very confident that an upheaval in the oil market would result in a stock market reaction that went in the right direction. Since there are so few of these significant events, and each event had its own special characteristics, such confidence would perhaps be unwarranted.

5.3 Oil price changes scaled by GARCH volatility

I consider a last transformation of oil price changes that has achieved popularity in the macroeconomics literature. Lee et al. (1995) suggest a scaled oil price measure where the oil price is divided by the conditional standard deviation of the oil price as measured by a GARCH model. These authors use quarterly data and include four quarters in the conditional mean equation. I use monthly data so in order to be consistent with their measure I include 12 lags in the conditional mean equation. Thus, I compute a GARCH(1,1) model based on the following conditional mean equation

$$\Delta o_t = \phi_0 + \sum_{i=1}^{12} \phi_i \Delta o_{t-i} + a_t, \tag{19}$$

and conditional variance equation

$$a_t = \sigma_t \varepsilon_t, \qquad \varepsilon_t \sim NID\,(0,1) \tag{20}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 a_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \tag{21}$$

The volatility-adjusted oil price increase is then given by[7]

$$\Delta o_t^{\ddagger} = \Delta o_t / \sigma_t. \tag{22}$$

The aforementioned authors add another nonlinear transformation to this measure. They set the volatility-adjusted measure equal to zero if the oil price decreased. This transformation effectively reduces the impact of oil price changes if they occur in a volatile environment and increase the effects in quiet markets. As before I also examine what happens when I use the positive and negative parts only

$$\Delta o_t^{\ddagger,+} = \max\left\{\Delta o_t^{\ddagger}, 0\right\} \tag{23}$$

$$\Delta o_t^{\ddagger,-} = \min\left\{\Delta o_t^{\ddagger}, 0\right\}. \tag{24}$$

I estimate the GARCH(1,1) model using oil price data starting in 1946:01.

Table 8 shows the results when forecasts are made using the volatility adjusted oil price changes. The surprising result is that the asymmetry between positive and negative changes has disappeared. The same six countries are predicted by each of the three volatility adjusted measures. Unreported IS and OOS graphs show that the predictability comes from exactly the exogenous episodes identified previously and hence renders also this an undesirable predictor for an investment strategy.

6 An anomaly or time-varying discount rates?

Driesprong et al. (2008) claim that higher oil prices imply increased risk in the economy; therefore, if oil price changes predict the stock market, then discount rates should increase rather than decrease after a positive oil price shock. Since this is not the case, they conclude that an underreaction to information, i.e. an anomaly is the most likely cause of the predictability.

My analysis above may provide an alternative interpretation. First, it seems most of the predictability is associated only with specific episodes, as outlined in Section 2.

[7] Lee et al. (1995) use a_t/σ_t as their measure instead of $\Delta o_t/\sigma_t$. I use the latter for compatibility with the rest of the results in the paper.

Following Campbell (1991) unexpected excess returns in period t+1 can be decomposed into changes in expectations about future dividend growth, future interest rates and future excess returns

$$r_{t+1}^{e} - E_t\left[r_{t+1}^{e}\right] = (E_{t+1} - E_t)\left[\sum_{j=1}^{\infty}\rho^{j-1}\Delta d_{t+j}\right] - (E_{t+1} - E_t)\left[\sum_{j=1}^{\infty}\rho^{j-1}r_{t+j}^{f}\right] - (E_{t+1} - E_t)\left[\sum_{j=2}^{\infty}\rho^{j-1}r_{t+j}^{e}\right], \tag{25}$$

where r_{t+1}^{e} denotes the log excess return from period t to period t+1, Δd_{t+1} denotes log real dividend growth over the same period and r_{t+1}^{f} denotes the risk-free rate over the same period. The parameter ρ comes from the loglinearization and is a number a little smaller than one. This decomposition of returns says that unexpected returns must be either news about cash flows, news about risk-free rates, news about excess returns or a combination thereof. If we consider the unexpected returns associated with the exogenous events exclusively, a combination of news to dividends, risk-free rates and excess returns seems likely. The oil shocks resulted in permanent cash flow effects on the real economy and influenced montary policy. Since several of the shocks were associated with recessions in the economy, they are also likely to have led to revisions in expectations about future discount rates.

The claim that higher oil prices mean a riskier economy, which should lead to a higher risk premium and vice versa, requires an interpretation of an oil price shock as a cost shock to the economy or as a disturbance which requires a reallocation of capital or labor, thus inducing inefficiencies. Recent research suggests yet another interpretation; an oil price increase could also be the effect of positive aggregate demand shocks to commodities (Kilian, 2006). In these cases, the oil prices may be procyclical: high when discount rates are low and low when discount rates are high, consistent with time-varying equity risk premia.

6.1 Why do oil price decreases predict better than increases?

From finance theory, we know that an individual who is risk averse and who strictly prefers more to less will only hold risky assets if at least one of the risky assets has a positive expected risk premium. In equilibrium, investors cannot expect a negative equity premium. If the expected equity premium were negative, investors would want to short stocks and invest in the risk-free asset until a positive expected equity premium was restored. Consistent with this idea, negative predictions of the equity premium can be truncated to zero (see e.g. Campbell and Thompson, 2008). The fact that negative oil price changes, which imply a positive equity premium prediction, have better forecasting performance is thus consistent with a rational model.

Furthermore, recent research argues that demand shocks have become increasingly important. Kilian (2006) shows that not only supply disruptions but also aggregate demand shocks is an important reason for oil price increases. This partly explains why negative oil price changes predict returns better than positive changes. The recent increases in oil prices are caused by high global demand for global commodities. Higher oil prices thus have two opposite effects on the economy. On the one hand, the high demand for commodities represents a stimulus to the economy. On the other hand, the high oil prices may dampen that stimulus. It is not always clear which of the two effects will dominate. Several of the upward spikes in the two figures on the left in Fig. 7 have not been accompanied by negative returns, thus reducing the predictability from oil price increases.

7 Conclusion

I have shown that most of the predictive power of changes in oil prices comes from distinct episodes in the oil market. These episodes are military conflicts in the Middle East and OPEC crises. This means that investors should not be too confident in an equity trading strategy based on signals from the oil market.

Another question of interest is whether it is the oil price changes or the events associated with the military conflicts that affect the economy. Two questions then arise. First, why do oil price changes lead equity returns? Second, the OPEC collapse in 1985-1986 is not associated with a military conflict. Hence, the predictability arising from this episode is oil-specific.

The paper includes data from 1973 through 2007. The financial crisis of 2008 has seen dramatic movements both in the equity indexes that I consider and in the oil price; in fact the oil price has constantly been in the financial news in 2008 because of its extreme rise and fall during the year. The predictability from oil price changes has weakened considerably during the financial crisis as both oil prices and stock prices have plummeted, again illustrating the risks involved in following an oil-based investment strategy.

Recent research (Kilian, 2006) provides a way to disentangle supply and demand shocks in the crude oil market. A natural extension of the results in this paper would be to see how different countries respond differentially to the supply shock, aggregate demand shock for commodities and the oil-specific demand shock that Kilian identify. This is left for future research.

References

Baker, Malcolm P. and Jeffrey Wurgler, "Investor Sentiment in the Stock Market," *Journal of Economic Perspectives*, 2007, *21* (2), 129–151.

Barsky, Robert B. and Lutz Kilian, "Oil and the Macroeconomy since the 1970s," *Journal of Economic Perspectives*, Fall 2004, *18* (4), 115–134.

Campbell, John Y, "A Variance Decomposition for Stock Returns," *Economic Journal*, March 1991, *101* (405), 157–79.

Campbell, John Y. and Robert J. Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, Fall 1988, *1* (3), 195–228.

_ and Samuel B. Thompson, "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies*, 2008, *21* (4), 1509–1531.

Cooper, Ilan and Richard Priestley, "Time-Varying Risk Premiums and the Output Gap," *Review of Financial Studies,* forthcoming, 2008.

Driesprong, Gerben, Ben Jacobsen, and Benjamin Maat, "Striking Oil: Another Puzzle?," *Jounal of Financial Economics*, August 2008 2008, *89* (2), 307–327.

Fama, Eugene F. and G. William Schwert, "Asset returns and inflation," *Journal of Financial Economics*, November 1977, *5* (2), 115–146.

_ and Kenneth R. French, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, October 1988, *22* (1), 3–25.

_ and _ , "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, November 1989, *25* (1), 23–49.

Goyal, Amit and Ivo Welch, "Predicting the Equity Premium with Dividend Ratios," *Management Science*, 2003, *49*, 639–654.

_ and _ , "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 2008, *21* (4), 1455–1508.

Hamilton, James D., "What is an Oil Shock?," *Journal of Econometrics*, 2003, *113* (2), 363–398.

_ , "Oil and the Macroeconomy," in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, second ed., Palgrave Macmillan, 2008.

Hong, Harrison and Jeremy C. Stein, "A Unified Theory of Underreaction, Momentum Trading and Overreaction in Asset Markets," *Journal of Finance*, December 1999, *54* (6), 2143–2184.

— , Walter Torous, and Rossen Valkanov, "Do industries lead stock markets?," *Journal of Financial Economics*, February 2007, *83* (2), 367–396.

Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," CEPR Discussion Papers 5994, C.E.P.R. Discussion Papers December 2006.

Kothari, S.P. and Jay Shanken, "Book-to-market, Dividend Yield, and Expected Market Returns: A Time-series Analysis," *Journal of Financial Economics*, 1997, *44*, 169–203.

Lamont, Owen, "Cash Flow and Investment: Evidence from Internal Capital Markets," *Journal of Finance*, 1997, *52* (1), 83–109.

— , "Earnings and Expected Returns," *Journal of Finance*, Oct 1998, *53* (5), 1563–1587.

Lee, Kiseok, Shawn Ni, and Ronald A. Ratti, "Oil Shocks and the Macroeconomy: The Role of Price Variability," *Energy Journal*, 1995, *16* (4), 39–56.

Mankiw, N. Gregory, *Macroeconomics*, Worth Publishers, 2003.

Mork, Knut Anton, "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy*, 1989, *97* (3), 740.

Newey, Whitney K and Kenneth D West, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, May 1987, *55* (3), 703–08.

Pollet, Joshua Matthew, "Predicting Asset Returns With Expected Oil Price Changes," 2004. Harvard University Working Paper.

Stambaugh, Robert F., "Predictive Regressions," *Journal of Financial Economics*, 1999, *54*, 375–421.

Table 1: Descriptive statistics

The top panel shows descriptive statistics for log real oil price changes and log real returns for the G7 countries, Norway and the World market. The bottom panel shows correlations below the diagonal and covariances on and above the diagonal for the log real returns on the MSCI country indices. The column labeled "Mean Excess" is the mean equity premium over the period. The rest of the statistics in the table are for the actual real log returns, not the excess returns. The means, standard deviations and medians have been annualized.

	Mean	Mean Excess	St. Dev.	Median	Min	Max	Skew	Kurt	AC(1)	AC(2)	JB
OIL	4.89	-	28.61	-0.32	-34.98	36.31	0.20	6.89	0.17	-0.05	267.1
MSCNDA	6.99	3.02	16.96	8.27	-21.86	16.22	-0.46	5.16	0.04	-0.05	96.2
MSFRNC	8.94	4.42	20.57	14.75	-22.01	22.06	-0.13	4.13	0.08	-0.01	23.5
MSGERM	8.50	4.44	19.55	10.17	-24.86	21.29	-0.49	5.38	0.04	0.02	115.9
MSITAL	6.59	1.77	24.02	4.31	-18.94	26.76	0.40	4.14	0.06	0.01	33.9
MSJPAN	4.04	2.10	17.97	4.53	-20.18	20.18	0.00	4.19	0.06	0.04	24.6
MSNWAY	10.09	4.55	25.06	15.82	-29.76	25.26	-0.27	4.08	0.14	-0.06	25.7
MSUTDK	7.82	3.61	20.01	11.38	-26.29	50.42	1.11	17.15	0.10	-0.10	3,590.1
MSUSAM	6.75	4.33	15.39	8.77	-21.43	16.64	-0.31	4.79	0.02	-0.02	62.5
MSWRLD	5.58	3.39	13.75	8.66	-19.78	13.65	-0.66	5.26	0.09	0.01	119.8

Stock return correlations and covariances

	MSCNDA	MSFRNC	MSGERM	MSITAL	MSJPAN	MSNWAY	MSUTDK	MSUSAM	MSWRLD
MSCNDA	**24.50**	**16.00**	**12.62**	**12.67**	**9.21**	**19.70**	**16.11**	**16.43**	**15.11**
MSFRNC	0.54	**35.46**	**22.58**	**22.03**	**11.62**	**22.99**	**19.60**	**15.51**	**16.65**
MSGERM	0.45	0.66	**32.76**	**18.91**	**10.49**	**19.35**	**15.82**	**13.84**	**15.20**
MSITAL	0.37	0.54	0.48	**47.01**	**12.76**	**18.00**	**16.61**	**10.95**	**14.09**
MSJPAN	0.36	0.38	0.35	0.36	**27.00**	**9.86**	**9.96**	**8.64**	**13.40**
MSNWAY	0.54	0.53	0.46	0.36	0.26	**53.50**	**19.48**	**17.09**	**16.93**
MSUTDK	0.57	0.58	0.49	0.43	0.34	0.47	**32.14**	**16.00**	**16.47**
MSUSAM	0.74	0.58	0.54	0.36	0.37	0.52	0.63	**19.95**	**16.33**
MSWRLD	0.76	0.70	0.66	0.51	0.64	0.58	0.72	0.91	**16.14**

Table 2: Predictive regressions for the G7 countries, Norway and the World market

The table shows the beta coefficient, the t-statistic under strict OLS assumptions, the GMM-corrected t-statistic using Newey and West (1987) and the adjusted R^2. The top panel shows the results for the actual log nominal and real returns; the bottom panel shows the results for the log nominal and real excess returns. The log nominal and real excess equity returns are identical. Hence, the only difference between the results in the bottom panel is that the left bottom panel uses nominal oil returns whereas the right bottom panel uses real oil returns.

	Nominal				Real			
Country	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
				Log returns				
MSCNDA	-0.03	-1.19	-1.33	0.34	-0.04	-1.36	-1.50	0.44
MSFRNC	-0.13	**-3.83**	**-3.31**	3.39	-0.14	**-4.09**	**-3.49**	3.85
MSGERM	-0.15	**-4.52**	**-4.34**	4.67	-0.16	**-4.75**	**-4.48**	5.14
MSITAL	-0.21	**-5.30**	**-4.40**	6.31	-0.21	**-5.39**	**-4.54**	6.51
MSJPAN	-0.06	-1.91	-1.84	0.87	-0.07	**-2.28**	**-2.13**	1.23
MSNWAY	-0.02	-0.49	-0.44	0.06	-0.02	-0.57	-0.50	0.08
MSUTDK	-0.12	**-3.65**	**-4.69**	3.10	-0.13	**-3.97**	**-4.79**	3.64
MSUSAM	-0.09	**-3.72**	**-4.27**	3.22	-0.11	**-4.21**	**-4.77**	4.08
MSWRLD	-0.09	**-4.11**	**-4.40**	3.89	-0.11	**-4.63**	**-4.86**	4.90
				Log excess returns				
MSCNDA	-0.03	-1.14	-1.26	0.31	-0.03	-1.02	-1.12	0.25
MSFRNC	-0.13	**-3.80**	**-3.27**	3.34	-0.13	**-3.72**	**-3.19**	3.22
MSGERM	-0.15	**-4.51**	**-4.30**	4.65	-0.15	**-4.43**	**-4.20**	4.50
MSITAL	-0.20	**-5.18**	**-4.31**	6.04	-0.20	**-5.09**	**-4.23**	5.86
MSJPAN	-0.06	-1.91	-1.84	0.87	-0.06	-1.82	-1.74	0.79
MSNWAY	-0.02	-0.42	-0.37	0.04	-0.02	-0.36	-0.32	0.03
MSUTDK	-0.12	**-3.67**	**-4.64**	3.13	-0.12	**-3.61**	**-4.54**	3.04
MSUSAM	-0.10	**-3.73**	**-4.25**	3.23	-0.09	**-3.65**	**-4.08**	3.09
MSWRLD	-0.09	**-4.12**	**-4.39**	3.91	-0.09	**-4.01**	**-4.20**	3.71

Table 3: Significance in predictive regressions by industry and country

The table shows the t-statistic for the regression coefficient on nominal oil price changes when using the Newey and West (1987) procedure to correct for heteroskedasticity and autocorrelation. In the top panel, oil price changes are lagged one period; in the bottom panel, oil price changes are contemporaneous. The table also displays the number of significant industries and the total number of industries for which there exists a full set of observations in each market. t-statistics are marked in bold. In this table, the stock returns are the nominal Datastream indices as I do not have industry data for the MSCI indices available.

	DSCNDA	DSFRNC	DSGERM	DSITAL	DSJPAN	DSNWAY	DSUTDK	DSUSAM	DSWRLD	Sign.	Tot.
Returns regressed on lagged change in oil prices											
TOTMK	-1.93	**-3.22**	**-4.39**	**-4.26**	-1.79	-0.39	**-4.60**	**-4.17**	**-3.29**	6	9
OILGS	0.03	-1.24	-	-	-0.31	0.75	**-2.71**	-1.18	-1.39	1	7
BMATR	**-2.10**	**-2.79**	**-3.22**	**-3.93**	-1.18	**-2.20**	**-2.72**	**-2.20**	-1.90	7	9
INDUS	**-2.07**	**-3.01**	**-3.31**	**-3.39**	**-2.42**	-0.89	**-3.65**	**-3.44**	**-3.06**	8	9
CNSMG	-	**-2.87**	**-3.62**	**-4.38**	**-2.48**	-	**-3.24**	**-3.66**	**-2.81**	7	7
HLTHC	**-2.52**	**-2.15**	**-3.52**	-	-0.82	-	**-3.30**	**-2.86**	**-2.65**	6	7
CNSMS	**-3.57**	**-3.42**	**-3.45**	**-2.78**	**-2.08**	-0.44	**-4.43**	**-4.28**	**-3.83**	8	9
TELCM	-	-	**-3.71**	**-5.54**	0.00	-	-	**-2.49**	**-2.17**	4	5
UTILS	-0.75	-	**-3.73**	**-3.36**	-0.97	-0.53	-	-0.91	-1.41	2	7
FINAN	-1.69	**-3.19**	**-3.59**	**-3.70**	-0.42	0.31	**-3.85**	**-2.93**	**-2.29**	6	9
TECNO	-1.24	**-3.17**	-	-	**-2.93**	-	**-3.23**	**-3.57**	**-3.96**	5	6
Significant	4	8	9	8	4	1	9	9	8	60	84
Total	9	9	9	8	11	7	9	11	11	84	-
Returns regressed on contemporaneous change in oil prices											
TOTMK	1.46	-1.18	-1.03	-0.29	-0.05	**3.07**	-1.29	-1.71	-0.75	1	9
OILGS	**4.54**	**4.09**	-	-	0.44	**6.76**	**3.43**	**4.00**	**3.71**	6	7
BMATR	-0.16	-0.97	-1.23	0.12	-0.32	1.18	-0.49	-1.59	-0.49	0	9
INDUS	-0.13	-1.12	-0.74	-0.25	0.63	-0.13	-0.87	-1.70	-0.48	0	9
CNSMG	-	-1.95	-1.30	-1.23	0.60	-	-1.00	**-3.04**	-0.81	1	7
HLTHC	-0.68	**-3.15**	-1.61	-	-0.60	-	**-3.94**	**-3.92**	**-3.11**	4	7
CNSMS	-1.57	-1.58	**-1.96**	0.34	-0.63	0.39	-1.54	**-3.41**	-1.86	2	9
TELCM	-	-	-0.34	-1.73	0.55	-	-	-1.31	-0.55	0	5
UTILS	-0.23	-	-1.24	0.32	**-2.86**	-0.63	-	-1.43	**-2.10**	2	7
FINAN	-1.79	**-2.46**	-1.16	-0.19	-0.43	-0.26	**-2.48**	**-3.40**	-1.70	3	9
TECNO	-0.40	-1.16	-	-	1.13	-	0.71	-0.47	0.04	0	6
Significant	1	3	1	0	1	2	3	5	3	19	84
Total	9	9	9	8	11	7	9	11	11	84	0

Table 4: Granger causality between stock returns and oil price changes

The table shows the p-values of F-tests of bivariate Granger causality using the number of lags specified in the header. The null hypothesis is that the variable on the right side of the arrow does not enter the reduced form equation for the variable on the left side of the arrow. The instances where one variable Granger causes the other at the 5% level are in bold.

	1 lag		6 lags		12 lags	
Country	$\Delta o \overset{G}{\nrightarrow} r$	$r \overset{G}{\nrightarrow} \Delta o$	$\Delta o \overset{G}{\nrightarrow} r$	$r \overset{G}{\nrightarrow} \Delta o$	$\Delta o \overset{G}{\nrightarrow} r$	$r \overset{G}{\nrightarrow} \Delta o$
MSCNDA	0.27	0.39	0.82	0.23	0.95	0.86
MSFRNC	**0.00**	0.97	**0.02**	0.85	0.24	0.74
MSGERM	**0.00**	0.77	**0.00**	0.26	0.09	0.53
MSITAL	**0.00**	0.93	**0.00**	0.94	**0.01**	0.55
MSJPAN	0.07	0.78	0.07	0.18	0.11	0.59
MSNWAY	0.45	0.36	0.35	0.15	0.52	0.30
MSUTDK	**0.00**	0.93	**0.01**	0.86	**0.02**	0.95
MSUSAM	**0.00**	0.20	**0.02**	0.29	**0.03**	0.84
MSWRLD	**0.00**	0.57	**0.01**	0.51	0.06	0.78

Table 5: Predictive regressions using positive or negative oil price changes

The table shows the beta coefficient, the t-statistic under strict OLS assumptions, the GMM-corrected t-statistic using Newey and West (1987) and the adjusted R^2. The left panel considers positive oil price changes only; the right panel considers negative oil price changes only. The rightmost column reports the p-value for the null hypothesis that the coefficients on positive and negative oil price changes are identical in a regression where they are both included as explanatory variables.

	Δo_t^+				Δo_t^-				
Country	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	$p(b^+ = b^-)$
MSCNDA	0.01	0.15	0.13	0.01	-0.09	-1.91	**-2.94**	0.87	0.08
MSFRNC	-0.12	**-2.23**	**-2.03**	1.18	-0.22	**-3.81**	**-3.99**	3.36	0.15
MSGERM	-0.18	**-3.43**	**-2.94**	2.74	-0.21	**-3.64**	**-3.78**	3.09	0.73
MSITAL	-0.15	**-2.45**	**-2.41**	1.41	-0.39	**-5.91**	**-5.80**	7.72	0.00
MSJPAN	-0.03	-0.60	-0.55	0.09	-0.12	**-2.42**	**-2.49**	1.38	0.11
MSNWAY	0.01	0.21	0.17	0.01	-0.06	-0.83	-0.74	0.17	0.38
MSUTDK	-0.15	**-2.92**	**-3.09**	2.01	-0.16	**-2.85**	**-4.22**	1.91	0.94
MSUSAM	-0.09	**-2.19**	**-2.04**	1.14	-0.16	**-3.72**	**-4.58**	3.22	0.16
MSWRLD	-0.08	**-2.17**	-1.84	1.12	-0.17	**-4.36**	**-6.12**	4.35	0.05

Table 6: Predictive regressions using maximum and minimum measures of returns

The table shows the beta coefficient, the t-statistic under strict OLS assumptions, the GMM-corrected t-statistic using Newey and West (1987) and the adjusted R^2. The left panel shows positive price changes only, the middle panel shows negative price changes only and the right panel shows both positive and negative oil price changes as described in Eqs. (12), (14), and (15). The upper panel uses a horizon of 12 months to calculate the maximum or minimum return whereas the lower panel uses 36 months.

	$\Delta o_t^{\dagger,+}$				$\Delta o_t^{\dagger,-}$				$\Delta o_t^{\dagger,\pm}$			
Country	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Considering max/min returns using horizon k=12												
MSCNDA	-0.03	-0.53	-0.59	0.07	-0.10	-1.33	**-2.26**	0.42	-0.06	-1.23	-1.74	0.36
MSFRNC	-0.10	-1.24	-1.27	0.36	-0.33	**-3.66**	**-5.86**	3.11	-0.18	**-3.22**	**-3.01**	2.43
MSGERM	-0.14	-1.84	-1.86	0.81	-0.20	**-2.27**	**-2.42**	1.22	-0.15	**-2.79**	**-3.15**	1.84
MSITAL	-0.07	-0.78	-0.99	0.14	-0.55	**-5.48**	**-8.26**	6.72	-0.26	**-4.00**	**-2.67**	3.70
MSJPAN	-0.06	-0.84	-0.67	0.17	-0.18	**-2.35**	**-2.70**	1.30	-0.10	**-2.11**	-1.61	1.05
MSNWAY	-0.13	-1.34	-1.81	0.43	0.01	0.06	0.06	0.00	-0.07	-0.94	-0.84	0.21
MSUTDK	-0.20	**-2.69**	**-2.71**	1.71	-0.20	**-2.29**	**-4.66**	1.25	-0.19	**-3.43**	**-4.32**	2.75
MSUSAM	-0.10	-1.73	**-2.12**	0.71	-0.16	**-2.45**	**-3.67**	1.42	-0.12	**-2.82**	**-4.29**	1.88
MSWRLD	-0.09	-1.74	-1.55	0.72	-0.20	**-3.34**	**-5.36**	2.61	-0.13	**-3.40**	**-3.68**	2.69
Considering max/min returns using horizon k=36												
MSCNDA	-0.05	-0.74	-0.88	0.13	-0.12	-1.51	**-2.19**	0.54	-0.08	-1.52	**-2.28**	0.55
MSFRNC	-0.10	-1.20	-1.23	0.34	-0.35	**-3.62**	**-6.16**	3.04	-0.20	**-3.20**	**-2.99**	2.40
MSGERM	-0.12	-1.55	-1.61	0.57	-0.15	-1.65	**-2.34**	0.65	-0.13	**-2.21**	**-2.68**	1.16
MSITAL	-0.05	-0.48	-0.61	0.06	-0.59	**-5.44**	**-7.54**	6.62	-0.26	**-3.77**	**-2.37**	3.30
MSJPAN	-0.04	-0.58	-0.46	0.08	-0.19	**-2.25**	**-2.55**	1.20	-0.10	-1.87	-1.36	0.83
MSNWAY	-0.17	-1.65	**-2.61**	0.65	0.05	0.39	0.46	0.04	-0.07	-0.98	-0.83	0.23
MSUTDK	-0.19	**-2.38**	**-2.33**	1.35	-0.19	**-2.03**	**-4.59**	0.97	-0.18	**-3.09**	**-3.75**	2.23
MSUSAM	-0.11	-1.77	**-2.22**	0.74	-0.18	**-2.46**	**-3.20**	1.43	-0.13	**-2.90**	**-4.37**	1.97
MSWRLD	-0.09	-1.57	-1.36	0.58	-0.20	**-3.10**	**-4.98**	2.26	-0.13	**-3.16**	**-3.32**	2.33

Table 7: Predictive regressions excluding the maximum and minimum measures of returns

The table shows the beta coefficient, the t-statistic under strict OLS assumptions, the GMM-corrected t-statistic using Newey and West (1987) and the adjusted R^2. The left panel shows positive price changes only, the middle panel shows negative price changes only and the right panel shows both positive and negative oil price changes, i.e. the three headers correspond to using the oil price measures in Eqs. (16)–(18) as the predictor. The upper panel uses a horizon of 12 months to calculate the maximum or minimum return whereas the lower panel uses 36 months.

	$\Delta\tilde{o}_t^{\dagger,+}$				$\Delta\tilde{o}_t^{\dagger,-}$				$\Delta\tilde{o}_t^{\dagger,\pm}$			
Country	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Considering max/min returns using horizon k=12												
MSCNDA	0.002	1.00	1.08	0.24	-0.001	-0.63	-0.47	0.09	0.000	0.33	0.29	0.03
MSFRNC	-0.001	-0.54	-0.57	0.07	0.004	**2.03**	1.66	0.98	0.002	1.13	0.97	0.31
MSGERM	-0.002	-1.26	-1.34	0.38	0.003	1.54	1.45	0.56	0.000	0.17	0.16	0.01
MSITAL	0.001	0.23	0.29	0.01	0.005	1.95	1.37	0.90	0.003	1.71	1.59	0.69
MSJPAN	-0.002	-1.40	-1.60	0.47	0.000	-0.21	-0.20	0.01	-0.002	-1.30	-1.45	0.40
MSNWAY	0.001	0.30	0.24	0.02	0.001	0.18	0.17	0.01	0.001	0.38	0.33	0.04
MSUTDK	-0.003	-1.79	**-2.18**	0.76	0.004	1.75	**1.99**	0.73	0.000	-0.09	-0.09	0.00
MSUSAM	-0.001	-0.74	-0.89	0.13	0.001	0.92	0.90	0.20	0.000	0.11	0.11	0.00
MSWRLD	-0.001	-1.05	-1.26	0.26	0.002	1.23	1.07	0.36	0.000	0.10	0.10	0.00
Considering max/min returns using horizon k=36												
MSCNDA	0.001	0.27	0.30	0.02	0.003	1.27	1.23	0.38	0.002	1.06	1.09	0.27
MSFRNC	-0.001	-0.62	-0.75	0.09	0.008	**2.88**	**2.35**	1.95	0.003	1.37	1.26	0.45
MSGERM	-0.004	-1.56	**-2.08**	0.58	0.003	1.09	1.08	0.29	-0.001	-0.57	-0.63	0.08
MSITAL	0.002	0.57	0.71	0.08	0.013	**4.00**	**3.24**	3.70	0.007	**3.09**	**3.00**	2.23
MSJPAN	-0.003	-1.27	-1.38	0.38	0.003	0.96	1.05	0.22	-0.001	-0.41	-0.44	0.04
MSNWAY	-0.001	-0.49	-0.40	0.06	0.001	0.20	0.18	0.01	-0.001	-0.27	-0.23	0.02
MSUTDK	-0.004	-1.88	**-2.05**	0.84	0.004	1.32	1.78	0.42	-0.001	-0.67	-0.71	0.11
MSUSAM	-0.002	-0.93	-1.18	0.20	0.003	1.56	1.81	0.58	0.000	0.26	0.29	0.02
MSWRLD	-0.002	-1.12	-1.40	0.30	0.004	1.86	1.89	0.82	0.000	0.29	0.31	0.02

Table 8: Predictive regressions using oil price changes scaled with GARCH volatility

The table shows the beta coefficient, the t-statistic under strict OLS assumptions, the GMM-corrected t-statistic using Newey and West (1987) and the adjusted R^2. The left panel is the change in oil price scaled by the GARCH volatility as in Eq. (22); the middle panel considers positive changes only and the right panel considers negative changes only, as described in Eqs. (23) and (24).

	$\Delta o_t^{\ddagger}$				$\Delta o_t^{\ddagger,+}$				$\Delta o_t^{\ddagger,-}$			
Country	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Considering max/min returns using horizon k=12												
MSCNDA	-0.001	-0.67	-0.62	0.11	-0.001	-0.30	-0.23	0.02	-0.003	-0.88	-1.13	0.19
MSFRNC	-0.007	**-3.18**	**-2.84**	2.37	-0.009	**-2.69**	**-2.47**	1.70	-0.011	**-2.46**	**-2.45**	1.43
MSGERM	-0.009	**-4.38**	**-4.09**	4.39	-0.012	**-3.86**	**-3.59**	3.46	-0.013	**-3.13**	**-3.34**	2.30
MSITAL	-0.009	**-3.58**	**-3.21**	2.99	-0.008	**-2.05**	**-1.98**	0.99	-0.020	**-4.10**	**-3.72**	3.88
MSJPAN	-0.003	-1.33	-1.27	0.42	-0.005	-1.74	-1.70	0.72	-0.001	-0.22	-0.25	0.01
MSNWAY	-0.001	-0.39	-0.34	0.04	0.002	0.51	0.41	0.06	-0.008	-1.43	-1.44	0.49
MSUTDK	-0.007	**-3.15**	**-3.46**	2.33	-0.009	**-2.86**	**-3.07**	1.93	-0.009	**-2.18**	**-2.57**	1.12
MSUSAM	-0.006	**-3.55**	**-3.64**	2.93	-0.007	**-2.89**	**-2.67**	1.96	-0.009	**-2.87**	**-3.23**	1.94
MSWRLD	-0.006	**-3.69**	**-3.16**	3.17	-0.007	**-3.18**	**-2.47**	2.37	-0.008	**-2.76**	**-3.07**	1.79

Figure 1: The figure plots the nominal and real WTI oil price. The nominal price is the solid line and the real price is the dotted line. The vertical bars mark important events that impacted the oil price.

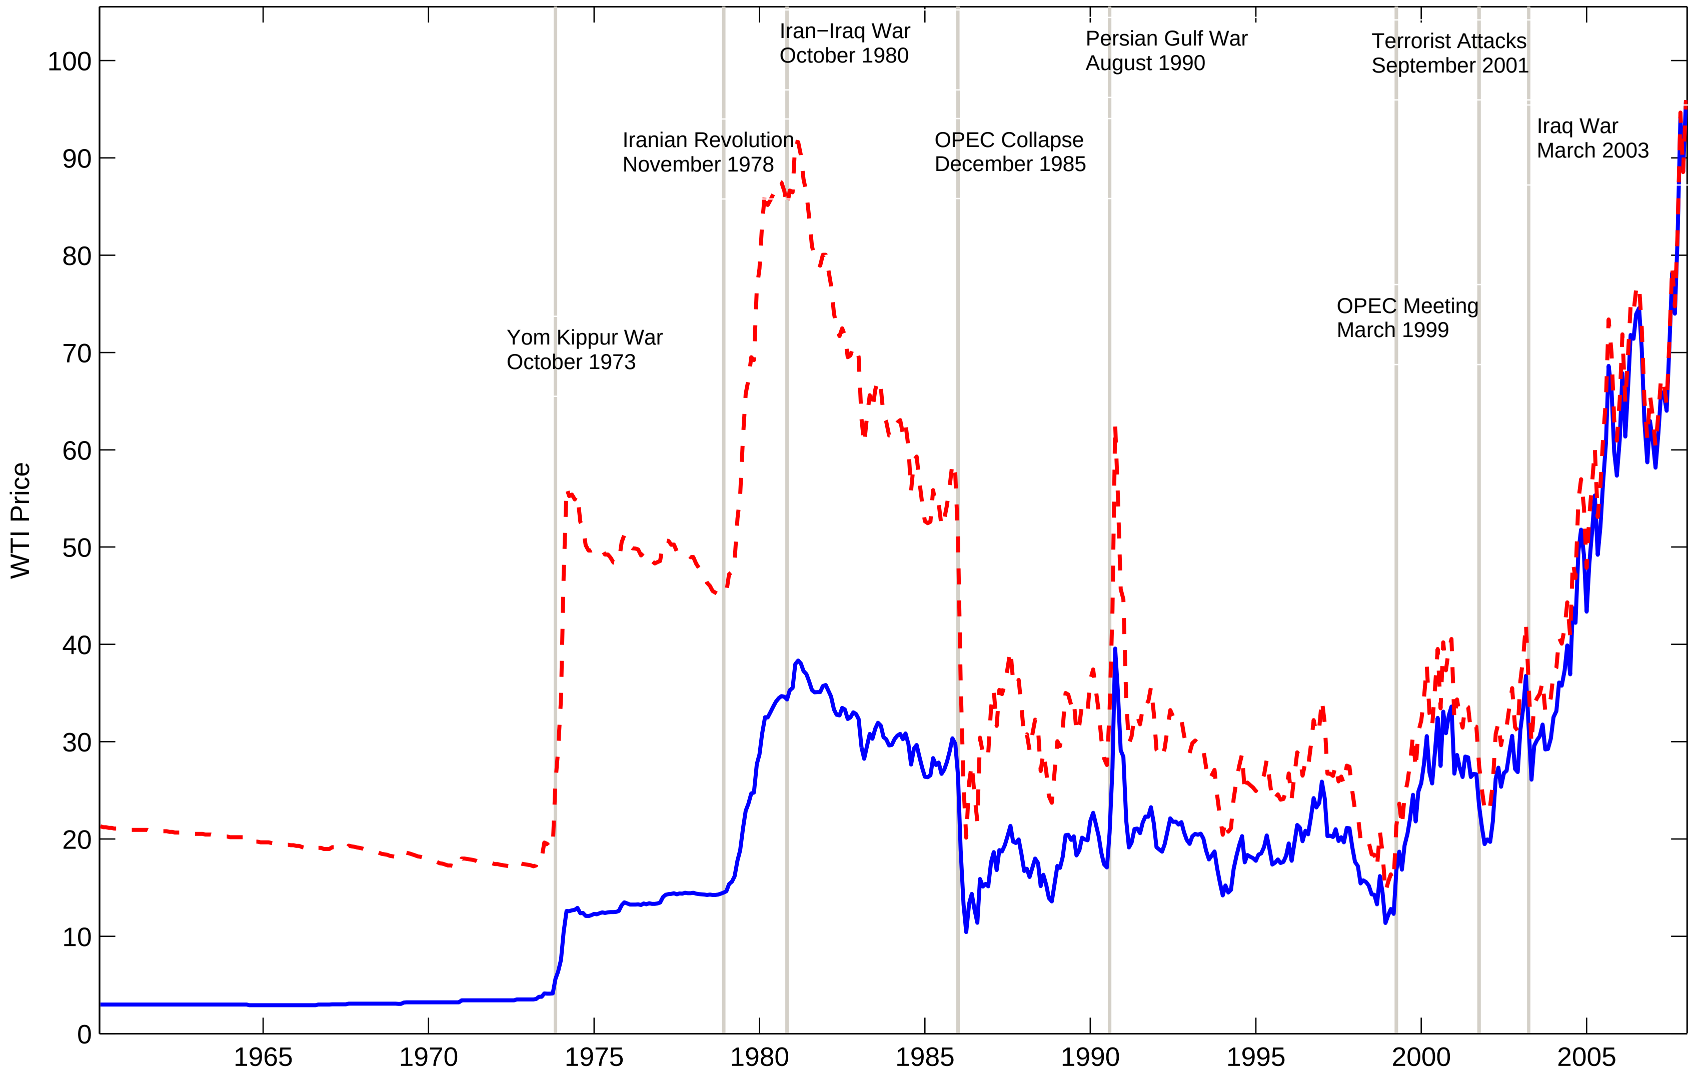

Figure 2: The figure plots the time-varying beta coefficient on lagged oil price changes in the predictive univariate regression. The upper panel shows important oil dates marked as grey bars whereas the lower panel shows NBER recession dates as grey bars.

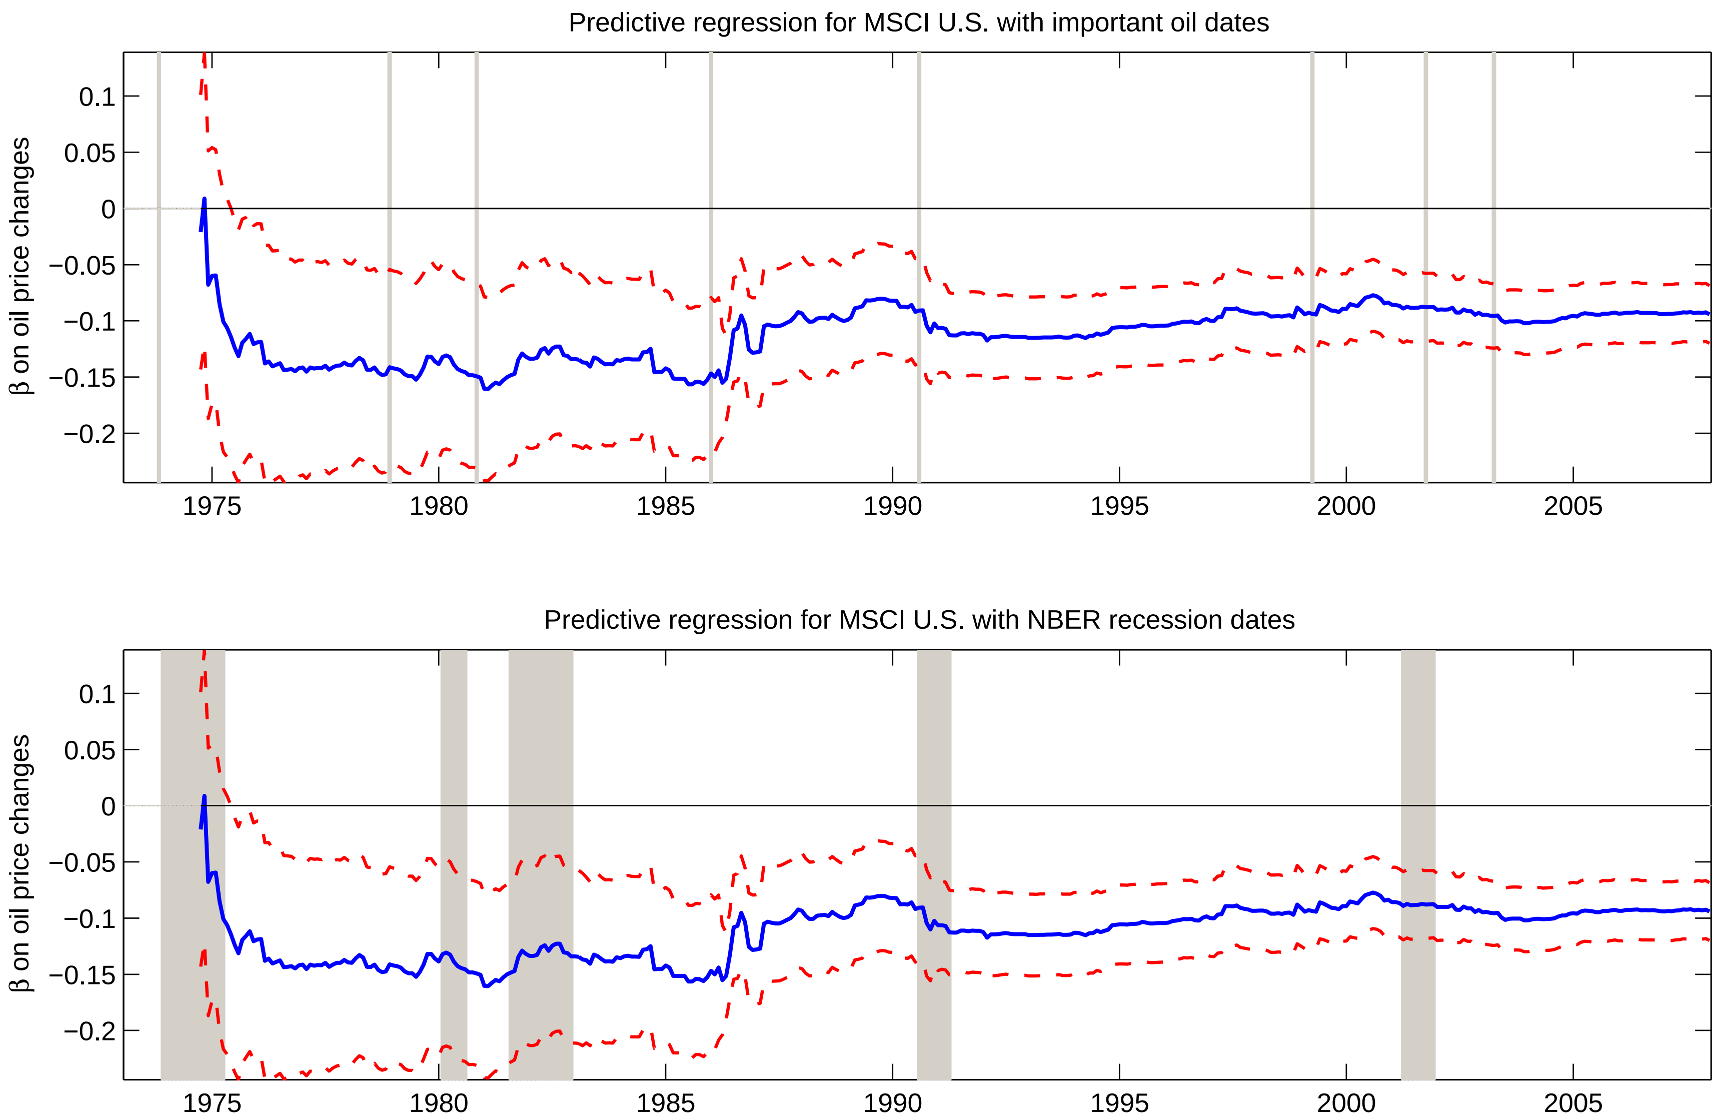

Figure 3: The figure plots the time-varying beta coefficients on lagged oil price changes in the predictive univariate regressions. The grey bars mark important oil dates.

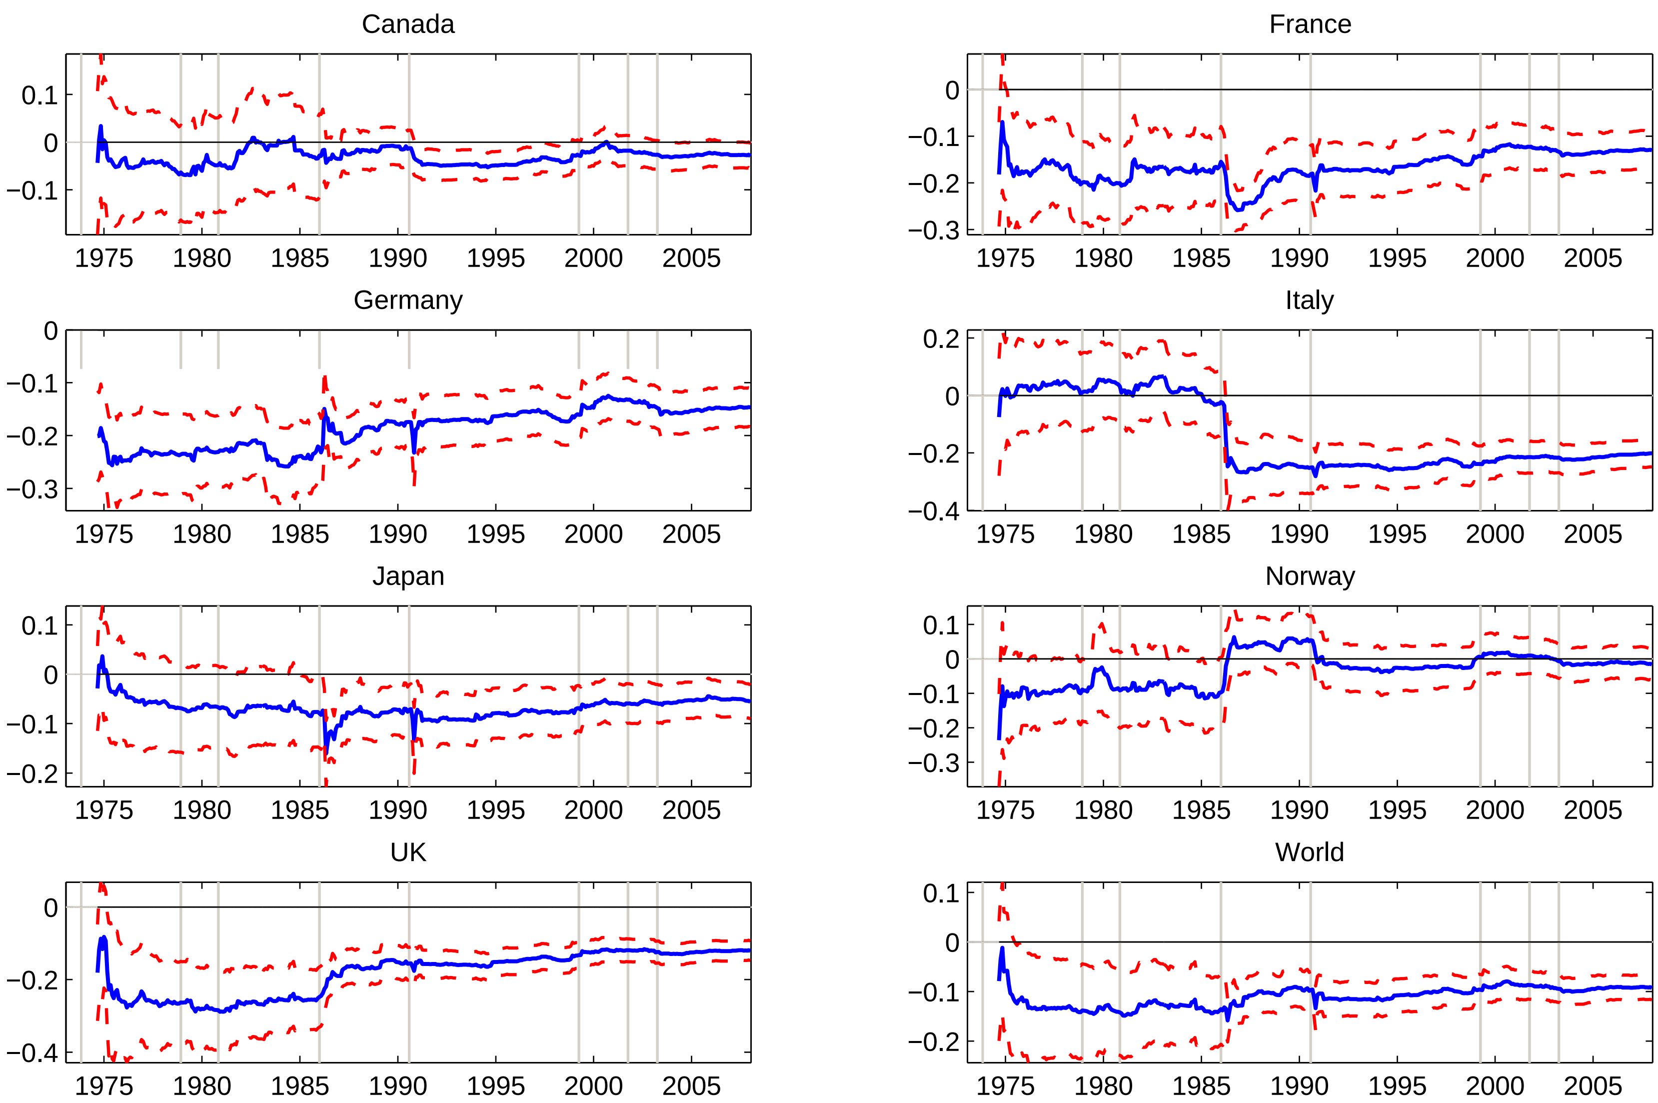

Figure 4: The figure plots the cumulative difference in the sum of squared errors between the prevailing mean and a predictive regression for the U.S. When the graph is increasing, the OLS model predicts better than a random walk.

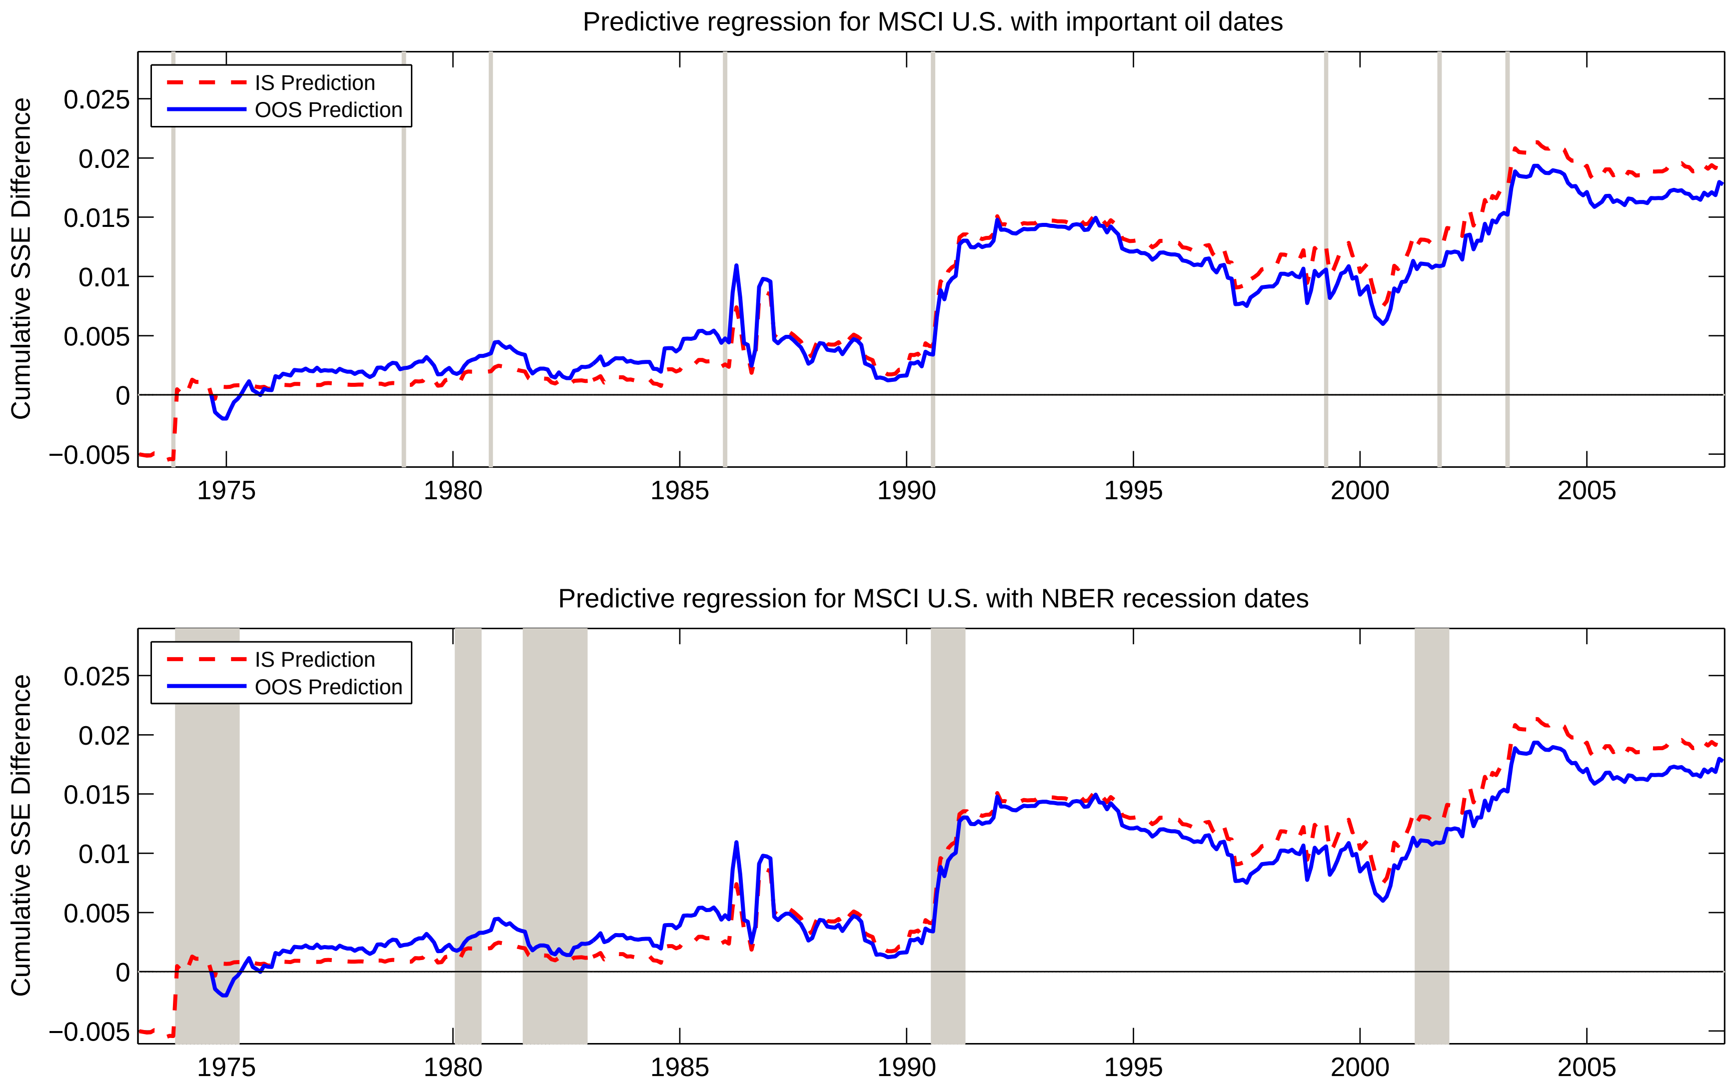

Figure 5: The figure plots the cumulative difference in the sum of squared errors between the prevailing mean and a predictive regression for the remaining countries in the sample. When the graph is increasing, the OLS model predicts better than a random walk.

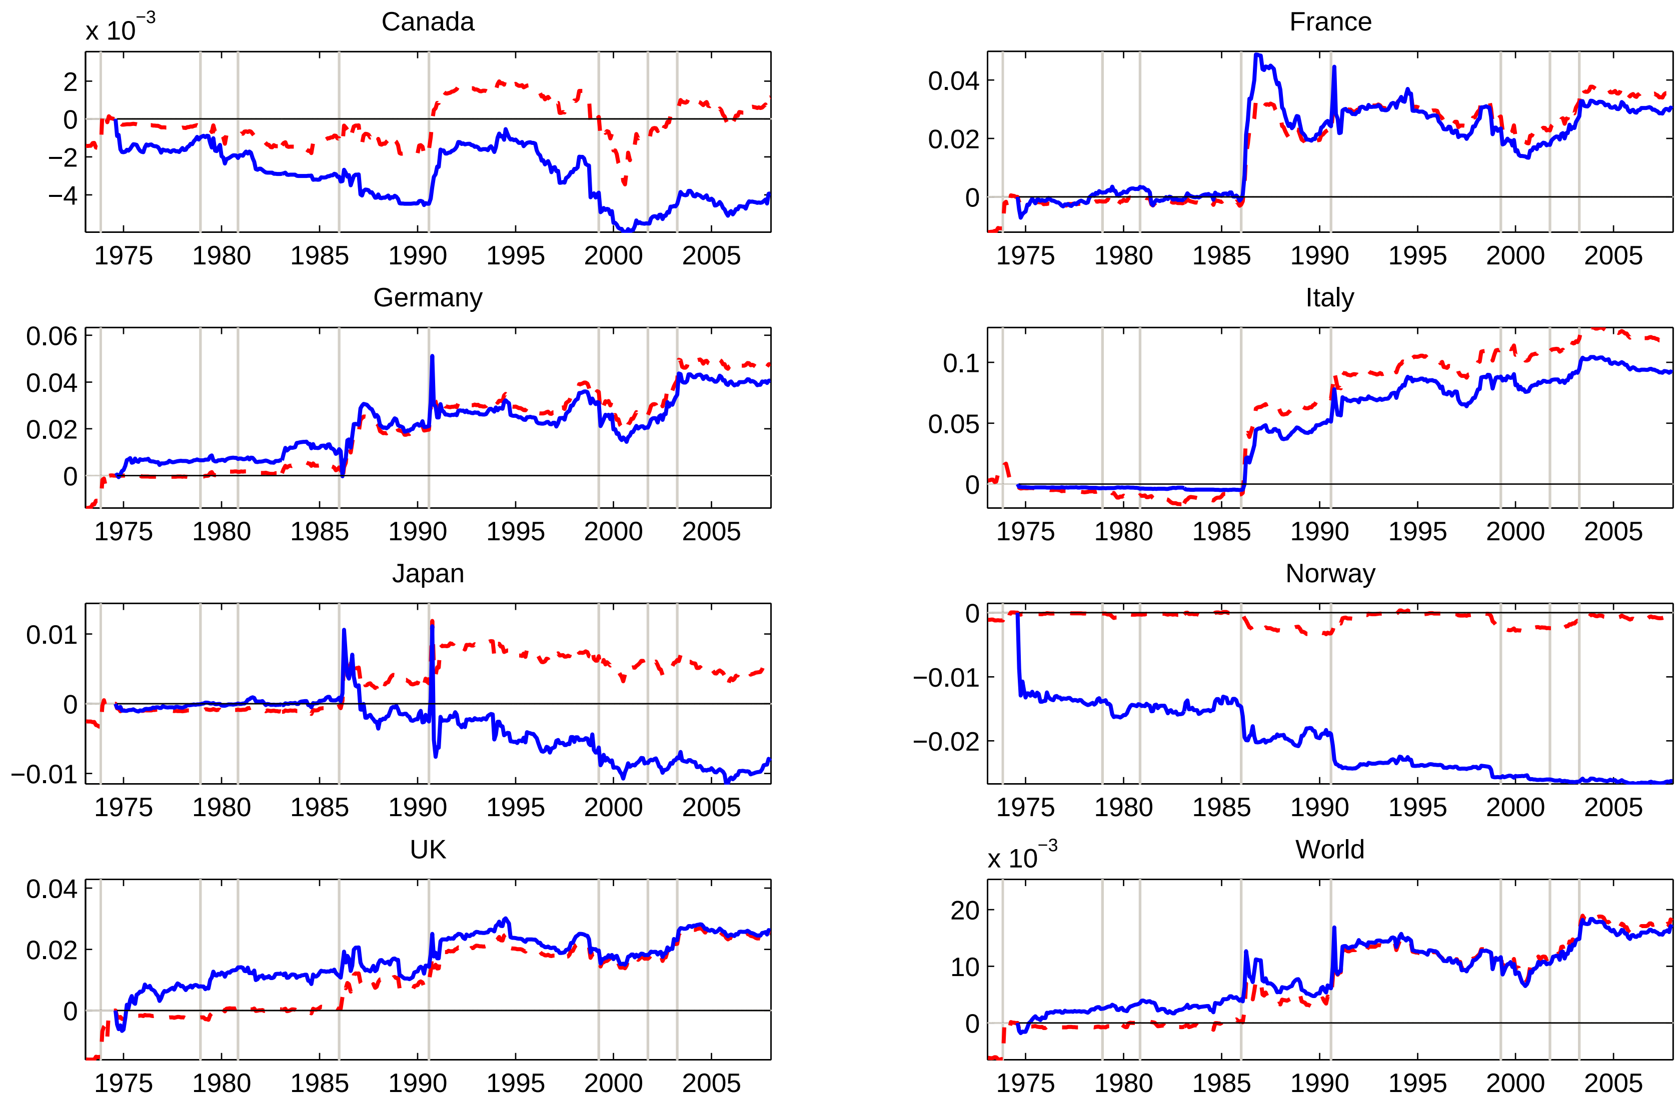

Figure 6: The figure plots the cumulative difference in the sum of squared errors between the prevailing mean and a predictive regression. In this figure the events identified in Section 2 have been removed from the sample. The x axis now corresponds to observation number after the removal. When the graph is increasing, the OLS model predicts better than a random walk.

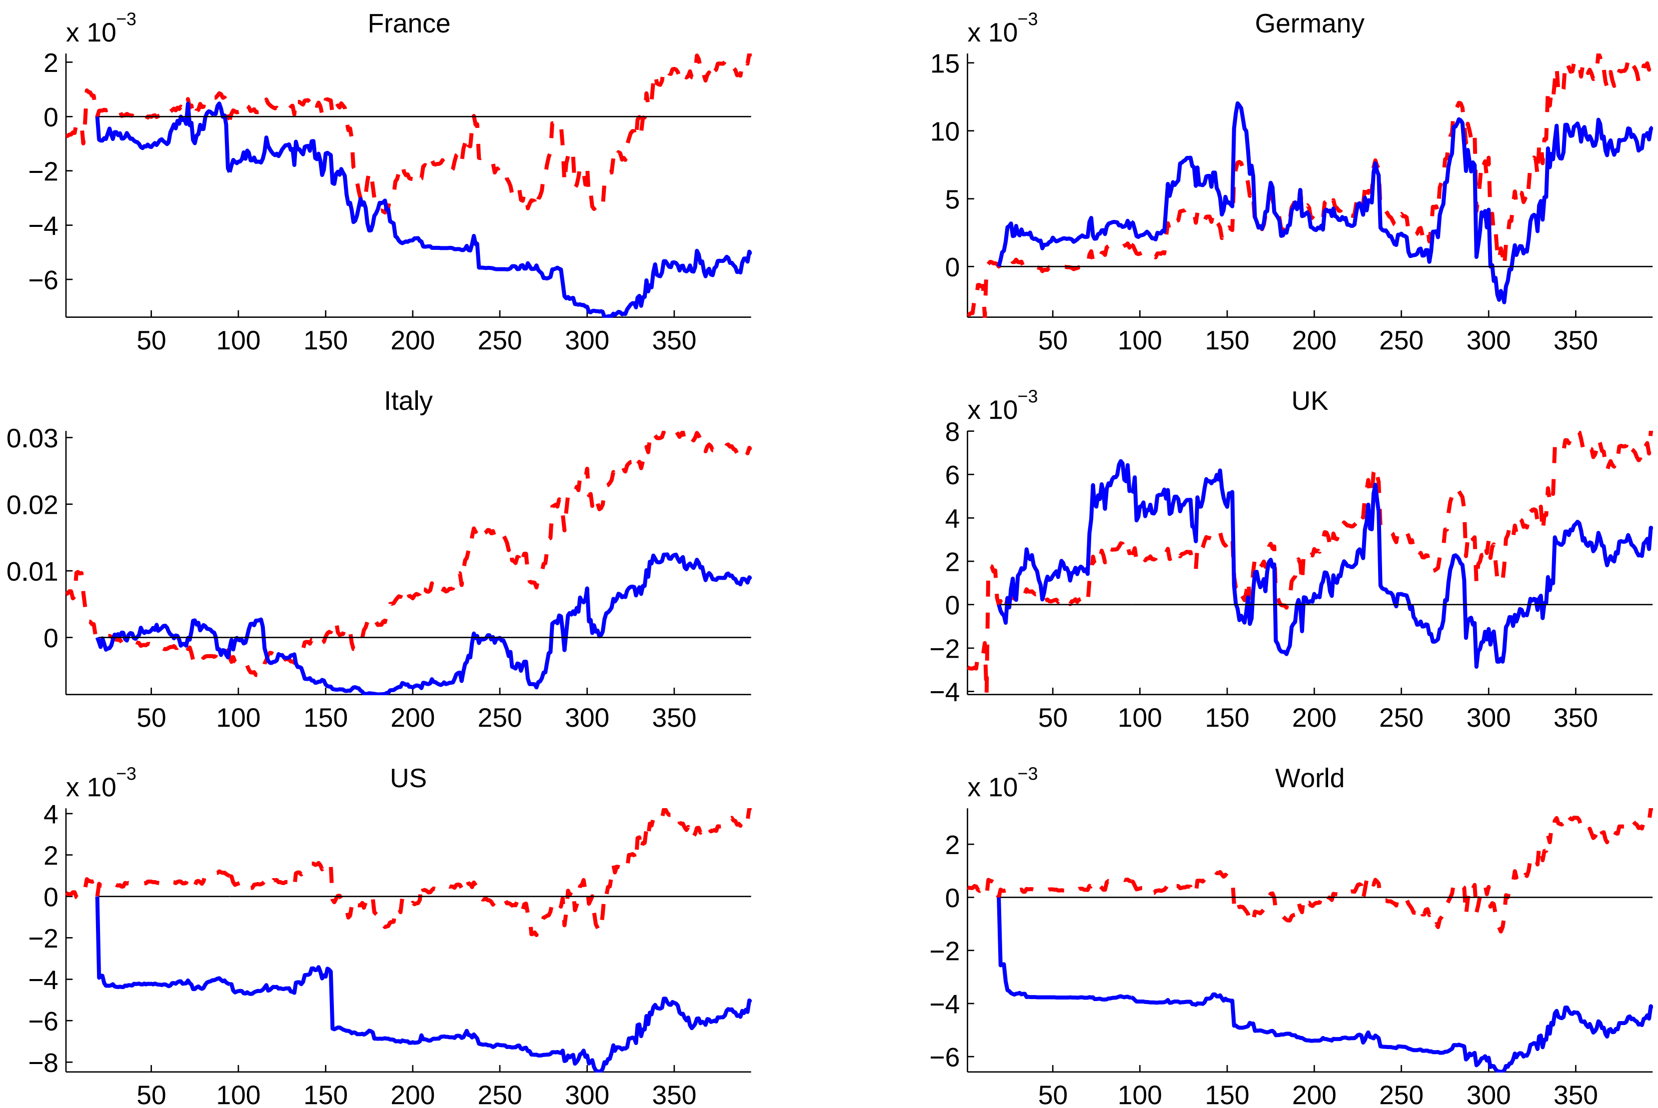

Figure 7: The figure plots the oil price measures corresponding to the max and min change in the oil price over a period of 12 and 36 months as described in Eqs. (12) and (14). The grey bars mark the important oil dates.