NHH

INSTITUTT FOR FORETAKSØKONOMI

DEPARTMENT OF FINANCE AND MANAGEMENT SCIENCE

FOR 13 2009
ISSN: 1500-4066
NOVEMBER 2009

Discussion paper

Oil Price Shocks and Stock Return Predictability

BY
LARS QVIGSTAD SØRENSEN

Norges Handelshøyskole
NORWEGIAN SCHOOL OF ECONOMICS AND BUSINESS ADMINISTRATION

Oil Price Shocks and Stock Return Predictability *

Lars Qvigstad Sørensen †

Norwegian School of Economics and Business Administration

First Draft: December 2008

This Draft: August 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 and its representation as an anomaly.

*Thanks to Jonas Andersson, Jørgen Haug, Halvor Hoddevik, Thore Johnsen, Ralph Koijen (discussant at EFA), Peter Nyberg, Richard Priestley, Per Östberg, and seminar participants at the Nordic Finance Network 2009 Workshop in Copenhagen, NHH Finance Seminar Series, and the FIBE 2009 Conference at NHH for helpful conversations and comments.

†E-mail: Lars.Sorensen@nhh.no. Address: NHH, Department of Finance and Management Science, Helleveien 30, 5045 Bergen, Norway.

1 Introduction

The predictability of asset returns is a controversial topic in financial economics. According to the Efficient Markets Hypothesis, all information in investors' information sets 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 investment to capital ratio (Cochrane, 1991), the consumption, wealth, income ratio (cay) (Lettau and Ludvigson, 2001), and the output gap (Cooper and Priestley, 2008). Traditionally, these variables have been thought to capture time-varying risk premia, 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), in which individual commodities are weighted by their respective world production quantities. Crude oil accounted for 40% of the entire index and Brent crude oil for another 15%, thus giving oil products a weight of around 55% of the entire commodity index. Given its large weight, it is not surprising that changes in oil prices may be an important factor behind fluctuations in the economy. In fact, some of the largest fluctuations in the U.S. economy have been preceded by political events in the Middle East that 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 U.S. 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. Given their importance for the macroeconomy, it is also plausible that oil price changes are important in explaining changes in stock prices, either through discount rate or cash flow effects. In two interesting recent articles, Pollet (2004) and Driesprong, Jacobsen, and Maat (2008) document that oil price changes predict stock returns. Driesprong et al. (2008) use linear predictive regressions to show that lagged oil price changes strongly predict returns, particularly in developed markets. They provide a number of arguments for why the predictability does not represent time-varying discount rates, and argue instead that it represents an underreaction to information, consistent with the story in Hong and Stein (1999). They conclude that the statistical predictability stemming from lagged oil price changes represents an anomaly that investors can profitably exploit.

The interpretation of the oil price predictability as an anomaly requires that it represents a stable relation between lagged oil price changes and subsequent stock returns that investors can identify ex ante, trade on and make risk-adjusted excess returns. Since the early 1970s the oil price has responded sharply to geopolitical events, particularly in the Middle East. I examine whether these events have contributed to the statistical predictability in a significant way and evaluate the predictability according to the criteria set forth in Goyal and Welch (2008). They write that the signal from a predictive regression could give an investor confidence in the signal if it showed (i) good in-sample (IS) and reasonable out-of-sample (OOS) performance over the sample period, (ii) an upward drift in both IS and OOS plots, and (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 as, historically, many of the most dramatic oil price movements have corresponded to 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 and times of crisis are the only periods with good predictability.

In order to assess the IS and OOS performance of lagged oil price changes as a predictor, I use the graphical diagnostic developed in Goyal and Welch (2003, 2008). The graphs

[1] The plots are explained in Section 4.

reveal that oil price changes seemingly predict well both IS and OOS.[2] However, a closer inspection of the graphs suggests that most of the predictability is driven by a few very significant events. This casts doubt on the oil strategy's virtue as an equity investment strategy. I construct a counterfactual scenario where I remove a few historic events that led to large oil price changes. This significantly weakens the predictability, particularly OOS. In subsequent sections I reinforce and formalize the graphical evidence by considering nonlinear transformations of the oil price changes. 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. These measures, constructed to pick out oil price changes that arise because of exogenous events, are henceforth labeled exogenous oil price changes. 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.

This paper illuminates the sources of the predictability of stock returns using lagged oil price changes as the predictor. Since most of the predictability is associated with wars and other extreme events, and since it disappears when the frequency is lowered to quarterly or annual, I argue that the predictability represents neither time-varying discount rates nor an anomaly. Rather, it is due to a few very significant extreme events that investors arguably could not have confidently profited from. A plausible counterfactual scenario removes much of the predictability, and non-linear transformations removes it completely, thus highlighting the importance of extreme events for the statistical forecastability. My conclusion is therefore different from Driesprong et al. (2008): I do not consider the statistical predictability as evidence of an anomaly, nor do I recommend a trading strategy based on the oil signal.

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, both IS and OOS. Section 5 performs regressions using

[2]This is in contrast to most of the predictors examined in Goyal and Welch (2008), which have performed poorly since the oil shocks in 1973.

nonlinear transformations of the oil price changes as the predictor. Section 6 concludes.

2 Data and descriptive statistics

2.1 The oil price

Throughout the analysis 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 obtained by deflating the nominal WTI crude price with the U.S. 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 when investors could reasonably a relation between oil prices, the real economy and the stock markets. The vertical 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, briefly explain below the most important of the episodes that are marked with a vertical 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

The main part of the analysis is carried out using the MSCI equity indexes from 1973 to 2007; all equity indexes are monthly total return indexes. 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. Therefore, it is 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 Datastream indexes are available from 1973, except for Norway, for which Datastream index returns are available from 1980.

2.3 Interest rates and CPI data

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. I use the 3-month bill rate for the U.S. to compute excess returns for the World market. Similarly, U.S. inflation is used to compute real World returns. Alternatively, bill rate and CPI baskets could have been constructed to obtain a World risk-free rate and inflation rate. However, the empirical analysis suggests that these assumptions are fairly innocuous as the results are similar for real and nominal, and actual and excess returns. Interest rates and CPI data are downloaded from Global Financial Data.

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. The real oil price is the most volatile of the time series with an annual volatility of 28.61% during the sample period. 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 tumultuous 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 for normality. 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

In this section I extend the results in Driesprong et al. (2008) by considering real actual returns and returns in excess of the risk-free rate, in addition to nominal actual returns. The sample period is also extended from ending in 2003 to ending in 2007. I start by estimating univariate predictive regressions for all G7 countries, Norway and the World market with a linear specification 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 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, which is 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 corrected for heteroskedasticity and autocorrelation (HAC) using the Newey and West (1987) procedure.[3] Following Newey and West I include autocovariances up to five lags in the monthly sample running from 1973:01 to 2007:12.[4] Table 2 reports the results for the univariate regressions in three cases. Panel A considers nominal and real log returns as measured in local currency. Panel B reports the results in local currency 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. Unreported results show that nominal increases in the oil price lead to statistically significant increases in inflation in all countries in the sample except Italy. This effect is not surprising and is partly mechanical through the energy component of the CPI. This helps explain why the negative relation between lagged oil price changes and stock returns is amplified when considering real rather than nominal variables. It is not clear whether the effect of oil price changes on international stock returns will be greater or smaller if these stock returns are hedged into U.S. dollars. To assess the currency effect, I compute actual real and nominal international stock returns in U.S. dollars according to the formula

$$r_{t+1}^{\text{usd}} = \ln\left[\left(1 + R_{t+1}^{\text{local}}\right) / \left(1 + R_{t+1}^{\text{fx}}\right)\right], \tag{2}$$

where r_{t+1}^{usd} is the real or nominal log return measured in U.S. dollars, R^{local} is the real or nominal return on the equity index measured in the local currency and R^{fx} is the return on the local currency versus U.S. dollars. Excess log returns are unaffected by the currency conversion as the currency returns cancel. There is a negative albeit statistically insignificant relation between oil price changes and exchange rate depreciation versus U.S. dollars for the countries in the sample, i.e. currencies have tended to weaken versus the dollar following an oil price increase. Panel C in Table 2 shows that the predictability is marginally weaker when both oil price changes and stock returns are measured in U.S. dollars. This is consistent with the observed exchange rate movements mentioned above.

[3] The White heteroskedasticity consistent estimator and the Hansen-Hodrick estimator both give similar results for the asymptotic covariance matrix of the coefficients.

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

In the remainder I focus on stock returns measured in local currency. 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. The results in Table 2 suggest that lagged oil price changes strongly predict returns in six out of the nine markets considered. However, one has to bear in mind that these nine markets 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.

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.[5] 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 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

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

theory seems to be that the ICB supersector Automobiles is strongly forecasted in all markets where it is available.[6] 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 appears to be no feedback from the stock markets to the oil price.

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

3.1 Robustness tests

The previous section shows that oil price changes lead stock returns and that this relation has been both statistically and economically significant. However, the importance of this finding would be diminished if the predictive ability of oil price changes were subsumed by other predictors used in the extant literature. This section assesses how robust the findings from the univariate regressions are by including several known predictors. The dividend-price ratio, measured as the moving sum of the last 12 months of dividends

[6] 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.

divided by price, has been shown to have predictive power for returns by Campbell and Shiller (1988) and Fama and French (1988) among others. The short rate has been shown to be a strong forecasting variable by for example Ang and Bekaert (2007). The term spread, measured as the difference between the 10-year government bond rate and the 3-month T-bill rate has been used to predict returns by e.g. Fama and French (1989). Inflation has been shown to predict returns by Fama and Schwert (1977). I also control for possible autocorrelation in returns by including the lagged equity premium. The results from these robustness tests are reported in Table 5. Interestingly, all countries that recorded a statistically significant coefficient on lagged oil price changes in the univariate regression, as reported in Table 2, continue to do so. The coefficient on oil price changes is significant in more countries than any of the other predictors that have been used in the extant literature. The statistical forecasting power of oil price changes is not subsumed by the other predictors.

It is also important to examine what happens to the predictability when changing the horizon over which stock returns and oil price changes are measured. There are several ways to change the frequency; I consider what happens both when I use overlapping and when I use non-overlapping returns. The non-overlapping results are straightforward. They are obtained by summing log returns and log oil price changes over a horizon, never using the same observation more than once, and then performing the regressions. Thus, the sample shrinks from T to T/h, where T is the number of months used in the estimation and h is the number of months over which returns are aggregated. The overlapping log excess returns from time $t-h$ to t are computed as

$$r^e_{t-h\to t} = \sum_{i=0}^{h} r^e_{t-i}. \tag{3}$$

This procedure induces serial correlation since the same lagged returns are included several times. To try to account for the serial correlation I correct the standard errors of the coefficient estimates using the Newey and West (1987) procedure with $2(h-1)$ lags. The vector of oil price changes is left unchanged so the regression results using overlapping observations are computed by estimating

$$r^e_{t-h\to t} = \alpha + \beta x_{t-h} + \varepsilon_{t+1}. \tag{4}$$

Table 7 reproduces Panel B of Table 2 but with quarterly and annual frequencies instead of monthly. The statistical significance disappears completely when considering non-overlapping returns but is mostly retained when considering overlapping returns, especially at the quarterly frequency. This could be explained by the historical tendency of a reversion following some of the largest oil price changes. Therefore, when using non-overlapping observations over a longer horizon the monthly movements in both oil prices and stock returns have canceled out. However, the overlapping procedure is constructed in such a way that these intermediate movements are not fully washed away. The evidence in Table 7 suggests that the predictability of equity premiums by oil price changes is mostly a short-horizon phenomenon. This is important evidence against the oil predictability describing time-varying expected returns, but is consistent with a slow reaction to information story, as suggested by Driesprong et al. (2008).

4 Out-of-sample evidence of predictability

The IS predictability in the previous section seems impressive. In this section I try to scrutinize the predictability results to find out when and why oil price changes predict returns. I examine both IS and OOS predictability using a newly developed 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.

$$E_t^{(pm)}\left[r_{t+1}^e\right] = \frac{1}{t}\sum_{i=1}^{t} r_i^e. \tag{5}$$

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

$$r_{t+1}^e = \alpha + \beta x_t + \varepsilon_{t+1}^{(x)}. \tag{6}$$

Superscript (x) means that a variable belongs to the OLS model and (pm) means that a variable belongs to the prevailing mean model. 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

$$E_t^{(x)}\left[r_{t+1}^e\right] = \hat{\alpha}_t + \hat{\beta}_t x_t. \tag{7}$$

This approach uses an "expanding window" where parameters are estimated recursively with data from the first data point in the sample up to the time of the forecast to generate one-step ahead forecasts. 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. The IS period before I start forecasting is thus 1973:03-1974:10 and I compute OOS forecasts from 1974:11 to 2007:12. 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

$$\mathrm{SE}^{(pm)}(t) = \left(r_t^e - E_{t-1}^{(pm)}\left[r_t^e\right]\right)^2 \tag{8}$$

$$\mathrm{SE}^{(x)}(t) = \left(r_t^e - E_{t-1}^{(x)}\left[r_t^e\right]\right)^2. \tag{9}$$

Subtracting Eq. (9) from Eq. (8) 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

$$\text{Net-SSE}(\tau) = \sum_{t=1}^{\tau} \mathrm{SE}^{(pm)}(t) - \mathrm{SE}^{(x)}(t), \tag{10}$$

where $\mathrm{SE}^{(pm)}$ is the squared OOS prediction error in year t when the prevailing mean is used as the forecast of the next return and $\text{SE}^{(x)}$ 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 IS 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 IS forecast errors are obtained using the entire sample under consideration.

The procedure described above entertains a different null hypothesis than the one used in Tables 2 and 3. In order to assess the OOS statistical significance of the alternative OLS model versus the prevailing mean, different tests must be performed. The MSE-F test developed by McCracken (2007) tests the null hypothesis that the prevailing mean model has a mean-squared error (MSE) that is less than or equal to the MSE of the OLS model. This statistic is given as

$$\text{MSE-F} = P \cdot \frac{\text{MSE}^{(pm)} - \text{MSE}^{(x)}}{\text{MSE}^{(x)}}, \tag{11}$$

where P is the number of one-step ahead forecasts and $\text{MSE}^{(pm)}$ and $\text{MSE}^{(x)}$ are the mean-squared errors from the prevailing mean and OLS model respectively.[7] I also report the Clark and McCracken (2001) encompassing statistic, where the null hypothesis is that the restricted model encompasses the unrestricted model and is given by

$$\text{ENC-NEW} = P \cdot \frac{\sum_{t=R}^{T} \left[\varepsilon_{t+1}^{(x)}\right]^2 - \varepsilon_{t+1}^{(x)} \varepsilon_{t+1}^{(pm)}}{\sum_{t=R}^{T} \left[\varepsilon_{t+1}^{(pm)}\right]^2}, \tag{12}$$

where R denotes the number of IS observations used before forecasts are made and $\varepsilon^{(x)}$ and $\varepsilon^{(pm)}$ are the residuals from the forecasting models. Goodness of fit is measured by IS and OOS R^2 estimates (not adjusted for degrees of freedom here). This statistic is computed as

$$R^2 = 1 - \left(1 - \frac{\text{MSE}^{(x)}}{\text{MSE}^{(pm)}}\right). \tag{13}$$

[7]All OOS statistics described in this section are estimated with a recursive scheme that uses more data as the forecasting moves forward in time.

Table 6 reports the OOS statistics discussed above. The first column displays the IS R^2 for comparison; these are identical to the ones in Table 2 Panel B using real returns.

Figure 4 graphs the IS and OOS performance of the predictive regression 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 dotted (IS) line starts before the solid (OOS) 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 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.[8]

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. The difference between the IS and OOS results is 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.

[8]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"*.

4.1 Predictability under a counterfactual scenario

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

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

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 $\mathbf{Z}$. Removing the oil dates results in a significant drop in the predictability. Only three out of nine equity premia, Germany, Italy and the United Kingdom, remain significantly predictable from oil price changes IS after removing these dates (the table has been omitted for brevity). This illustrates 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

This section supports and formalizes the graphical evidence in Section 4, that predictability is largely driven by a few significant events. There exists a large literature that examines 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. The positive effects on the economy following the oil price decrease after the OPEC collapse in 1985 and 1986 were not of the same magnitude as the adverse effects following the large oil price increases in the 1970s. This led Mork (1989) to consider a nonlinear specification in which only oil price increases impact the economy. Elaborating on these ideas, Hamilton (2003) further examines the functional form for oil and the macroeconomy. However, no consensus has been reached regarding the effects oil price changes have on the stock markets. The analysis in Section 3 was conducted with a linear specification, as 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). The examination of functional forms is another means of identifying and isolating the events that cause the statistical predictability. If the true functional form is nonlinear, then inference based on a linear regression can be biased. 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

$$r_{t+1} = f\left(\Delta o_t\right) + \varepsilon_{t+1}, \tag{14}$$

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

$$r_{t+1} = \alpha + \boldsymbol{\beta}'\tilde{\mathbf{x}}_t + \varepsilon_{t+1}, \tag{15}$$

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 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

$$\Delta o_t^{+} = \begin{cases} \Delta o_t & , \quad \Delta o_t > 0 \\ 0 & , \quad \Delta o_t \leq 0 \end{cases} . \tag{16}$$

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. Therefore, a priori, it may be just as likely that the markets react positively to an oil price decrease. I therefore also consider the negative equivalent

$$\Delta o_t^{-} = \begin{cases} \Delta o_t & , \quad \Delta o_t < 0 \\ 0 & , \quad \Delta o_t \geq 0 \end{cases} . \tag{17}$$

Table 8 shows that 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.

5.2 Extreme changes in the oil price over a horizon

The oil market since 1973 has been characterized by extreme price fluctuations corresponding to important geopolitical events, as illustrated in Fig. 1. I therefore examine whether these extreme events have had a particular influence on the predictability of stock returns using lagged oil price changes. Hamilton (2003) proposes a measure which is targeted at gauging the effects of these exogenous oil price shocks on GDP. He compares today's oil price against the maximum oil price attained over a prespecified horizon. 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 occurred. 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, often referred to as the net oil price increase, is

$$\Delta o_t^{\dagger,+} = \begin{cases} \ln\left(o_t/o_{t-1}^{\max}\right) & , \quad o_t > o_{t-1}^{\max} \\ 0 & , \quad o_t \leq o_{t-1}^{\max} \end{cases}, \tag{18}$$

where the max oil price is defined as

$$o_{t-1}^{\max} = \max\left\{o_{t-1}, \ldots, o_{t-k}\right\}, \tag{19}$$

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

$$\Delta o_t^{\dagger,-} = \begin{cases} \ln\left(o_t/o_{t-1}^{\min}\right) & , \quad o_t < o_{t-1}^{\min} \\ 0 & , \quad o_t \geq o_{t-1}^{\min} \end{cases}, \tag{20}$$

where $o_{t-1}^{\min}$ is defined as in Eq. (19) with the min operator instead of the max. Figure 7 plots the oil price measures described in Eqs. (18) and (20). The figure shows that the minimum and maximum transformations to a large extent pick out the oil price changes that are caused by exogenous events and set the oil price changes not caused by these events to zero. I also define a third measure that considers extreme oil price changes, whether they be positive or negative

$$\Delta o_t^{\dagger,\pm} = \begin{cases} \ln\left(o_t/o_{t-1}^{\max}\right) & , \quad o_t > o_{t-1}^{\max} \\ 0 & , \quad o_{t-1}^{\min} \leq o_t \leq o_{t-1}^{\max} \\ \ln\left(o_t/o_{t-1}^{\min}\right) & , \quad o_t < o_{t-1}^{\min} \end{cases}. \tag{21}$$

The results from regressions using the three measures above as the explanatory variables are reported in Table 9. 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 8, the predictability from positive oil price changes are almost non-existent. Only two out of nine countries are predicted by $\Delta o_t^{\dagger,+}$ when $k = 12$ and three markets 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 $\Delta o_t^{\dagger,-}$. The forecasting ability of $\Delta o_t^{\dagger,\pm}$ 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.

I now define measures that truncate the extreme oil price changes, as captured by the max oil price increases and decreases, to zero. I define these as the log change in the real oil price when the exogenous measures $\Delta o_t^{\dagger,+}$, $\Delta o_t^{\dagger,-}$ and $\Delta o_t^{\dagger,\pm}$ above are zero. The one corresponding to the maximum is thus

$$\Delta\tilde{o}_t^{\dagger,+} = \begin{cases} 0 & , \quad o_t > o_{t-1}^{\max} \\ \Delta o_t & , \quad o_t \leq o_{t-1}^{\max} \end{cases} , \tag{22}$$

the one corresponding to the minimum is

$$\Delta\tilde{o}_t^{\dagger,-} = \begin{cases} 0 & , \quad o_t < o_{t-1}^{\min} \\ \Delta o_t & , \quad o_t \geq o_{t-1}^{\min} \end{cases} , \tag{23}$$

and the one corresponding to both extremes is

$$\Delta\tilde{o}_t^{\dagger,\pm} = \begin{cases} 0 & , \quad o_t > o_{t-1}^{\max} \\ \Delta o_t & , \quad o_{t-1}^{\min} \leq o_t \leq o_{t-1}^{\max} \\ 0 & , \quad o_t < o_{t-1}^{\min} \end{cases} . \tag{24}$$

The results from estimating the predictive regressions using these measures are reported in Table 10. 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 $\Delta\tilde{o}_t^{\dagger,\pm}$ remove the most extreme changes that are due to geopolitical influences. Together, Tables 9 and 10 then establish a very interesting fact: the asset return predictability comes from the exogenous changes in oil prices. Once these have been truncated to zero, the remaining 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.

Driesprong et al. (2008) report that their results are robust to exclusion of the most extreme five percent of oil price changes. Unreported results, obtained by trimming the data at different percentiles, confirm their finding for the extended sample period ending in 2007. Compared to Table 2, the coefficient point estimates are similar, but less precisely estimated, so the t-statistics are somewhat lower. However, Table 10 shows that truncating the extreme oil price changes to zero according to Eq. (24) removes the predictability in all markets when the horizon is set to 12 months, and only the Italian equity premium has a significant t-statistic when the horizon is 36 months. Hence, trimming the data and truncating extreme values has different effects. The extensions of Hamilton's net oil price measure represent quantitative dummy variables corresponding to rare shocks. They capture the empirical observation that oil price shocks that occur in an otherwise calm environment have a stronger impact on stock returns than those that occur in a volatile environment.

5.3 Oil price changes scaled by GARCH volatility

Another transformation 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. Lee et al. (1995) 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

$$\Delta o_t = \phi_0 + \sum_{i=1}^{12} \phi_i \Delta o_{t-i} + a_t, \tag{25}$$

and conditional variance equation

$$a_t = \sigma_t \varepsilon_t, \qquad \varepsilon_t \sim NID\,(0,1) \tag{26}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 a_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \tag{27}$$

The volatility-adjusted oil price increase is then given by[9]

$$\Delta o_t^{\ddagger} = \Delta o_t / \sigma_t. \tag{28}$$

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

$$\Delta o_t^{\ddagger,+} = \max\left\{\Delta o_t^{\ddagger}, 0\right\} \tag{29}$$

$$\Delta o_t^{\ddagger,-} = \min\left\{\Delta o_t^{\ddagger}, 0\right\}. \tag{30}$$

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

Table 11 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.

5.4 Tests for asymmetry

The results presented thus far suggest that the functional form for the forecasting regression is both asymmetric and non-linear. However, Tables 8 to 11 do not provide formal statistical tests of non-linearity and asymmetry of the slope coefficients. One way to

[9] 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.

test for the presence of non-linearity is to estimate a regression equation which includes linear symmetric oil price changes and a censored regressor as an additional explanatory variable. The regression equation is then

$$r_{t+1} = \alpha + \beta_1 x_t + \beta_2 \tilde{x}_t + \varepsilon_{t+1}, \tag{31}$$

where x denotes the log change in the oil price and $\tilde{x}$ denotes the censored regressor. Panel A of Table 12 reports the p-values corresponding to the Wald tests of the null hypothesis that β_2 is zero, where $\tilde{x}$ is given as the extreme measures as outlined in Eqs. (18), (20) and (21), and the scaled measures as in Eqs. (28) to (30) respectively. The table shows that the presence of non-linearity is best captured by the extreme oil price decreases given by Eq. (20). As discussed above, this surprising result is in contrast to the effect on macroeconomic variables. There appears to be strong non-linear effects in France, Italy, Japan and for the World market. In contrast, the scaled oil price measures are mostly insignificant when included as additional regressors.

The test for equality of the parameter estimates is conducted by running the regression

$$r_{t+1} = \alpha + \beta_1 \tilde{x}_t^+ + \beta_2 \tilde{x}_t^- + \varepsilon_{t+1}, \tag{32}$$

where the explanatory variables are the censored positive and negative values of the linear, extreme and scaled measures. I have computed the robust Wald statistics using GMM estimators corrected with the Newey and West (1987) procedure. In the current setting, this is equivalent to testing the significance of the t-statistic on the censored explanatory variable. Panel B shows that the asymmetric effect is particularly strong in Italy; asymmetric effects measured as measured with Eq. (32) are also present for Canada, France and the World market.

6 Conclusion

The extant literature on the predictability of the equity premium usually explains statistical forecastability either as time-varying discount rates or as an anomaly. Driesprong et al. (2008) argue that the predictability stemming from lagged oil price changes represents an anomaly: *"the predictability of stock market returns based on oil price changes* *does qualify as truly anomalous as it cannot be attributed to time-varying risk premia"*. They conclude that the results point in the direction of a true market inefficiency. This paper argues that the statistical forecastability of international equity premiums represents neither time-varying discount rates nor an anomaly. 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 predominantly 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. A 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.

This 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 plummeting oil prices were followed by abysmal stock returns, thus disobeying the estimated negative relation between the two quantities. This most recent OOS evidence again illustrates the risks involved in following an oil-based investment strategy, and the fragility of the oil and stock returns relation.

I show that there is a strong statistical predictability arising mainly from a few important events, and that this statistical forecastability is not subsumed by other known predictors. Other rational predictors are unlikely to explain returns arising from these extreme events, and will consequently perform poorly as predictors in those periods. As such, accounting for the oil shocks by controlling for the major oil price movements may prove valuable when assessing the performance of various predictors in future research.

References

Ang, A. and G. Bekaert (2007): "Stock Return Predictability: Is it There?" *Review of Financial Studies*, 20, 651–707.

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

Barsky, R. B. and L. Kilian (2004): "Oil and the Macroeconomy since the 1970s," *Journal of Economic Perspectives*, 18, 115–134.

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

Clark, T. E. and M. W. McCracken (2001): "Tests of Equal Forecast Accuracy and Encompassing for Nested Models," *Journal of Econometrics*, 105, 85–110.

Cochrane, J. H. (1991): "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations," *Journal of Finance*, 46, 209–238.

Cooper, I. and R. Priestley (2008): "Time-Varying Risk Premiums and the Output Gap," *Review of Financial Studies,* forthcoming.

Driesprong, G., B. Jacobsen, and B. Maat (2008): "Striking Oil: Another Puzzle?" *Jounal of Financial Economics*, 89, 307–327.

Fama, E. F. and K. R. French (1988): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.

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

Fama, E. F. and G. W. Schwert (1977): "Asset Returns and Inflation," *Journal of Financial Economics*, 5, 115–146.

Goyal, A. and I. Welch (2003): "Predicting the Equity Premium with Dividend Ratios," *Management Science*, 49, 639–654.

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

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

——— (2008): "Oil and the Macroeconomy," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan, second ed.

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

Hong, H., W. Torous, and R. Valkanov (2007): "Do industries lead stock markets?" *Journal of Financial Economics*, 83, 367–396.

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

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

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

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

Lettau, M. and S. Ludvigson (2001): "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance, forthcoming*, 56, 815–849.

Mankiw, N. G. (2003): *Macroeconomics*, Worth Publishers.

McCracken, M. W. (2007): "Asymptotics for Out of Sample Tests of Granger Causality," *Journal of Econometrics*, 140, 719–752.

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.

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

Pollet, J. M. (2004): "Predicting Asset Returns With Expected Oil Price Changes," Harvard University Working Paper.

Stambaugh, R. F. (1999): "Predictive Regressions," *Journal of Financial Economics*, 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. Column three is the log 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. The covariances are quoted in basis points. The sample period is 1973 to 2007.

	Mean	Mean Excess	St. Dev.	Median	Min	Max	Skew	Kurt	AC(1)	AC(2)	JB
Oil	4.89	3.67	28.61	-0.32	-34.98	36.31	0.20	6.89	0.17	-0.05	267
Canada	5.52	3.02	17.15	8.24	-24.67	15.03	-0.77	6.05	0.04	-0.05	205
France	6.80	4.42	20.63	14.66	-24.86	19.93	-0.41	4.37	0.09	-0.02	45
Germany	6.54	4.44	19.83	10.13	-28.58	19.30	-0.86	6.33	0.05	0.02	246
Italy	3.75	1.77	23.75	4.30	-20.99	23.72	0.10	3.84	0.06	0.01	13
Japan	2.42	2.10	17.97	5.87	-22.60	18.38	-0.26	4.35	0.07	0.05	37
Norway	6.89	4.40	25.34	15.72	-35.33	22.52	-0.63	4.99	0.14	-0.06	97
U.K.	5.87	3.62	19.64	11.33	-30.50	40.83	0.10	11.47	0.09	-0.09	1,255
U.S.	5.55	4.33	15.47	8.74	-24.12	15.39	-0.57	5.48	0.02	-0.02	131
World	4.61	3.39	13.92	8.62	-22.04	12.79	-0.90	6.08	0.09	0.01	223

Stock return correlations and covariances

	Canada	France	Germany	Italy	Japan	Norway	U.K.	U.S.	World
Canada	**24.50**	**16.00**	**12.62**	**12.67**	**9.21**	**19.70**	**16.11**	**16.43**	**15.11**
France	0.54	**35.46**	**22.58**	**22.03**	**11.62**	**22.99**	**19.60**	**15.51**	**16.65**
Germany	0.45	0.66	**32.76**	**18.91**	**10.49**	**19.35**	**15.82**	**13.84**	**15.20**
Italy	0.37	0.54	0.48	**47.01**	**12.76**	**18.00**	**16.61**	**10.95**	**14.09**
Japan	0.36	0.38	0.35	0.36	**27.00**	**9.86**	**9.96**	**8.64**	**13.40**
Norway	0.54	0.53	0.46	0.36	0.26	**53.50**	**19.48**	**17.09**	**16.93**
U.K.	0.57	0.58	0.49	0.43	0.34	0.47	**32.14**	**16.00**	**16.47**
U.S.	0.74	0.58	0.54	0.36	0.37	0.52	0.63	**19.95**	**16.33**
World	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 between lagged log oil price changes and stock returns, the corresponding t-statistic under strict OLS assumptions, the GMM-corrected t-statistic using Newey and West (1987) and the R^2. Panel A shows the results for the actual log nominal and real returns. Panel B 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 Panel B is that the left panel uses nominal oil returns whereas the right panel uses real oil returns. Panel C shows the results when the international stock returns have been hedged into U.S. dollars. The sample period is 1973 to 2007.

	Nominal				Real			
	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Panel A: Log returns in local currency								
Canada	-0.03	-1.19	-1.33	0.34	-0.04	-1.37	-1.51	0.45
France	-0.13	**-3.82**	**-3.31**	3.39	-0.14	**-4.08**	**-3.49**	3.85
Germany	-0.15	**-4.52**	**-4.34**	4.68	-0.16	**-4.75**	**-4.48**	5.15
Italy	-0.21	**-5.29**	**-4.39**	6.31	-0.21	**-5.38**	**-4.54**	6.51
Japan	-0.06	-1.92	-1.84	0.88	-0.07	**-2.28**	**-2.13**	1.24
Norway	-0.02	-0.49	-0.44	0.06	-0.02	-0.57	-0.50	0.08
U.K.	-0.12	**-3.65**	**-4.69**	3.11	-0.13	**-3.97**	**-4.80**	3.65
U.S.	-0.10	**-3.73**	**-4.27**	3.24	-0.11	**-4.22**	**-4.77**	4.11
World	-0.09	**-4.12**	**-4.40**	3.91	-0.11	**-4.64**	**-4.87**	4.92
Panel B: Log excess returns in local currency								
Canada	-0.03	-1.14	-1.27	0.31	-0.03	-1.03	-1.13	0.25
France	-0.13	**-3.79**	**-3.27**	3.34	-0.13	**-3.72**	**-3.18**	3.22
Germany	-0.15	**-4.51**	**-4.30**	4.66	-0.15	**-4.43**	**-4.20**	4.50
Italy	-0.20	**-5.17**	**-4.31**	6.04	-0.20	**-5.09**	**-4.23**	5.86
Japan	-0.06	-1.92	-1.85	0.88	-0.06	-1.82	-1.74	0.79
Norway	-0.02	-0.42	-0.38	0.04	-0.02	-0.36	-0.33	0.03
U.K.	-0.12	**-3.67**	**-4.64**	3.14	-0.12	**-3.61**	**-4.54**	3.04
U.S.	-0.10	**-3.74**	**-4.26**	3.26	-0.09	**-3.65**	**-4.09**	3.11
World	-0.09	**-4.13**	**-4.39**	3.93	-0.09	**-4.01**	**-4.20**	3.72
Panel C: Log returns in USD								
Canada	-0.03	-0.82	-0.88	0.16	-0.03	-0.96	-1.03	0.22
France	-0.12	**-3.25**	**-2.40**	2.48	-0.13	**-3.46**	**-2.54**	2.80
Germany	-0.14	**-4.07**	**-3.41**	3.84	-0.15	**-4.27**	**-3.52**	4.20
Italy	-0.20	**-4.84**	**-3.41**	5.34	-0.21	**-4.88**	**-3.47**	5.42
Japan	-0.03	-0.93	-0.70	0.21	-0.05	-1.21	-0.90	0.35
Norway	-0.02	-0.34	-0.36	0.03	-0.02	-0.42	-0.43	0.04
U.K.	-0.11	**-2.88**	**-3.28**	1.95	-0.11	**-3.15**	**-3.52**	2.33
U.S.	-0.10	**-3.73**	**-4.27**	3.24	-0.11	**-4.22**	**-4.77**	4.11
World	-0.09	**-3.65**	**-3.54**	3.10	-0.10	**-4.14**	**-3.95**	3.95

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. The sample period is 1973 to 2007, except for Norway where the sample starts in 1980.

	Canada	France	Germany	Italy	Japan	Norway	U.K.	U.S.	World	Sign.	Tot.
					Returns regressed on lagged change in oil prices						
Total Market	-1.93	**-3.22**	**-4.39**	**-4.26**	-1.79	-0.39	**-4.60**	**-4.17**	**-3.29**	6	9
Oil and Gas	0.03	-1.24	-	-	-0.31	0.75	**-2.71**	-1.18	-1.39	1	7
Basic Materials	**-2.10**	**-2.79**	**-3.22**	**-3.93**	-1.18	**-2.20**	**-2.72**	**-2.20**	-1.90	7	9
Industrials	**-2.07**	**-3.01**	**-3.31**	**-3.39**	**-2.42**	-0.89	**-3.65**	**-3.44**	**-3.06**	8	9
Consumer Goods	-	**-2.87**	**-3.62**	**-4.38**	**-2.48**	-	**-3.24**	**-3.66**	**-2.81**	7	7
Health Care	**-2.52**	**-2.15**	**-3.52**	-	-0.82	-	**-3.30**	**-2.86**	**-2.65**	6	7
Consumer Services	**-3.57**	**-3.42**	**-3.45**	**-2.78**	**-2.08**	-0.44	**-4.43**	**-4.28**	**-3.83**	8	9
Telecommunications	-	-	**-3.71**	**-5.54**	0.00	-	-	**-2.49**	**-2.17**	4	5
Utilities	-0.75	-	**-3.73**	**-3.36**	-0.97	-0.53	-	-0.91	-1.41	2	7
Financials	-1.69	**-3.19**	**-3.59**	**-3.70**	-0.42	0.31	**-3.85**	**-2.93**	**-2.29**	6	9
Technology	-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						
Total Market	1.46	-1.18	-1.03	-0.29	-0.05	**3.07**	-1.29	-1.71	-0.75	1	9
Oil and Gas	**4.54**	**4.09**	-	-	0.44	**6.76**	**3.43**	**4.00**	**3.71**	6	7
Basic Materials	-0.16	-0.97	-1.23	0.12	-0.32	1.18	-0.49	-1.59	-0.49	0	9
Industrials	-0.13	-1.12	-0.74	-0.25	0.63	-0.13	-0.87	-1.70	-0.48	0	9
Consumer Goods	-	-1.95	-1.30	-1.23	0.60	-	-1.00	**-3.04**	-0.81	1	7
Health Care	-0.68	**-3.15**	-1.61	-	-0.60	-	**-3.94**	**-3.92**	**-3.11**	4	7
Consumer Services	-1.57	-1.58	**-1.96**	0.34	-0.63	0.39	-1.54	**-3.41**	-1.86	2	9
Telecommunications	-	-	-0.34	-1.73	0.55	-	-	-1.31	-0.55	0	5
Utilities	-0.23	-	-1.24	0.32	**-2.86**	-0.63	-	-1.43	**-2.10**	2	7
Financials	-1.79	**-2.46**	-1.16	-0.19	-0.43	-0.26	**-2.48**	**-3.40**	-1.70	3	9
Technology	-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	-

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. The sample period is 1973 to 2007.

	1 lag		6 lags		12 lags	
	$\Delta o \overset{G}{\nrightarrow} r$	$r \overset{G}{\nrightarrow} \Delta o$	$\Delta o \overset{G}{\nrightarrow} r$	$r \overset{G}{\nrightarrow} \Delta o$	$\Delta o \overset{G}{\nrightarrow} r$	$r \overset{G}{\nrightarrow} \Delta o$
Canada	0.27	0.39	0.82	0.23	0.95	0.86
France	**0.00**	0.97	**0.02**	0.85	0.24	0.74
Germany	**0.00**	0.77	**0.00**	0.26	0.09	0.53
Italy	**0.00**	0.93	**0.00**	0.94	**0.01**	0.55
Japan	0.07	0.78	0.07	0.18	0.11	0.59
Norway	0.45	0.36	0.35	0.15	0.52	0.30
U.K.	**0.00**	0.93	**0.01**	0.86	**0.02**	0.95
U.S.	**0.00**	0.20	**0.02**	0.29	**0.03**	0.84
World	**0.00**	0.57	**0.01**	0.51	0.06	0.78

Table 5: Robustness tests of the predictability

The table tests whether the predictability of equity premiums by lagged changes in the oil price is robust to the inclusion of forecasting variables that have been used in the extant literature. The columns report the least squares estimates, t-statistics (in parentheses), and adjusted R^2. The standard errors used in the computation of the t-statistics are GMM-corrected using Newey and West (1987). The predictors of the equity premium are all lagged one period and include the oil price change, equity premium, dividend-price ratio, first difference in 3-month T-bill rate, inflation, and term premium. The sample period is 1973 to 2007. *Significant at 10% level. **Significant at 5% level.

	Δo	r^e	DP	Δr_f	ΔCPI	TERM	$\overline{R}^2$(%)
Canada	-0.03	0.03	0.16	-0.10	-0.63	1.36	0.30
	(-1.23)	(0.59)	(0.56)	**(-1.86)***	(-0.94)	(0.56)	
France	-0.12	0.07	0.04	0.03	-0.25	3.62	2.76
	(-3.02)**	(1.34)	(0.21)	(0.45)	(-0.26)	(1.17)	
Germany	-0.14	0.03	0.16	-0.08	-0.62	1.52	3.60
	(-3.76)**	(0.52)	(0.55)	(-0.89)	(-0.75)	(0.48)	
Italy	-0.20	0.05	0.19	0.01	-0.51	3.01	5.45
	(-4.21)**	(1.23)	(0.63)	(0.20)	(-0.73)	(1.06)	
Japan	-0.05	0.06	0.69	0.10	-0.73	0.60	0.76
	(-1.58)	(1.05)	(1.51)	(0.73)	**(-1.83)***	(0.13)	
Norway	-0.02	0.14	0.57	0.00	-1.05	3.74	2.16
	(-0.50)	**(2.63)****	**(1.83)***	(-0.18)	(-1.35)	**(1.81)***	
U.K.	-0.10	0.09	1.05	-0.05	-1.44	-1.62	6.53
	(-3.71)**	(1.39)	**(2.14)****	(-0.92)	**(-3.21)****	(-1.15)	
U.S.	-0.08	-0.03	0.29	-0.16	-1.71	0.30	6.02
	(-3.31)**	(-0.55)	(1.56)	**(-2.41)****	**(-2.89)****	(0.14)	
World	-0.08	0.04	0.42	-0.13	-1.99	-0.51	7.35
	(-3.50)**	(0.70)	**(2.00)****	**(-2.92)****	**(-3.54)****	(-0.26)	

Table 6: OOS predictability tests

The table shows in-sample (R^2_{IS}) and out-of-sample (R^2_{OOS}) goodness of fit measures. In addition, the MSE-F statistic from McCracken (2007) and the encompassing statistic from Clark and McCracken (2001) are reported. The sample period is 1973 to 2007. ***Significant at 1% level.

	$R^2_{IS}(\%)$	$R^2_{OOS}(\%)$	MSE-F	ENC-NEW
Canada	0.25	-0.31	-0.98	0.02
France	3.22	2.42	10.61***	16.84***
Germany	4.50	3.01	12.38***	18.98***
Italy	5.86	5.13	21.6***	25.55***
Japan	0.79	-0.71	-2.83	2.10
Norway	0.03	-0.86	-3.08	-1.20
U.K.	3.04	2.37	9.87***	16.9***
U.S.	3.11	2.38	10.65***	14.08***
World	3.72	2.89	12.36***	17.16***

Table 7: Overlapping and non-overlapping quarterly and annual returns

This table replicates Table 2 using both non-overlapping and overlapping returns at quarterly and annual frequencies. The GMM-corrected t-statistics are computed using the Newey and West (1987) procedure with 2(h-1) lags, where h is the degree of overlap. The sample period is 1973 to 2007.

	Non-overlapping				Overlapping			
	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
				Panel A: Quarterly returns				
Canada	0.04	0.82	0.86	0.50	-0.02	-0.46	-0.48	0.05
France	0.00	-0.07	-0.10	0.00	-0.16	**-2.47**	**-2.35**	1.45
Germany	0.03	0.52	0.70	0.20	-0.14	**-2.33**	**-2.82**	1.30
Italy	-0.05	-0.67	-0.74	0.33	-0.26	**-3.58**	**-2.86**	3.01
Japan	0.01	0.18	0.22	0.02	-0.10	-1.81	-1.74	0.78
Norway	-0.06	-0.84	-1.21	0.53	-0.07	-0.87	-0.83	0.18
U.K.	0.00	0.08	0.12	0.00	-0.14	**-2.36**	**-2.41**	1.32
U.S.	0.01	0.12	0.15	0.01	-0.09	-1.93	**-2.01**	0.89
World	0.01	0.13	0.16	0.01	-0.11	**-2.65**	**-2.69**	1.67
				Panel B: Annual returns				
Canada	0.08	0.83	0.98	2.18	-0.05	-0.41	-0.28	0.04
France	0.09	0.56	0.56	1.01	-0.35	**-2.39**	-1.55	1.39
Germany	0.16	0.99	0.89	3.08	-0.12	-0.85	-1.02	0.18
Italy	0.19	1.11	1.16	3.85	-0.30	-1.79	-1.58	0.78
Japan	-0.04	-0.32	-0.32	0.33	-0.24	-1.91	-1.69	0.90
Norway	0.10	0.49	0.54	0.78	-0.14	-0.77	-0.53	0.15
U.K.	0.15	1.37	1.15	5.73	-0.41	**-3.39**	-1.44	2.76
U.S.	0.07	0.72	0.56	1.65	-0.22	**-2.34**	-1.50	1.34
World	0.05	0.58	0.47	1.06	-0.24	**-2.61**	-1.68	1.66

Table 8: 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 R^2. The left panel considers positive oil price changes only; the right panel considers negative oil price changes only. The sample period is 1973 to 2007.

	Δo_t^+				Δo_t^-			
	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Canada	0.01	0.15	0.13	0.01	-0.09	-1.91	**-2.94**	0.87
France	-0.12	**-2.23**	**-2.03**	1.18	-0.22	**-3.81**	**-3.99**	3.36
Germany	-0.18	**-3.43**	**-2.94**	2.74	-0.21	**-3.64**	**-3.78**	3.09
Italy	-0.15	**-2.45**	**-2.41**	1.41	-0.39	**-5.91**	**-5.80**	7.72
Japan	-0.03	-0.60	-0.55	0.09	-0.12	**-2.42**	**-2.49**	1.38
Norway	0.01	0.21	0.17	0.01	-0.06	-0.83	-0.74	0.17
U.K.	-0.15	**-2.92**	**-3.09**	2.01	-0.16	**-2.85**	**-4.22**	1.91
U.S.	-0.09	**-2.19**	**-2.04**	1.14	-0.16	**-3.72**	**-4.58**	3.22
World	-0.08	**-2.17**	-1.84	1.12	-0.17	**-4.36**	**-6.12**	4.35

Table 9: 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 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. (18), (20) and (21). The upper panel uses a horizon of 12 months to calculate the maximum or minimum return whereas the lower panel uses 36 months. The sample period is 1973 to 2007.

	$\Delta o_t^{\dagger,+}$				$\Delta o_t^{\dagger,-}$				$\Delta o_t^{\dagger,\pm}$			
	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Panel A: max/min returns using horizon k=12												
Canada	-0.03	-0.53	-0.59	0.07	-0.10	-1.33	**-2.26**	0.42	-0.06	-1.23	-1.74	0.36
France	-0.10	-1.24	-1.27	0.36	-0.33	**-3.66**	**-5.86**	3.11	-0.18	**-3.22**	**-3.01**	2.43
Germany	-0.14	-1.84	-1.86	0.81	-0.20	**-2.27**	**-2.42**	1.22	-0.15	**-2.79**	**-3.15**	1.84
Italy	-0.07	-0.78	-0.99	0.14	-0.55	**-5.48**	**-8.26**	6.72	-0.26	**-4.00**	**-2.67**	3.70
Japan	-0.06	-0.84	-0.67	0.17	-0.18	**-2.35**	**-2.70**	1.30	-0.10	**-2.11**	-1.61	1.05
Norway	-0.13	-1.34	-1.81	0.43	0.01	0.06	0.06	0.00	-0.07	-0.94	-0.84	0.21
U.K.	-0.20	**-2.69**	**-2.71**	1.71	-0.20	**-2.29**	**-4.66**	1.25	-0.19	**-3.43**	**-4.32**	2.75
U.S.	-0.10	-1.73	**-2.12**	0.71	-0.16	**-2.45**	**-3.67**	1.42	-0.12	**-2.82**	**-4.29**	1.88
World	-0.09	-1.74	-1.55	0.72	-0.20	**-3.34**	**-5.36**	2.61	-0.13	**-3.40**	**-3.68**	2.69
Panel B: max/min returns using horizon k=36												
Canada	-0.05	-0.74	-0.88	0.13	-0.12	-1.51	**-2.19**	0.54	-0.08	-1.52	**-2.28**	0.55
France	-0.10	-1.20	-1.23	0.34	-0.35	**-3.62**	**-6.16**	3.04	-0.20	**-3.20**	**-2.99**	2.40
Germany	-0.12	-1.55	-1.61	0.57	-0.15	-1.65	**-2.34**	0.65	-0.13	**-2.21**	**-2.68**	1.16
Italy	-0.05	-0.48	-0.61	0.06	-0.59	**-5.44**	**-7.54**	6.62	-0.26	**-3.77**	**-2.37**	3.30
Japan	-0.04	-0.58	-0.46	0.08	-0.19	**-2.25**	**-2.55**	1.20	-0.10	-1.87	-1.36	0.83
Norway	-0.17	-1.65	**-2.61**	0.65	0.05	0.39	0.46	0.04	-0.07	-0.98	-0.83	0.23
U.K.	-0.19	**-2.38**	**-2.33**	1.35	-0.19	**-2.03**	**-4.59**	0.97	-0.18	**-3.09**	**-3.75**	2.23
U.S.	-0.11	-1.77	**-2.22**	0.74	-0.18	**-2.46**	**-3.20**	1.43	-0.13	**-2.90**	**-4.37**	1.97
World	-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 10: 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 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. (22) to (24) 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. The sample period is 1973 to 2007.

	$\Delta\tilde{o}_t^{\dagger,+}$				$\Delta\tilde{o}_t^{\dagger,-}$				$\Delta\tilde{o}_t^{\dagger,\pm}$			
	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Panel A: max/min returns using horizon k=12												
Canada	0.002	1.00	1.08	0.24	-0.001	-0.63	-0.47	0.09	0.000	0.33	0.29	0.03
France	-0.001	-0.54	-0.57	0.07	0.004	**2.03**	1.66	0.98	0.002	1.13	0.97	0.31
Germany	-0.002	-1.26	-1.34	0.38	0.003	1.54	1.45	0.56	0.000	0.17	0.16	0.01
Italy	0.001	0.23	0.29	0.01	0.005	1.95	1.37	0.90	0.003	1.71	1.59	0.69
Japan	-0.002	-1.40	-1.60	0.47	0.000	-0.21	-0.20	0.01	-0.002	-1.30	-1.45	0.40
Norway	0.001	0.30	0.24	0.02	0.001	0.18	0.17	0.01	0.001	0.38	0.33	0.04
U.K.	-0.003	-1.79	**-2.18**	0.76	0.004	1.75	**1.99**	0.73	0.000	-0.09	-0.09	0.00
U.S.	-0.001	-0.74	-0.89	0.13	0.001	0.92	0.90	0.20	0.000	0.11	0.11	0.00
World	-0.001	-1.05	-1.26	0.26	0.002	1.23	1.07	0.36	0.000	0.10	0.10	0.00
Panel B: max/min returns using horizon k=36												
Canada	0.001	0.27	0.30	0.02	0.003	1.27	1.23	0.38	0.002	1.06	1.09	0.27
France	-0.001	-0.62	-0.75	0.09	0.008	**2.88**	**2.35**	1.95	0.003	1.37	1.26	0.45
Germany	-0.004	-1.56	**-2.08**	0.58	0.003	1.09	1.08	0.29	-0.001	-0.57	-0.63	0.08
Italy	0.002	0.57	0.71	0.08	0.013	**4.00**	**3.24**	3.70	0.007	**3.09**	**3.00**	2.23
Japan	-0.003	-1.27	-1.38	0.38	0.003	0.96	1.05	0.22	-0.001	-0.41	-0.44	0.04
Norway	-0.001	-0.49	-0.40	0.06	0.001	0.20	0.18	0.01	-0.001	-0.27	-0.23	0.02
U.K.	-0.004	-1.88	**-2.05**	0.84	0.004	1.32	1.78	0.42	-0.001	-0.67	-0.71	0.11
U.S.	-0.002	-0.93	-1.18	0.20	0.003	1.56	1.81	0.58	0.000	0.26	0.29	0.02
World	-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 11: 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 R^2. The left panel is the change in oil price scaled by the GARCH volatility as in Eq. (28); the middle panel considers positive changes only and the right panel considers negative changes only, as described in Eqs. (29) and (30). The sample period is 1973 to 2007.

	$\Delta o_t^{\ddagger}$				$\Delta o_t^{\ddagger,+}$				$\Delta o_t^{\ddagger,-}$			
	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$	b	t_{OLS}	t_{NW}	$R^2(\%)$
Canada	-0.001	-0.67	-0.62	0.11	-0.001	-0.30	-0.23	0.02	-0.003	-0.88	-1.13	0.19
France	-0.007	**-3.18**	**-2.84**	2.37	-0.009	**-2.69**	**-2.47**	1.70	-0.011	**-2.46**	**-2.45**	1.43
Germany	-0.009	**-4.38**	**-4.09**	4.39	-0.012	**-3.86**	**-3.59**	3.46	-0.013	**-3.13**	**-3.34**	2.30
Italy	-0.009	**-3.58**	**-3.21**	2.99	-0.008	**-2.05**	**-1.98**	0.99	-0.020	**-4.10**	**-3.72**	3.88
Japan	-0.003	-1.33	-1.27	0.42	-0.005	-1.74	-1.70	0.72	-0.001	-0.22	-0.25	0.01
Norway	-0.001	-0.39	-0.34	0.04	0.002	0.51	0.41	0.06	-0.008	-1.43	-1.44	0.49
U.K.	-0.007	**-3.15**	**-3.46**	2.33	-0.009	**-2.86**	**-3.07**	1.93	-0.009	**-2.18**	**-2.57**	1.12
U.S.	-0.006	**-3.55**	**-3.64**	2.93	-0.007	**-2.89**	**-2.67**	1.96	-0.009	**-2.87**	**-3.23**	1.94
World	-0.006	**-3.69**	**-3.16**	3.17	-0.007	**-3.18**	**-2.47**	2.37	-0.008	**-2.76**	**-3.07**	1.79

Table 12: Tests for non-linearity and asymmetry

The table shows p-values for various Wald tests. Panel A reports p-values corresponding to the null hypothesis that the coefficient on the variable in each column is zero, when included as an additional regressor. The columns correspond to the maximum measures (Eqs. (18), (20) and (21)) and scaled measures (Eqs. (28) to (30)) respectively. Panel B reports p-values corresponding to the null hypothesis that the coefficient on positive and negative changes are identical. The columns correspond to tests for equality of the coefficients on Eqs. (16) and (17), Eqs. (18) and (20), and Eqs. (29) and (30) respectively. All F-statistics have been GMM-corrected using Newey and West (1987). P-values less than 0.1 are in bold. The sample period is 1973 to 2007.

Panel A: Wald test for significance of coefficients						
	$\Delta o_t^{\dagger,+}$	$\Delta o_t^{\dagger,-}$	$\Delta o_t^{\dagger,\pm}$	$\Delta o_t^{\ddagger}$	$\Delta o_t^{\ddagger,+}$	$\Delta o_t^{\ddagger,-}$
Canada	0.99	0.13	0.29	0.81	0.73	0.76
France	0.42	**0.00**	0.36	0.82	0.82	0.99
Germany	0.59	0.74	0.80	0.22	0.27	0.77
Italy	**0.00**	**0.00**	0.55	0.44	0.22	0.24
Japan	0.93	**0.07**	0.43	0.83	0.50	0.11
Norway	**0.07**	0.78	0.28	0.87	0.42	**0.08**
U.K.	0.38	0.16	0.14	0.66	0.54	0.73
U.S.	0.81	0.20	0.57	0.29	0.59	0.48
World	0.72	**0.01**	0.43	0.53	0.62	0.85
Panel B: Wald test for equality of coefficients						
	Δo_t	$\Delta o_t^{\dagger}$	$\Delta o_t^{\ddagger}$			
Canada	**0.09**	0.38	0.64			
France	0.14	**0.02**	0.91			
Germany	0.77	0.62	0.84			
Italy	**0.01**	**0.00**	**0.06**			
Japan	0.18	0.26	0.18			
Norway	0.45	0.26	0.13			
U.K.	0.96	0.91	0.76			
U.S.	0.19	0.37	0.79			
World	**0.08**	0.13	0.98			

Figure 1: Oil prices and important oil events

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.

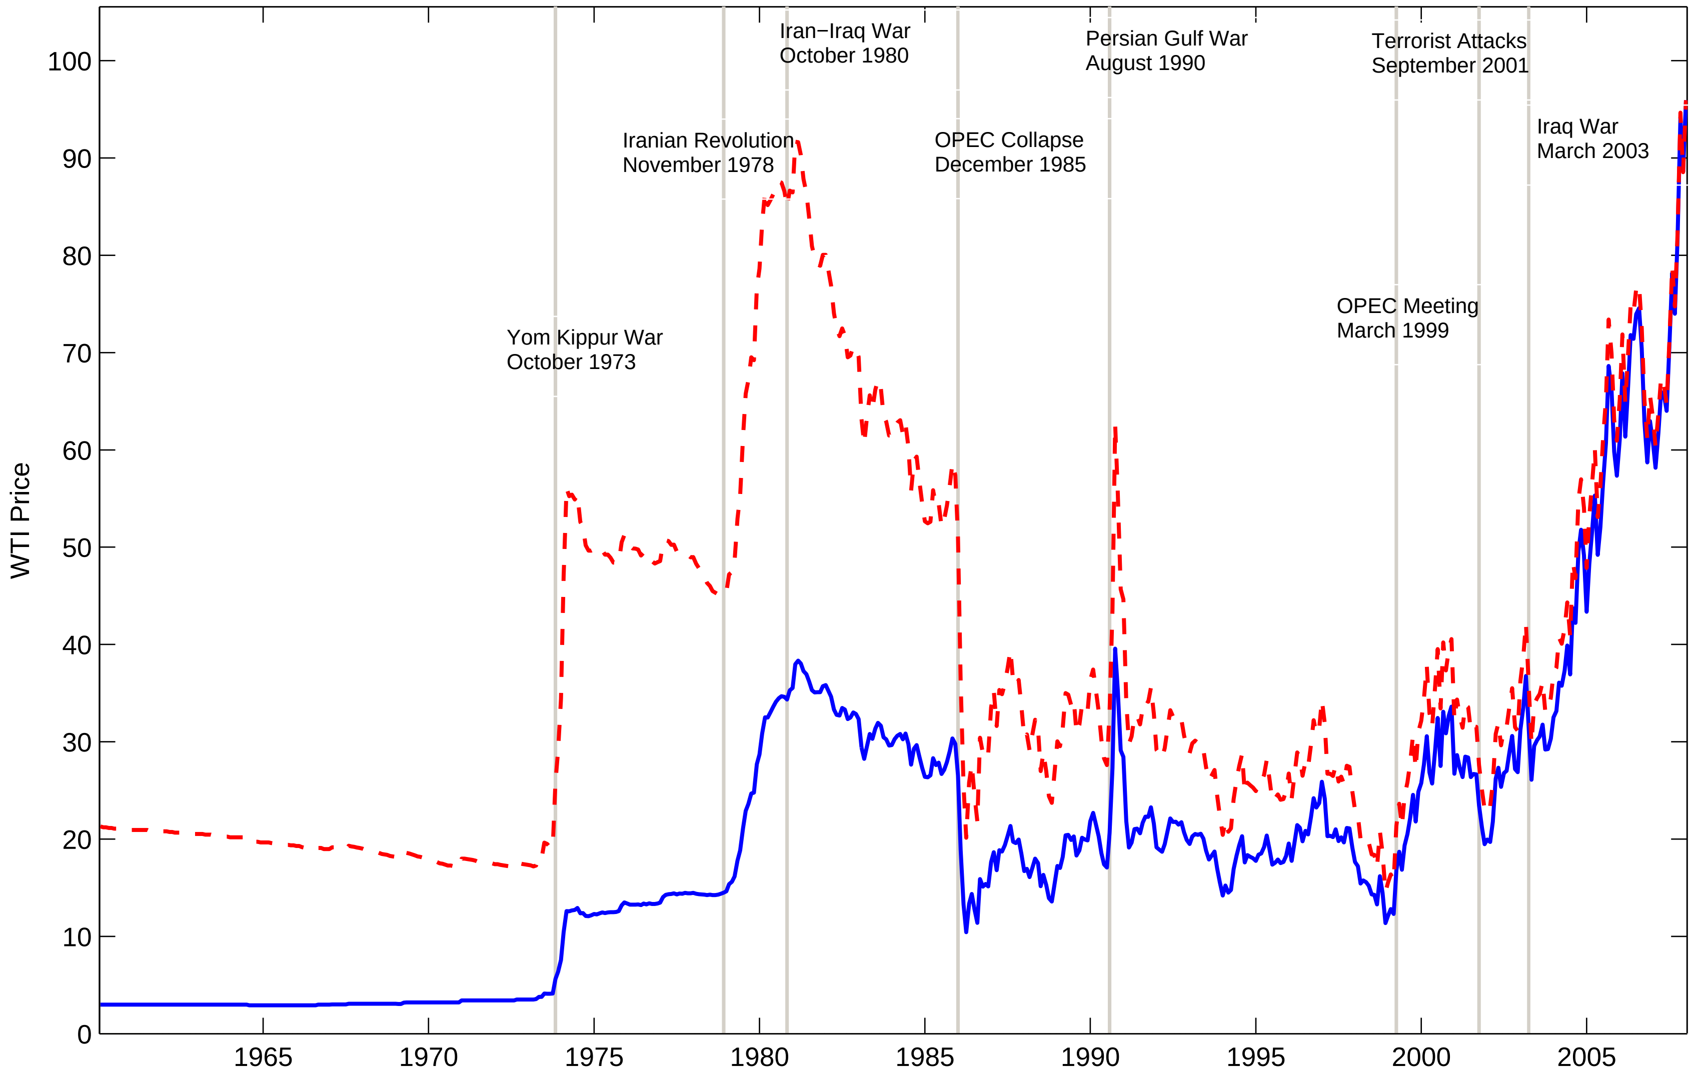

Figure 2: Recursive beta coefficients in the U.S.

The figure plots the recursively estimated beta coefficient on lagged oil price changes in the predictive univariate regression. The upper panel shows important oil dates marked as vertical bars whereas the lower panel shows NBER recession dates as vertical bars.

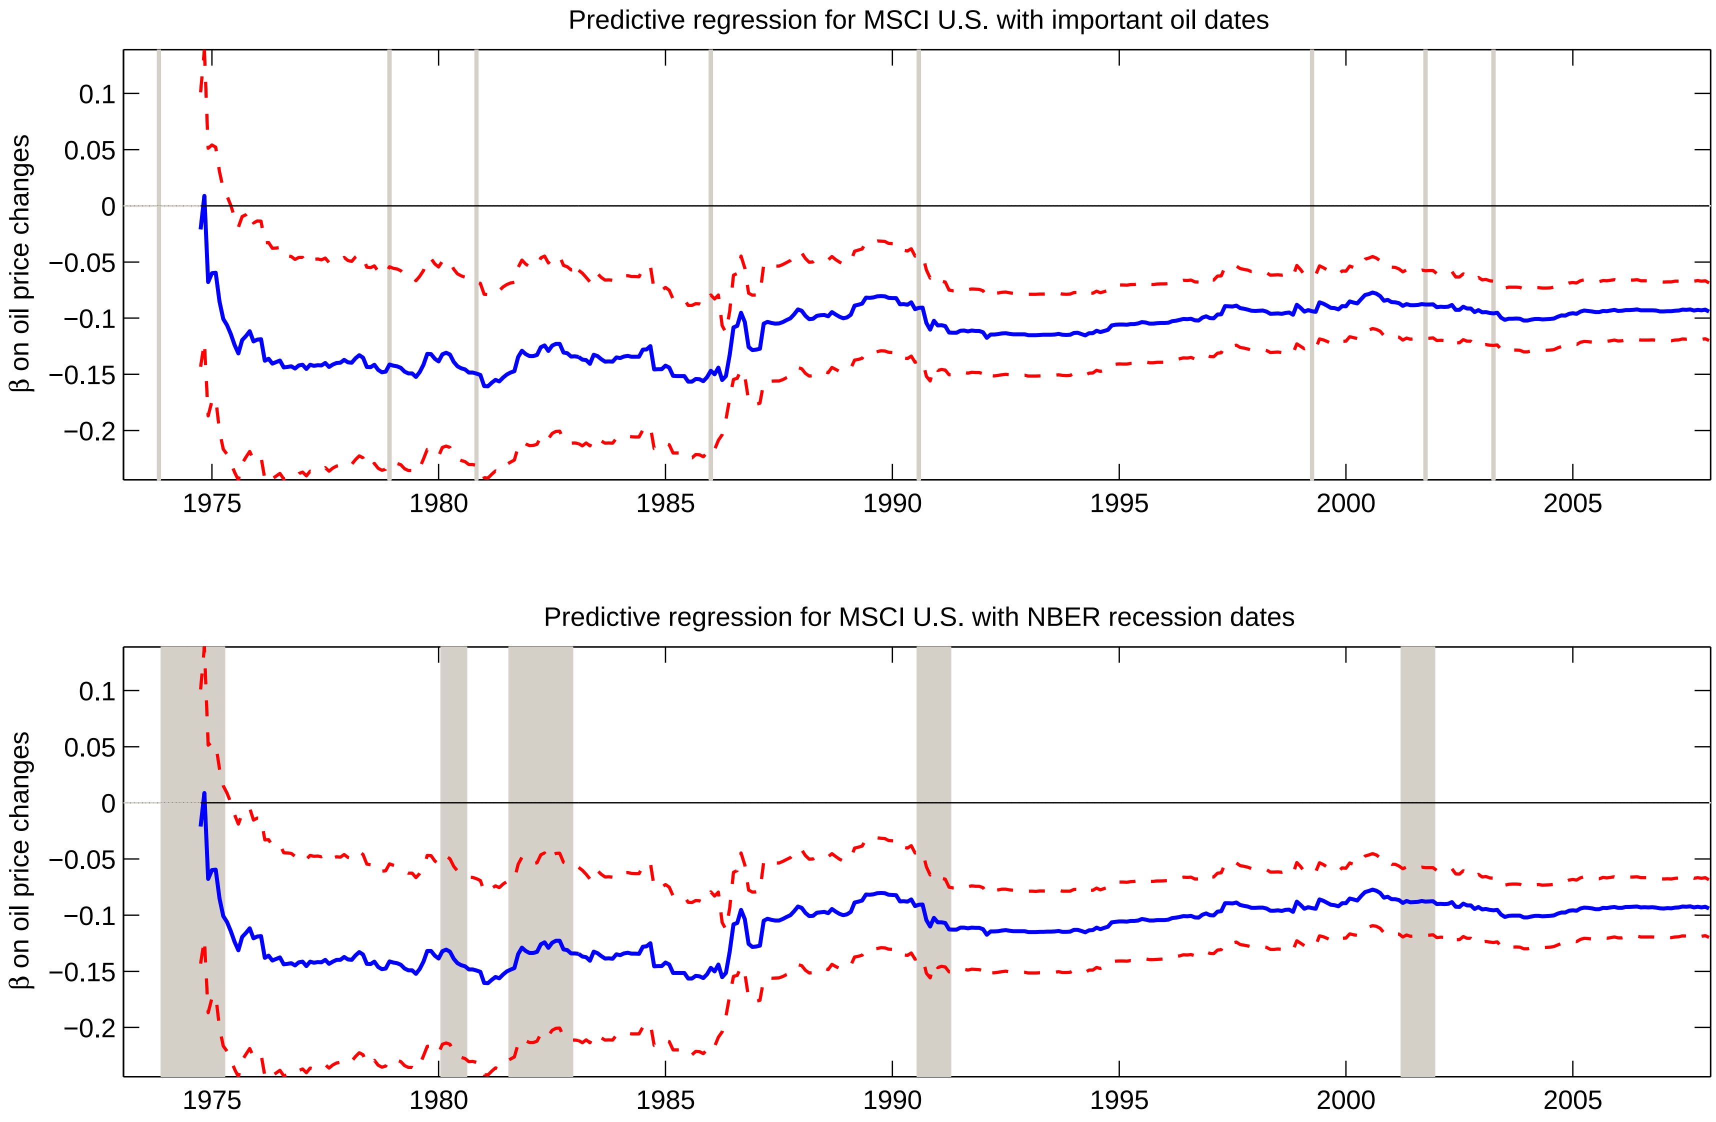

Figure 3: International recursive beta coefficients

The figure plots the recursively estimated beta coefficients on lagged oil price changes in the predictive univariate regressions. The vertical bars mark important oil dates.

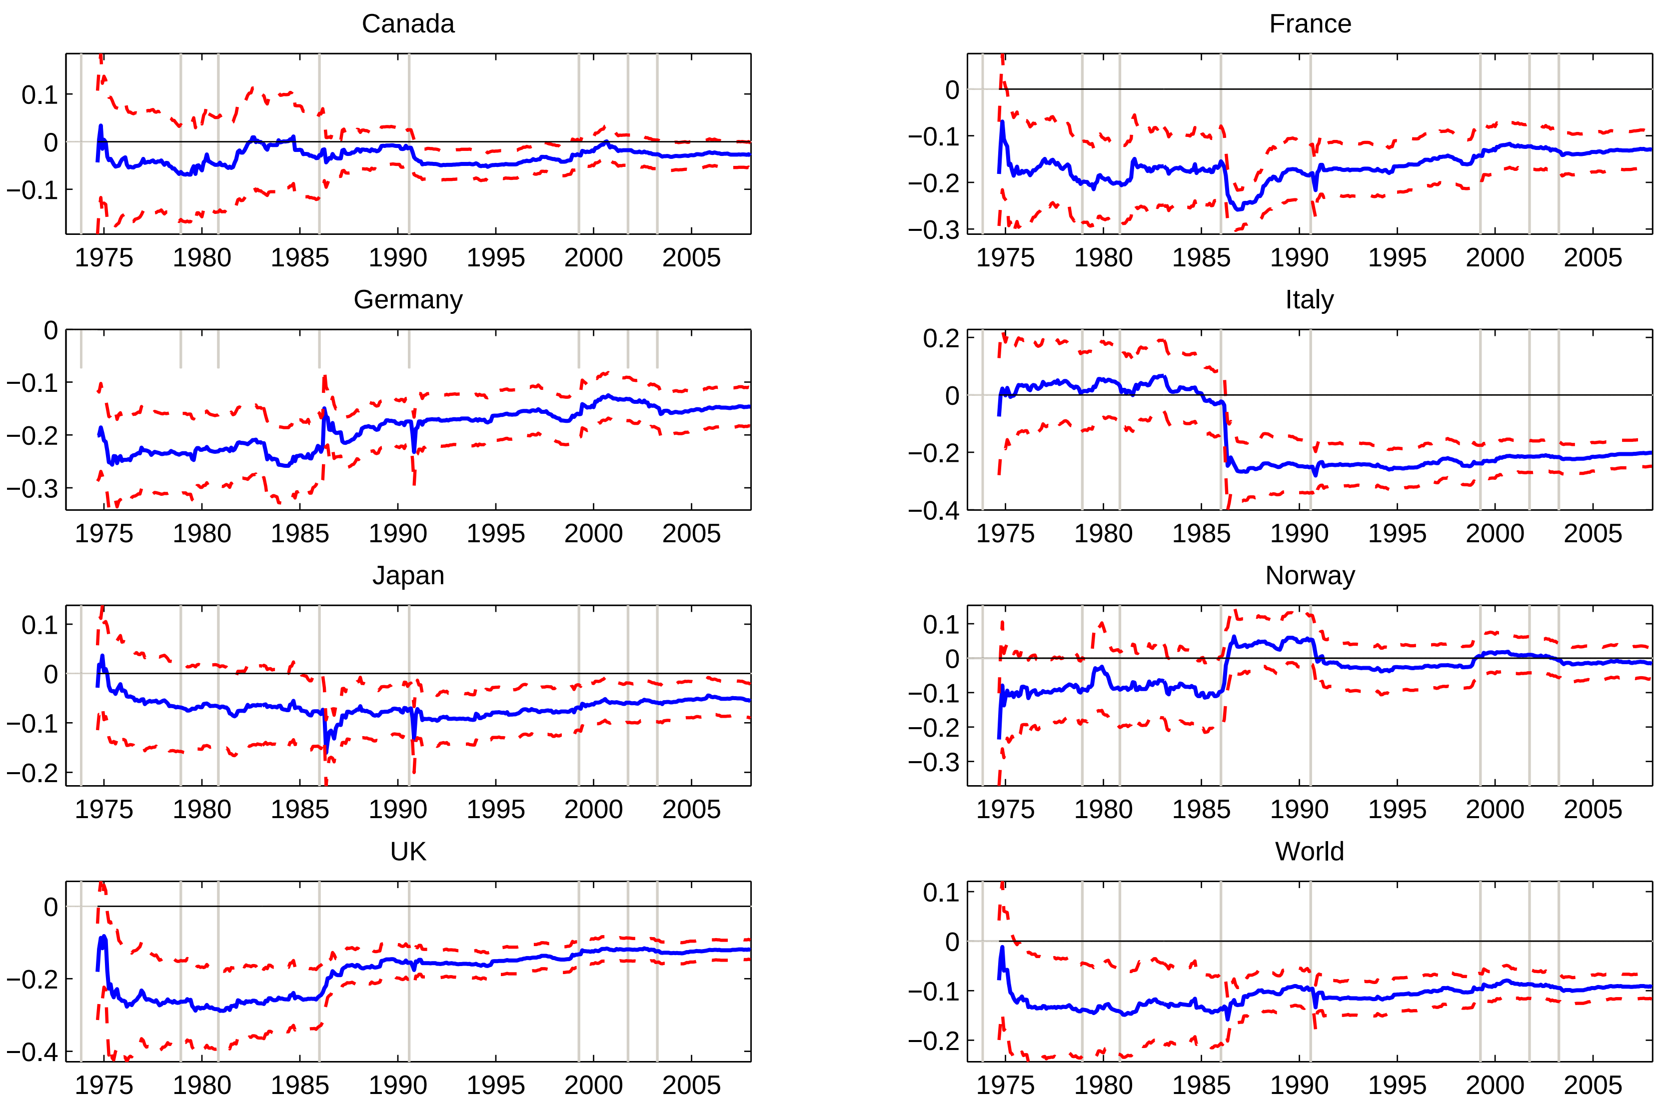

Figure 4: Cumulative sum of squared errors graphs in the U.S.

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.

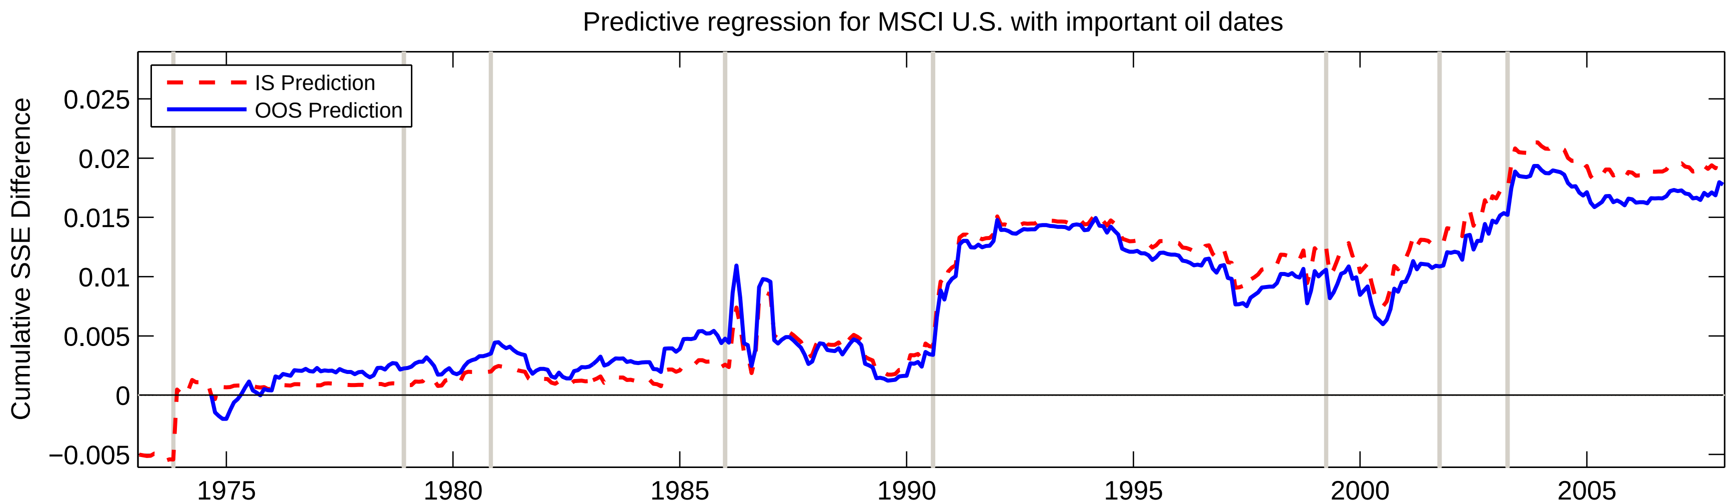

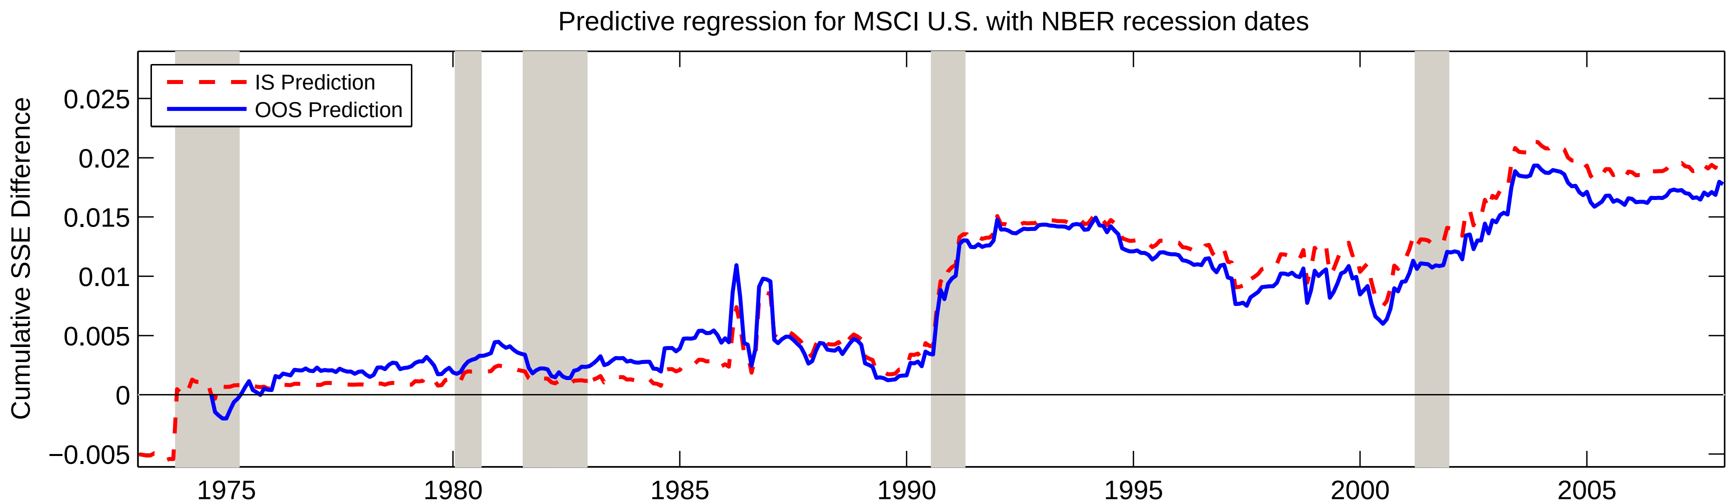

Figure 5: International cumulative sum of squared errors graphs

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.

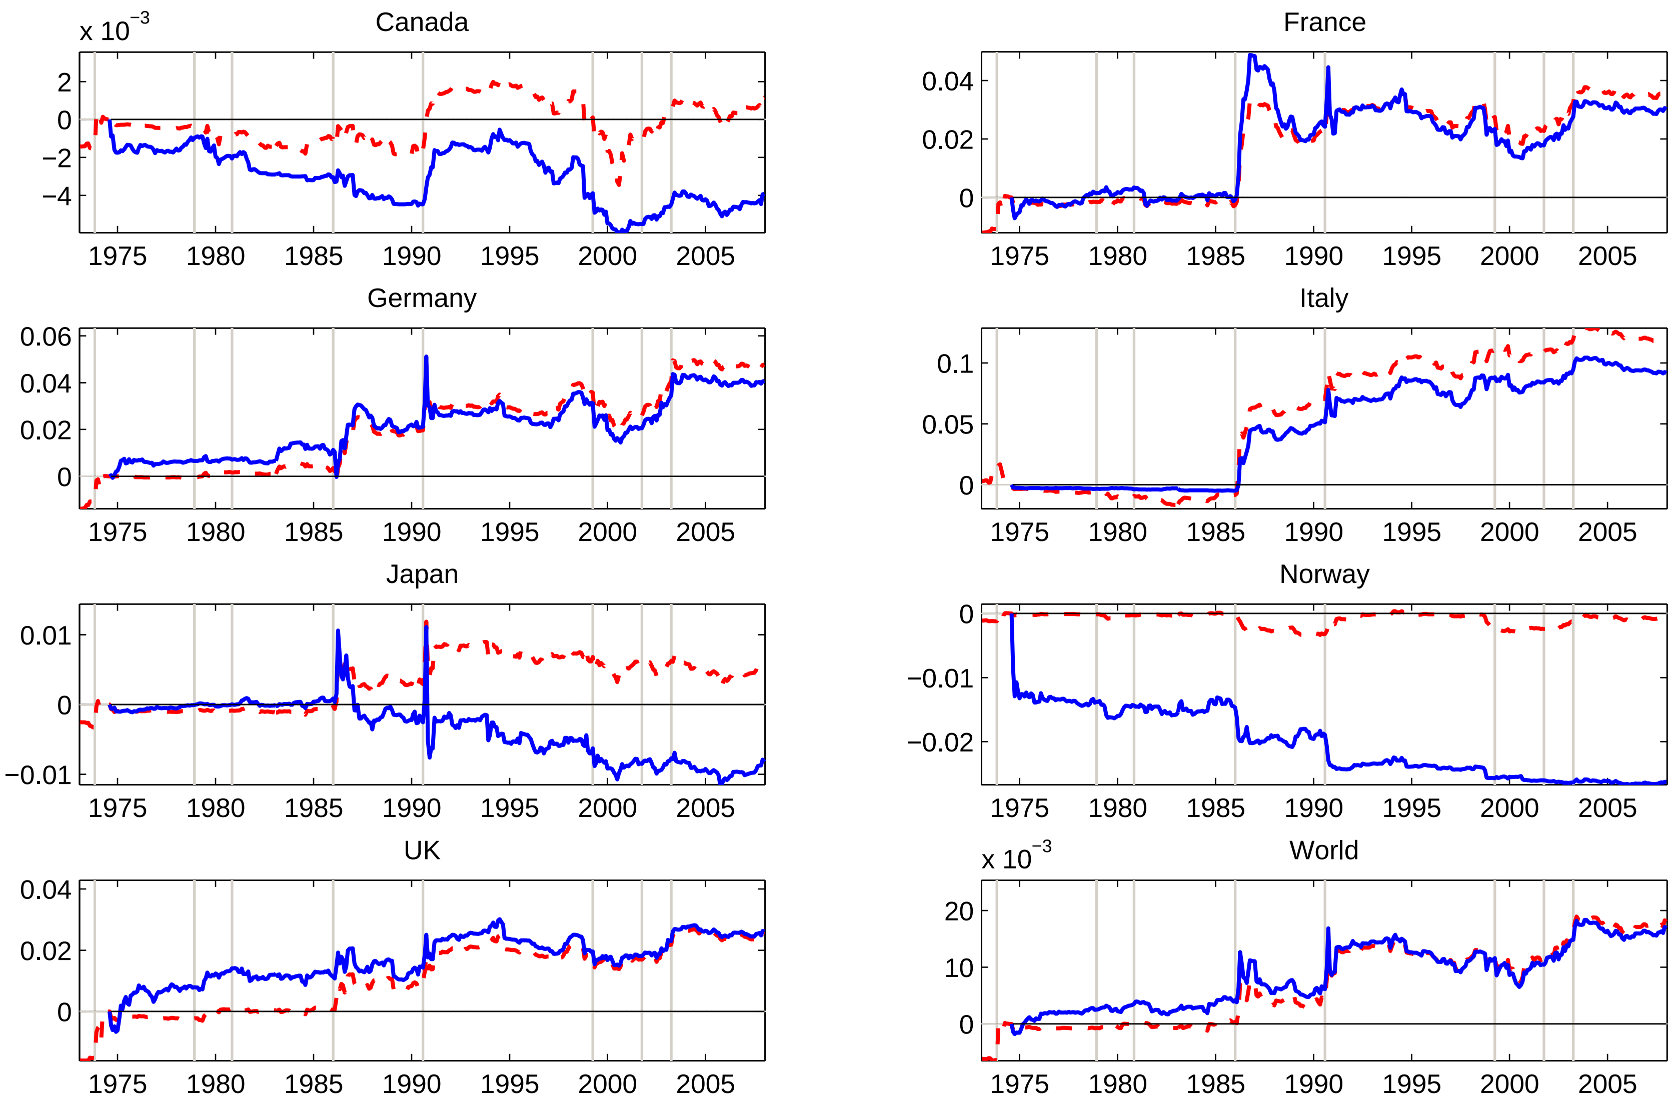

Figure 6: Cumulative sum of squared errors graphs after removing "oil dates"

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.

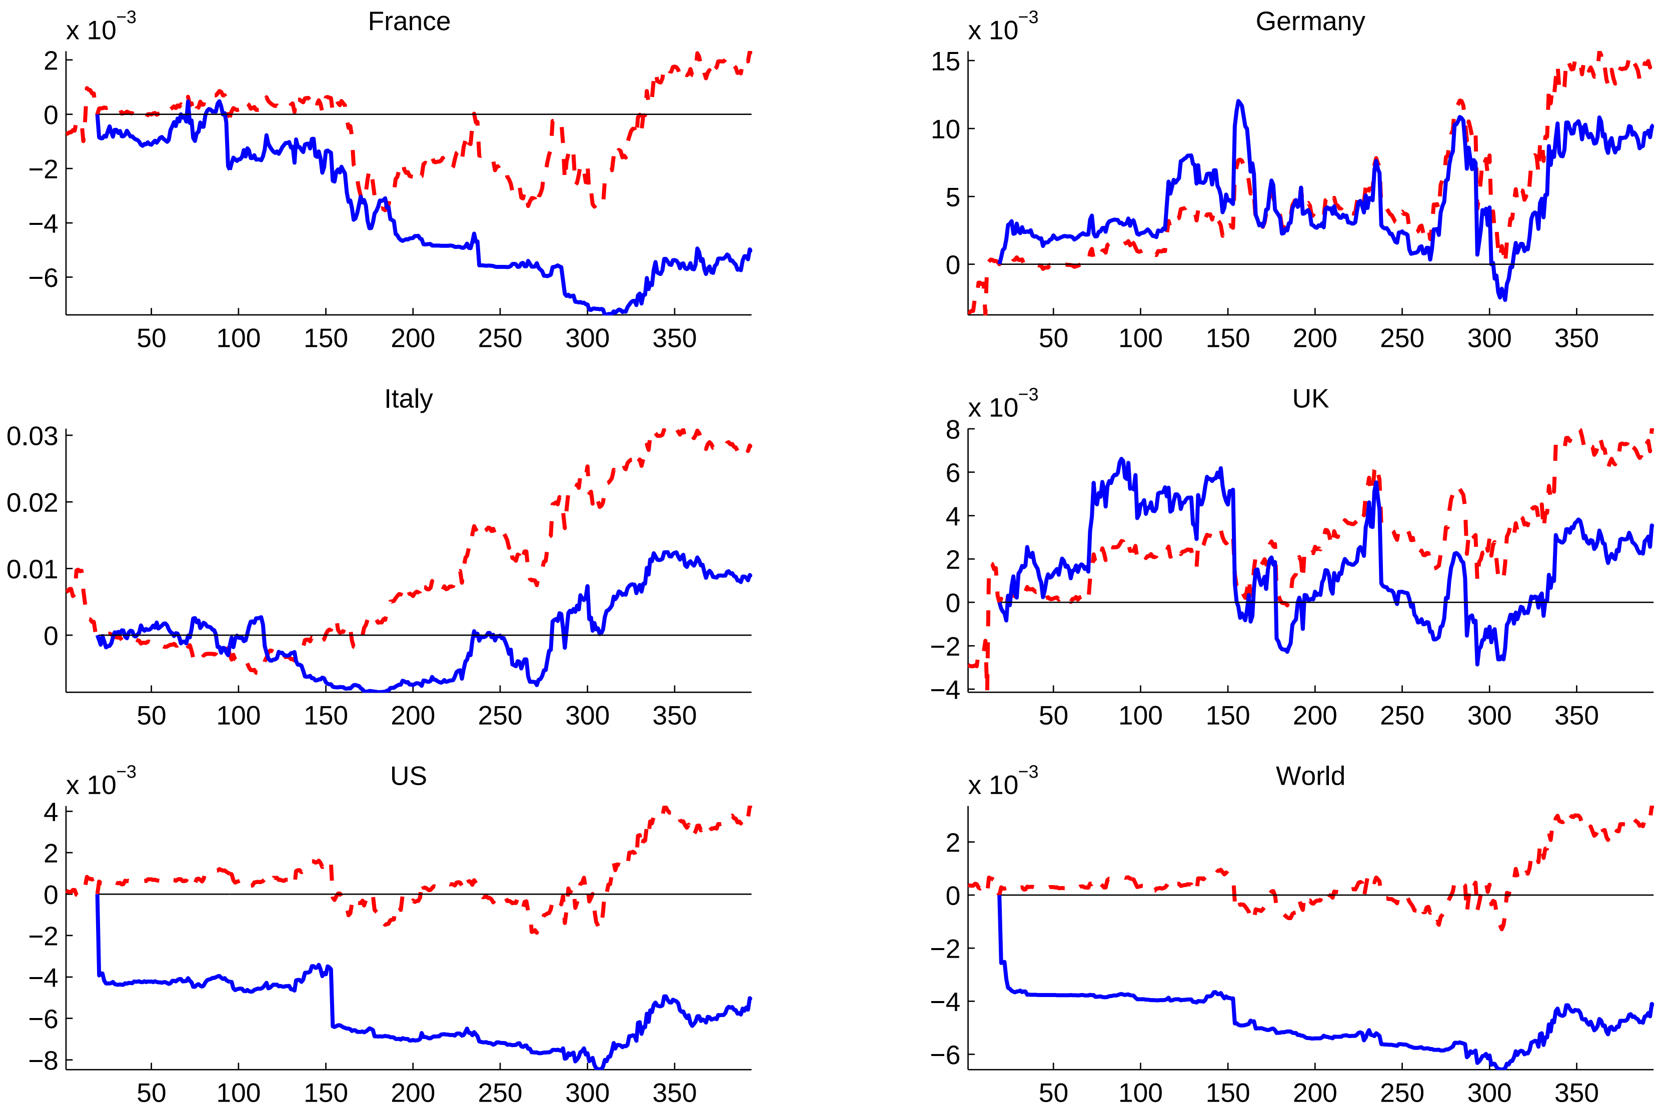

Figure 7: Max oil price increases and decreases

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. (18) and (20). The vertical bars mark the important oil dates.