Energy Economics 136 (2024) 107659

Contents lists available at ScienceDirect

Energy Economics

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

Betting on war? Oil prices, stock returns, and extreme geopolitical events

Knut Nygaard [a,*,1,2], Lars Qvigstad Sørensen [b,1,2]

[a] Oslo Metropolitan University, Oslo Business School, Norway
[b] Storebrand Asset Management, Norway

ARTICLE INFO

JEL classification:
G11
G14
G15
G17

Keywords:
Return predictability
Oil prices
International stock markets
Market efficiency
Stock returns

ABSTRACT

We show that the ability of oil price changes to predict stock returns is limited to periods of extreme geopolitical unrest. Four events generate most of the predictability: the 1973 Arab-Israel war, the 1986 OPEC collapse, the 1990/91 Persian gulf war, and the 2003 invasion of Iraq. We also find that a market-timing trading strategy based on oil price changes typically generates insignificant abnormal returns, contradicting previously published results. Our findings serve as an example of how a significant predictor in a time series forecasting regression may not be a useful or profitable market-timing signal.

1. Introduction

After months of escalating geopolitical tensions, Russia launched a full-scale military invasion of Ukraine on February 24, 2022. Oil prices increased by close to 30% over the first two months of 2022, while stock markets were down. This phenomenon of a negative correlation between oil price changes and stock returns during periods of extreme geopolitical unrest is not unique to the Russia-Ukraine war. Extreme geopolitical events often coincide with a jump in oil prices and a simultaneous drop in stock markets, as documented by, e.g., Wolfers and Zitzewitz (2009) and Omar et al. (2017). In this article, we examine whether extreme geopolitical events also correspond to periods where *lagged* oil price changes are negatively correlated with stock returns.

Using past oil price changes to forecast stock returns is well-established in the literature, see, e.g., Smyth and Narayan (2018) for an overview. Driesprong et al. (2008), hereafter referred to as DJM, established this oil puzzle: Global stock returns tend to be lower if the oil price increased in the previous month and higher if the oil price fell in the previous month. They posit this predictive relation as an exploitable and profitable stock market anomaly, possibly stemming from gradual information diffusion and investor underreaction (Hong and Stein, 1999). In contrast, we investigate whether the ability of oil price changes to predict stock returns is limited to periods of well-known oil price shocks following extreme geopolitical events, rather than being evidence of a more general market inefficiency.

Recent research indicates that the predictive relationship from oil price changes to stock returns attenuates from around when DJM's sample period ends in 2003, see, e.g., Jiang et al. (2018), Kim et al. (2019), Cederburg et al. (2023), and Goyal et al. (2023). As long as DJM's initial oil puzzle stands unapproached, this attenuation could be interpreted as investors learning about market inefficiency from the academic literature, similarly to what McLean and Pontiff (2016) show is the case for certain other stock market anomalies. However, we contend that it is more likely due to a lack of geopolitical events jointly disrupting oil and stock markets with the same force as in the DJM sample, after 2003.

Our main finding is that oil price shocks coinciding with extreme geopolitical events make an outsized contribution to the ability of oil price changes to forecast stock returns. We find that lagged oil price changes no longer predict stock returns in the DJM sample period in the counterfactual scenario where the extreme events did not occur. Four events generate most of the predictability: the 1973 Arab-Israel war, the 1986 OPEC collapse, the 1990/91 Persian gulf war, and the

* Corresponding author.
E-mail addresses: knut.nygaard@oslomet.no (K. Nygaard), lars.qvigstad.sorensen@storebrand.no (L.Q. Sørensen).

[1] The opinions expressed in this paper are solely those of the authors and are not intended to represent the view of Storebrand Asset Management. We received valuable comments from B. Espen Eckbo, Sturla Fjesme, Ben Jacobsen, Kam-Ming Wan and the referees. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The authors have no conflict of interest to disclose.

[2] All authors contributed equally to this research project.

https://doi.org/10.1016/j.eneco.2024.107659
Received 19 June 2023; Received in revised form 14 May 2024; Accepted 23 May 2024
Available online 28 May 2024

2003 invasion of Iraq. These four major events all occurred in the DJM sample period, which starts in October 1973 and ends in April 2003.

We establish the substantial influence of these extreme events using the graphical tools proposed by Goyal and Welch (2003, 2008). These graphs plot the cumulative difference over time between the prediction errors using two different models to forecast next month's stock return: the unconditional mean stock market return (the benchmark) and a regression using lagged oil price changes as the explanatory variable (the candidate forecasting model). A useful candidate forecasting model generates an upwards drift in these graphs, as it should generally predict next period's stock market return better than the benchmark thereby yielding a smaller prediction error. The extreme events correspond to sharp increases in both the in-sample (IS) and out-of-sample (OOS) predictive performance of the oil model. Outside these extreme events, the oil model is not much better than the unconditional mean stock market return at predicting stock market returns. We conclude that DJM's oil puzzle stems from a very small number of extreme geopolitical events in the sample period they investigated.

We also show that DJM's proposed trading strategy generates insignificant alphas in all but two of the G7 and World stock markets in the DJM sample period even without excluding any observations. We provide further detail as to how DJM's reported results seem incorrect based on their reported beta estimates. We would expect the market beta to be of a similar magnitude to the fraction of time that the oil strategy is invested in the stock market. DJM report market beta estimates of 0.5 or less for all indices, which is far below our estimates.[3] In our implementation for the US, for example, the trading strategy implies a stock market investment in more than three quarters of the sample period, which is hard to reconcile with DJM's reported beta estimate for the US of 0.43. Our results highlight that a significant predictor in a time series forecasting regression does not necessarily constitute a useful or profitable market-timing signal.

Our work relates to several strands of the literature. Firstly, our paper is related to the literature that examines sources of stock market predictability. Harvey et al. (2016) and Hou et al. (2020) survey firm characteristics associated with the cross-section of returns, and Rapach and Zhou (2013) review time series predictors of stock returns. Recent research has applied data reduction methods that extract any valuable information from collections of anomalies, see, e.g., Feng et al. (2020) and Dong et al. (2022). A complementary task is to examine what drives anomalies, which is what, e.g., Lochstoer and Tetlock (2020) do for a small set of cross-sectional anomalies. This paper undertakes a similar investigation into the source of predictability for the time-series anomaly that the long-established oil puzzle represents. Secondly, we contribute to the literature that examines how oil price shocks impact stock returns. A large literature establishes that oil price shocks significantly impact stock returns, see, e.g., Jones and Kaul (1996) and Kilian and Park (2009). We build on this literature and identify which oil price shocks matter most, and determine how much they matter, in the prediction of stock returns. Thirdly, our paper addresses the literature about how geopolitical risk influences oil and stock markets. Caldara and Iacoviello (2022) construct newspaper-based indices of geopolitical risk (GPR) and show that increased GPR is associated with a decline in both stock prices and oil prices. Cunado et al. (2020) argue that a decline in oil prices due to an increase in GPR is primarily due to reduced demand for oil. Our paper contributes to this literature by isolating extreme geopolitical events, rather than fluctuations in general geopolitical sentiment, and we establish how these events contribute to the predictive relationship from oil price changes to stock returns.

[3] We are grateful for the cordial communication we have had with Ben Jacobsen, the corresponding author of DJM, where we were informed that the data file used to generate the results in their paper is no longer accessible, which unfortunately prohibits a closer examination of their data and estimation procedure.

The remainder of this paper is organized as follows. Section 2 describes the identification of extreme geopolitical events. Section 3 details our data collection and presents summary statistics. Section 4 replicates DJM's results of oil prices predicting stock returns and investigates the counterfactual scenario where the extreme events did not occur. Section 5 tests the oil model's predictive performance relative to the unconditional mean stock return. Section 6 analyzes a market-timing trading strategy based on the oil price. Section 7 concludes.

2. Identifying extreme geopolitical events

In this section, we describe the identification of extreme geopolitical events. We focus on oil price shocks coinciding with geopolitical trigger events that have been highlighted in the extant literature. In this, we rely primarily on two authoritative reviews of the causes of major oil price shocks, by leading experts in energy economics and the oil market: Hamilton (2013) and Baumeister and Kilian (2016). These reviews highlight eight geopolitical trigger events from 1973: (1) the 1973 Arab-Israel war and Arab embargo, (2) the 1978/79 Iranian revolution, (3) the 1980 outbreak of the Iran-Iraq war, (4) the 1986 OPEC collapse, (5) the 1990/91 Persian Gulf war, (6) the 2002 Venezuelan oil strike, (7) the 2003 invasion of Iraq, and (8) the 2011 Libyan civil war.[4] To this list of events we add the Russian invasion of Ukraine in 2022, which has already received attention in the literature for its disruption to oil markets, see, e.g., Zhang et al. (2023).

We offer a detailed description of each of the nine events in the Online Appendix, and summarize the events in Table 1. Fig. 1 depicts these events together with nominal and real oil prices from 1969 to 2022. Each event corresponds to a sharp change in the oil price. However, the events could not have been identified solely based on oil price changes; there are several other large jumps in the oil price outside of these events. Moreover, events 1 and 5 overlap with NBER recession dates.

Whereas there is agreement on the identification of events, there seems to be some disagreement on the relative importance of each event and the aspects of the events that impact oil prices. Hamilton (2003) argues that changes to global oil supply caused by extreme geopolitical events are critical in understanding oil price shocks. In contrast, Kilian (2009) posits that changes to precautionary demand driven by uncertainty about future oil supply following extreme geopolitical events are often more important than supply shocks for determining oil price changes. He defines precautionary demand as reflecting the convenience yield from inventory holdings of oil that can serve as insurance against oil supply disruptions. We do not take a position on the relative importance of these two types of influence on oil prices. Instead, we include all the nine events identified, whether they change global oil supply, cause a shift in precautionary demand, or both. For example, the surprising collapse of OPEC in 1986 both increased global oil supply and reduced precautionary demand for oil, as documented by Kilian and Murphy (2014).

3. Data description

Our analysis closely follows DJM, using data starting in October 1973. While the DJM sample ends in April 2003, we update all data series to December 2022. Considering lagged oil price changes in the analysis results in the loss of one observation, yielding a total of 354 observations for the DJM sample and 590 for our extended sample. We use month-end prices for oil and stock market data.

[4] Hamilton (2013) and Baumeister and Kilian (2016) comment on four further periods of fluctuation in oil prices: the East Asian Crisis in 1997–1998, resumed growth in 1999–2000, growing demand and stagnant supply 2003–2008, and tensions in Iran in 2012. We do not include these events in our study as we have not been able to identify comparable geopolitical trigger events during these periods.

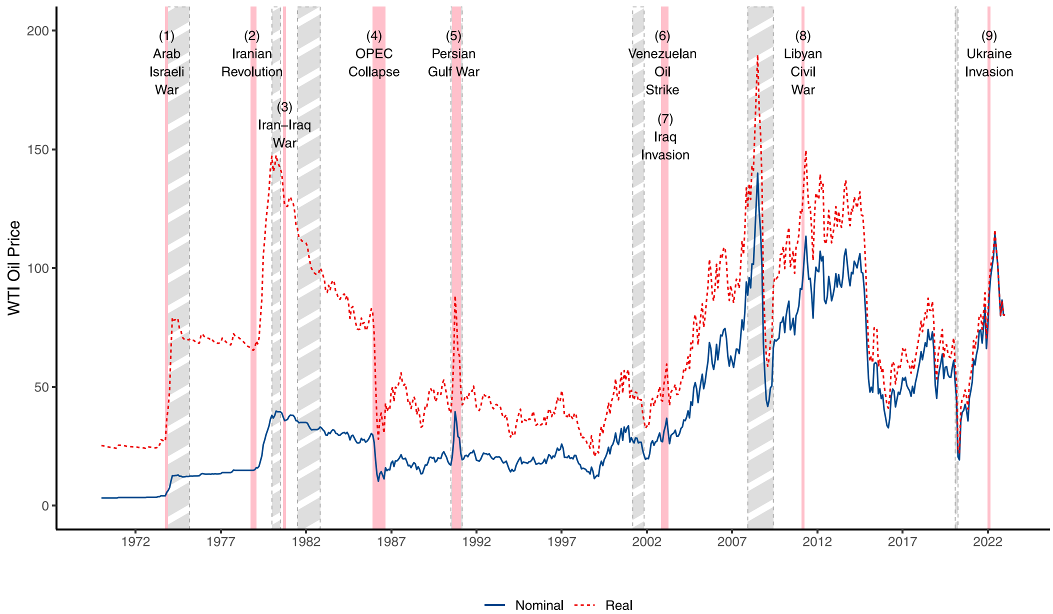

Fig. 1. Oil prices and extreme geopolitical events. The Figure shows the nominal and real oil prices from year-end 1969 to year-end 2022. The dashed line represents the oil price in 2022 dollars and the solid line shows the nominal price. The light red bars show the extreme events listed in Table 1. The gray dashed bars indicate NBER recession dates.

Table 1
Identification of extreme geopolitical events.

Event	Start date	End date	Hamilton (2013)	Baumeister and Kilian (2016)
(1) Arab-Israel war and Arab Embargo	October 1973	November 1973	x	x
(2) Iranian revolution	October 1978	January 1979	x	x
(3) Iran-Iraq War	September 1980	October 1980	x	x
(4) OPEC collapse	December 1985	August 1986	x	x
(5) Persian Gulf War	August 1990	January 1991	x	x
(6) Venezuelan oil strike	November 2002	December 2002	x	x
(7) Invasion of Iraq	January 2003	March 2003	x	x
(8) Libyan civil war	February 2011	March 2011		x
(9) Russian invasion of Ukraine	January 2022	February 2022		

The Table lists the oil price shocks coinciding with extreme geopolitical events identified during our sample period 1973 through to 2022. The Online Appendix offers a detailed description of each of these events.

Stock market returns are computed from total return local currency MSCI indices which are gross of fees and value-weighted. We focus our analysis on the G7 countries: Canada, France, Germany, Italy, Japan, United Kingdom, and United States.[5] Focusing on the G7 countries is common in international asset pricing studies, see, e.g., Cooper and Priestley (2009) and Henkel et al. (2011). International stock markets are often strongly positively correlated, so examining all markets may not add much independent evidence of a statistical relation. These seven countries constitute close to 90% of the market capitalization in the MSCI World index.[6] As of year-end 2022, stocks listed on United States exchanges represent 68.1% of the MSCI World index, Japan 6.2%, United Kingdom 4.4%, France 3.4%, Canada 3.5%, Germany 2.3%, and Italy 0.7%. The data source for the MSCI indices is *Bloomberg*.[7]

DJM use several oil prices in their analysis, but only West Texas Intermediate (WTI) and Arab Light have data available for their entire sample period. Arab Light is produced by Saudi Arabia, and was the main global benchmark in oil markets in the 1960s and 1970s. From the introduction of the NYMEX WTI oil futures in 1983, WTI became the main global benchmark in oil markets (Yergin, 2012). We use WTI denominated in USD/barrel as our measure of the oil price, and report results for Arab Light in the Online Appendix to provide comparability with DJM. Using WTI as the benchmark measure of the oil price is standard in the literature, see, e.g., Alquist and Kilian (2010) and Hamilton (2013). The data source for WTI is *Global Financial Data* and *EIA*, and the data source for Arab Light is *Bloomberg*.[8]

Like DJM, we collect 3-month government bill rates from IMF, obtained from the *FRED* database, for each of our included countries. However, these IMF data series were discontinued at different times for different countries across our extended sample period. We use the 3-month government bill rates collected from *Bloomberg* the month after the IMF series is discontinued. The IMF rates are monthly averages,

[5] DJM include analysis of MSCI indices in 18 countries for their full sample period 1973 to 2003 and report that lagged oil price changes is a significant predictor in 11 of them. They supplement some of their analysis with 30 additional countries that have MSCI index data from 1988 onwards, four of which are significantly predicted by lagged oil price changes.

[6] The MSCI World index covers 23 countries classified as developed markets and captures 85% of each country's free float-adjusted market capitalization (MSCI, 2022).

[7] The Bloomberg codes for the indices we use are GDDLCA, GDDLFR, GDDLGR, GDDLIT, GDDLJN, GDDLUK, GDDLUS, and GDDUWI, each with the Index suffix.

[8] Since January 2, 1986, the Cushing, OK WTI Spot Price FOB (Dollars per Barrel) is available from EIA. From this series, we have extracted the last price in each month to construct a monthly series. Prior to 1986, we have used WTI prices from Global Financial Data. Arab light was downloaded from Bloomberg using the code PGCRARLT Index, which has data covering the entire sample period.

Table 2
Descriptive statistics and stock market correlations.

	Canada (1)	France (2)	Germany (3)	Italy (4)	Japan (5)	UK (6)	US (7)	World (8)	WTI (9)	Arab (10)
Panel A: Descriptive statistics										
Mean	0.72	0.80	0.64	0.62	0.43	0.84	0.83	0.75	0.50	0.56
Std. Dev.	4.70	5.73	5.70	6.67	5.20	5.30	4.53	4.41	9.84	10.78
Min	−24.31	−24.62	−28.67	−25.42	−23.64	−30.02	−23.85	−20.99	−78.20	−69.58
Max	15.34	20.89	19.02	24.22	18.27	43.50	16.37	13.73	61.50	83.79
Panel B: Bivariate correlations										
France	0.60									
Germany	0.52	0.73								
Italy	0.47	0.64	0.59							
Japan	0.42	0.48	0.47	0.46						
UK	0.61	0.64	0.56	0.50	0.41					
US	0.76	0.66	0.63	0.48	0.48	0.65				
World	0.77	0.71	0.68	0.57	0.67	0.70	0.90			
WTI	−0.01	−0.09	−0.11	−0.12	−0.04	−0.10	−0.08	−0.08		
Arab	−0.03	−0.09	−0.12	−0.13	−0.05	−0.08	−0.11	−0.08	0.80	

Columns (1) to (8) in Panel A report descriptive statistics for total return local currency MSCI indices, which are gross of fees and value-weighted, for the G7 countries and the MSCI world index. Columns (9) and (10) in Panel A give descriptive statistics for the lagged change in the log of WTI and Arab Light oil prices, respectively. The data frequency is monthly and prices are end-of-month. The sample period is October 1973 to December 2022 (590 observations). Panel B reports the bivariate correlations.

whereas the Bloomberg rates, in contrast, are month-end.[9] For each country, we use the 3-month government bill rate observed in month $t-1$ as the proxy for the risk-free interest rate available to an investor during month t. The US bill rate is used as a proxy for the risk-free rate for an investor in the MSCI World index.[10]

3.1. Summary statistics

Table 2 offers summary statistics of the data used in our analysis for the full sample period from October 1973 to December 2022. Panel A shows descriptive statistics for the monthly log return for the MSCI indices and oil price changes. The World index has the lowest standard deviation, consistent with the diversification benefits from portfolio theory. Panel A also exhibits that both WTI and Arab Light oil prices are highly volatile with an annualized standard deviation exceeding 30%. Panel B tabulates the bivariate correlation coefficients. As expected, the index returns are strongly positively correlated. The returns on the US index and the World index have a correlation coefficient of 90%, which is not surprising given that the US stock market capitalization makes up two thirds of the World index market capitalization. WTI and Arab Light are less than perfectly correlated, with a correlation coefficient of 80%.

4. Replicating DJM and removing extreme events

We follow DJM, and use the regression

$$r_t = \mu + \gamma o_{t-1} + \varepsilon_t \tag{1}$$

to forecast stock market returns, where r_t is the current period log return on a stock market index and o_{t-1} is the previous period's change in the log oil price. With a monthly data frequency, the stock market return over, e.g., the month of November is compared to the oil price change over the month of October. If γ is significantly different from zero, we reject the null hypothesis of no lagged oil price effect for that particular stock index. DJM only report coefficient estimates from the oil model when using Arab Light to predict stock returns. They report t-statistics based on White (1980) standard errors when using either Arab Light or WTI to predict stock returns. Our main comparison with DJM is therefore in terms of the t-statistics (White standard errors) from the oil model when using WTI to predict stock returns.[11]

Column (1) of Panel A in Table 3 reproduces DJM's t-statistics from the oil model using WTI as the oil price, collected from Table 4 in their published paper. For six of the eight stock market indices, including the World index, they detect a statistically significant oil effect. Columns (2) to (7) present our findings when replicating DJM's results from the oil model. Column (5) reports our estimated t-statistics with White standard errors, closely matching DJM's published results in Column (1).

To assess the impact of the extreme events listed in Table 1, we re-estimate the oil model in a counterfactual scenario in which these extreme events did not occur. To construct this counterfactual scenario, we remove all months relating to these events from both the oil price change and stock return series. After removing these event months, we lag oil price changes and perform the predictive regressions. Treating all extreme events in this way removes 28 observations from the DJM sample period of 354 observations (8% of their sample).[12]

Columns (8) to (13) of Panel A in Table 3 show the results of removing the extreme events. In this counterfactual scenario, lagged oil price changes no longer predict stock market returns. No stock market is significantly predicted at the 5% level with White standard errors, and only the UK stock market is significant predicted at the 5% level with Newey–West standard errors.[13] Importantly, the largest markets,

[9] Codes for IMF rates in *FRED*: INTGSTCAM193N, INTGSTFRM193N, INTGSTDEM193N, INTGSTITM193N, INTGSTJPM193N, INTGSTGBM193N, INTGSTUSM193N. The codes for *Bloomberg* rates: GTCAD3M, GTFRF3M, GTDEM3M, GTITL3M, GTJPY3M, GTGBP3M, GB3, each with the Govt suffix.

[10] Dividing the annualized discount yield by 12 on a 3-month government bill is an approximation, as the exact return realized by an investor is the difference between the 3-month bill's purchase price and the sales price after one month divided by the purchase price. We do not have data which allows us to compute international risk-free returns more accurately, and DJM have not specified how they have computed the risk-free returns.

[11] We also supply Newey and West (1987) standard errors, as they correct for autocorrelation in addition to heteroskedasticity. In each estimation, we use floor$[4(T/100)^{2/9}]$ number of lags, where T is the number of observations, as suggested in Newey and West (1994).

[12] The implication of this procedure is that there will be a disconnect in time between the lagged oil price change and the subsequent stock market return at the points where we remove observations. For example, Event 4 (the OPEC collapse from December 1985 to August 1986) removes December 1985 to August 1986, inclusive. Therefore, the lagged oil price change used to predict the stock return in September 1986 is the oil price change observed during the month of November 1985. An alternative approach would be to first lag the oil price changes and then remove event months. Unreported results show this makes little difference to the results.

[13] The Online Appendix offers similar results when instead using Arab Light as the oil price.

Table 3
Predictive stock market regressions using WTI oil prices.

	DJM	Entire sample						Removing extreme events					
	t_{white} (1)	γ (2)	t_{ols} (3)	t_{nw} (4)	t_{white} (5)	R^2 (%) (6)	n (7)	γ (8)	t_{ols} (9)	t_{nw} (10)	t_{white} (11)	R^2 (%) (12)	n (13)
Panel A: 1973:10–2003:04													
Canada	−1.11	−0.04	−1.15	−1.32	−1.31	0.38	354	0.02	0.46	0.53	0.50	0.06	326
France	−3.36	−0.13	−3.46	−3.01	−3.28	3.30	354	−0.01	−0.20	−0.21	−0.22	0.01	326
Germany	−3.31	−0.16	−4.33	−4.09	−3.43	5.05	354	−0.09	−2.07	−1.86	−1.92	1.31	326
Italy	−4.74	−0.22	−5.17	−4.41	−4.85	7.06	354	−0.10	−1.74	−1.67	−1.94	0.93	326
Japan	−1.21	−0.06	−1.85	−1.75	−1.20	0.97	354	0.04	0.87	1.04	0.87	0.23	326
UK	−3.58	−0.13	−3.61	−4.86	−3.95	3.58	354	−0.08	−1.64	−2.11	−1.79	0.82	326
US	−3.47	−0.10	−3.49	−4.05	−3.48	3.34	354	−0.05	−1.30	−1.49	−1.36	0.52	326
World	−3.08	−0.09	−3.55	−3.53	−2.76	3.45	354	−0.02	−0.53	−0.57	−0.53	0.09	326
Panel B: 2003:05–2022:12													
Canada		0.02	1.01	0.77	0.67	0.44	236	0.02	1.11	0.84	0.73	0.53	232
France		0.02	0.69	0.62	0.61	0.20	236	0.02	0.84	0.78	0.75	0.31	232
Germany		0.01	0.50	0.40	0.37	0.11	236	0.02	0.64	0.52	0.48	0.18	232
Italy		0.03	1.04	1.03	1.04	0.46	236	0.04	1.12	1.14	1.13	0.54	232
Japan		0.02	0.56	0.51	0.56	0.14	236	0.02	0.80	0.72	0.77	0.28	232
UK		0.01	0.45	0.50	0.44	0.08	236	0.01	0.41	0.46	0.41	0.07	232
US		0.01	0.51	0.33	0.31	0.11	236	0.02	0.71	0.46	0.42	0.22	232
World		0.02	0.66	0.48	0.45	0.19	236	0.02	0.82	0.60	0.55	0.29	232
Panel C: 1973:10–2022:12													
Canada		0.00	−0.25	−0.23	−0.23	0.01	590	0.02	1.05	0.98	0.89	0.20	558
France		−0.05	−2.14	−1.52	−1.92	0.77	590	0.01	0.46	0.47	0.48	0.04	558
Germany		−0.06	−2.67	−2.03	−2.14	1.20	590	−0.02	−0.74	−0.62	−0.63	0.10	558
Italy		−0.08	−3.00	−1.88	−2.53	1.51	590	−0.01	−0.23	−0.21	−0.25	0.01	558
Japan		−0.02	−0.88	−0.76	−0.69	0.13	590	0.03	1.15	1.10	1.10	0.24	558
UK		−0.06	−2.55	−2.43	−2.72	1.10	590	−0.02	−0.84	−1.01	−1.02	0.13	558
US		−0.04	−2.04	−1.60	−1.61	0.70	590	0.00	−0.22	−0.17	−0.16	0.01	558
World		−0.03	−1.84	−1.39	−1.39	0.57	590	0.01	0.40	0.31	0.30	0.03	558

The Table reports results from the regression $r_t = \mu + \gamma o_{t-1} + \epsilon_t$, where r_t is the current period return on a stock market index and o_{t-1} is last period's change in the log oil price. Stock market index returns are computed from local currency total return MSCI indices using monthly data. Panel A reports the results from this predictive regression for the DJM sample period, which is October 1973 to April 2003. Panel B reports the results from May 2003 to December 2022. Panel C reports results for the full sample period October 1973 to December 2022. Column (1) reproduces the t-statistics (White standard errors) reported in DJM. Column (2) to (7) report the γ coefficient, corresponding t-statistics (standard OLS, Newey–West, and White) and R^2 when including all observations in the sample period. Columns (8) to (13) report the same when re-running the regressions after removing the observations relating to the extreme events that fall within the different sample periods.

the World and US indices, are far from significantly related to past oil price changes. These results show that the oil puzzle documented by DJM is critically dependent on the occurrence of the extreme events.

Next, we investigate whether some events were instrumental in generating the predictability of stock market returns by lagged oil price changes. Columns (1) through (7) in Panel A in Table 4 give the percent change in t-statistic when removing each extreme event separately during the DJM sample period. Four events each reduce the t-statistic averaged across markets by more than 10%: the 1973 Arab-Israel war, the 1986 OPEC collapse, the 1990/91 Persian gulf war, and the 2003 invasion of Iraq. One potential explanation for the relatively large impact of the three wars is that in addition to a possible second order effect of these events through oil markets to stock markets, these wars may also have had a significant first order effect on stock returns due to the increased geopolitical risk that they induced. For example, Omar et al. (2017) show that periods of extreme geopolitical unrest are associated with a reduction in equity prices and increase in the price of government bonds. As for the 1986 OPEC collapse, the dramatic economic impact of that event has received extensive attention in the literature, see, e.g., Gately (1986) and Lamont (1997).

Column (11) in Panel A in Table 4 shows that excluding only these four major events reduces the t-statistic by 95% on average. In the Online Appendix we reproduce the predictive regressions from Table 3 when the counterfactual scenario is that only these four major events did not occur. No stock market is significantly predicted by lagged oil price changes when we exclude these four events.[14]

[14] Column (10) in Panel A in Table 4 shows the percentage change in t-statistic when removing all events from the estimation. This corresponds to the percentage difference between Columns (4) and (10) in Panel A in Table 3. The average reduction in t-statistic across the G7 and World markets is 89%.

4.1. The oil model in the post-DJM and full sample periods

Columns (2) to (7) of Panel B in Table 3 report the results when we estimate the oil model for the period following the DJM sample period, that is, from May 2003 to December 2022. The estimated slope coefficients are now all positive and insignificantly different from zero, reversing the findings of DJM. Columns (8) to (13) give the results when removing the two extreme events identified in this sample period, which yielded no discernible differences compared to the results with all dates included.

Columns (2) to (7) of Panel C in Table 3 offer parameter estimates for the full sample period from October 1973 to December 2022. The slope coefficients in the German and UK stock markets are significantly different from zero at the 5% level (Newey–West standard errors). However, when we exclude the extreme events, Columns (8) to (13) show that the γ estimates are now all close to zero and none is statistically different from zero.

Columns (1) through (9) in Panel B in Table 4 give the percent change in t-statistic when separately removing each extreme event from the full sample period October 1973 to December 2022. The two events after the DJM sample period ends in 2003, the 2011 Libyan revolution and the 2022 Russian invasion of Ukraine, were less important to the predictive relation than the four events already highlighted in the DJM period.

We further investigate the stability of the slope coefficient estimates from the oil model by calculating coefficient estimates for both an expanding sample window and a rolling five-year window. Fig. 2 plots the estimated slope coefficients using an expanding window as a solid line. The dashed lines indicate 95% confidence intervals based on Newey–West standard errors. The dotted vertical line shows the end of the DJM sample period. For all stock markets (except Italy),

Table 4
Percent change in t-statistic when removing individual events.

	1973: Arab-Israel war (1)	1978/79: Iranian revolution (2)	1980: Iran-Iraq war (3)	**1986: OPEC collapse** (4)	**1990/91: Persian Gulf war** (5)	2002: Venezuela oil strike (6)	**2003: Invasion of Iraq** (7)	2011: Libyan revolution (8)	2022: Russian invasion of Ukraine (9)	Remove all events (10)	Remove four boldfaced events (11)
Panel A: 1973:10–2003:04											
Canada	−32.2	1.2	5.1	−23.9	−63.2	−9.0	−14.1			−139.8	−149.6
France	−12.0	−1.1	−1.3	−37.0	−23.2	−5.1	−17.5			−92.9	−102.8
Germany	−9.1	−0.4	0.5	−20.4	−18.5	−1.7	−16.5			−54.6	−68.1
Italy	−6.7	1.3	−4.4	−17.0	−23.4	−3.4	−8.4			−62.1	−62.9
Japan	−23.1	−1.4	−2.1	−35.4	−51.5	−8.2	−4.8			−159.7	−148.4
UK	−3.4	0.7	−0.8	−23.7	−21.7	−7.2	−10.8			−56.6	−62.9
US	−7.7	−0.7	1.1	−19.8	−25.7	−4.3	−11.3			−63.1	−70.8
World	−12.9	−0.4	−0.2	−34.0	−24.8	−5.2	−14.1			−83.8	−92.4
Average	−13.4	−0.1	−0.3	−26.4	−31.5	−5.5	−12.2			−89.1	−94.7
Panel B: 1973:10–2022:12											
Canada	−121.6	4.2	17.1	−85.3	−231.6	−32.5	−50.1	−10.0	−14.6	−487.9	−489.7
France	−11.5	−1.1	−1.2	−58.4	−20.7	−6.5	−18.4	0.0	−4.2	−125.1	−128.7
Germany	−9.4	−0.4	0.5	−25.3	−16.5	−4.8	−20.8	0.8	−4.0	−70.5	−82.2
Italy	−5.4	0.8	−3.7	−37.7	−21.3	−3.9	−8.8	0.1	−1.9	−90.1	−88.1
Japan	−33.2	−1.7	−2.7	−63.9	−74.1	−11.7	−7.3	−7.1	−11.1	−261.2	−223.5
UK	−9.5	1.1	−0.9	−21.8	−19.1	−8.7	−11.2	0.2	0.1	−62.7	−70.6
US	−14.0	−0.6	1.5	−25.2	−35.2	−6.5	−13.8	0.7	−7.4	−90.1	−90.9
World	−20.8	−0.5	−0.2	−47.8	−27.3	−7.9	−17.4	1.1	−7.4	−121.3	−122.7
Average	−28.2	0.2	1.3	−45.7	−55.7	−10.3	−18.5	−1.8	−6.3	−163.6	−162.0

The Table shows the percent change in the t-statistic (Newey–West) for the γ coefficient in the regression $r_t = \mu + \gamma o_{t-1} + \varepsilon_t$ when we remove each of the nine events in Table 1 separately, and then combinations of events; r_t is the current period return on a stock market index and o_{t-1} is last period's change in the log oil price. Stock market index returns are computed from local currency total return MSCI indices using monthly data. Columns (1) through (9) correspond to results for each of the nine events separately. Column (10) gives results when we remove all events, and Column (11) gives results when we remove the four events (in boldface) with the highest individual reduction in t-statistic. Panel A reports the results for the DJM sample period, which is October 1973 to April 2003. Panel B reports results for the full sample period October 1973 to December 2022.

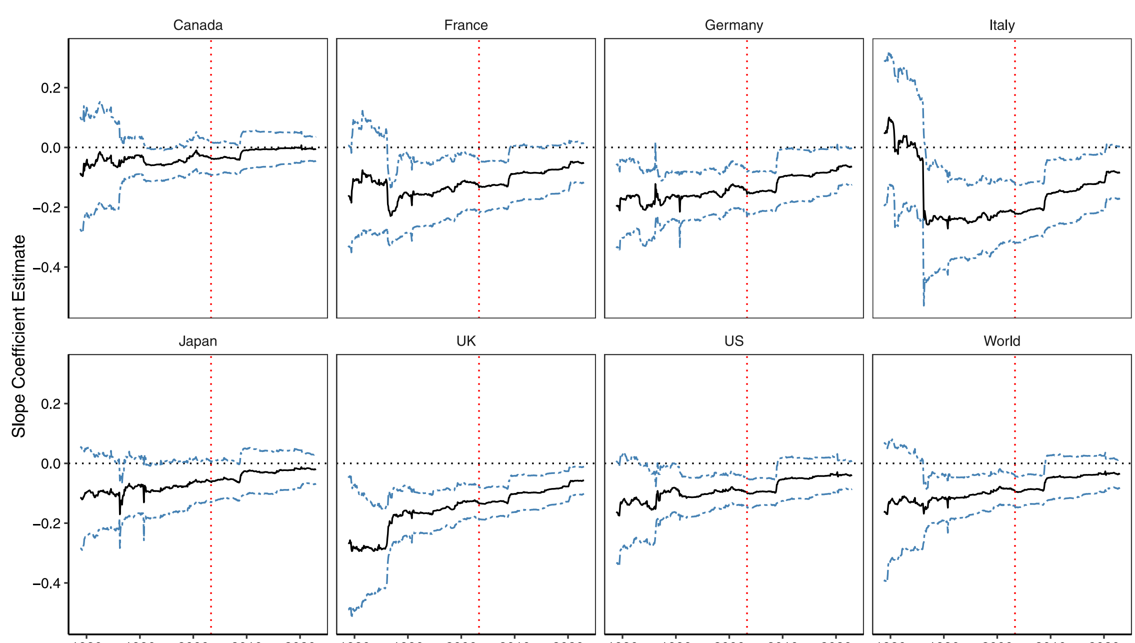

Fig. 2. Oil model slope coefficient estimates based on expanding window. The Figure plots the estimated gamma coefficient from the regression $r_t = \mu + \gamma o_{t-1} + \varepsilon_t$ based on an expanding estimation window. r_t is the current period return on a stock market index and o_{t-1} is last period's change in oil price. Stock market index returns are computed from local currency total return MSCI indices using monthly data. The dashed lines show the 95% confidence interval using Newey–West standard errors. The dotted vertical line indicates the end of the DJM sample period.

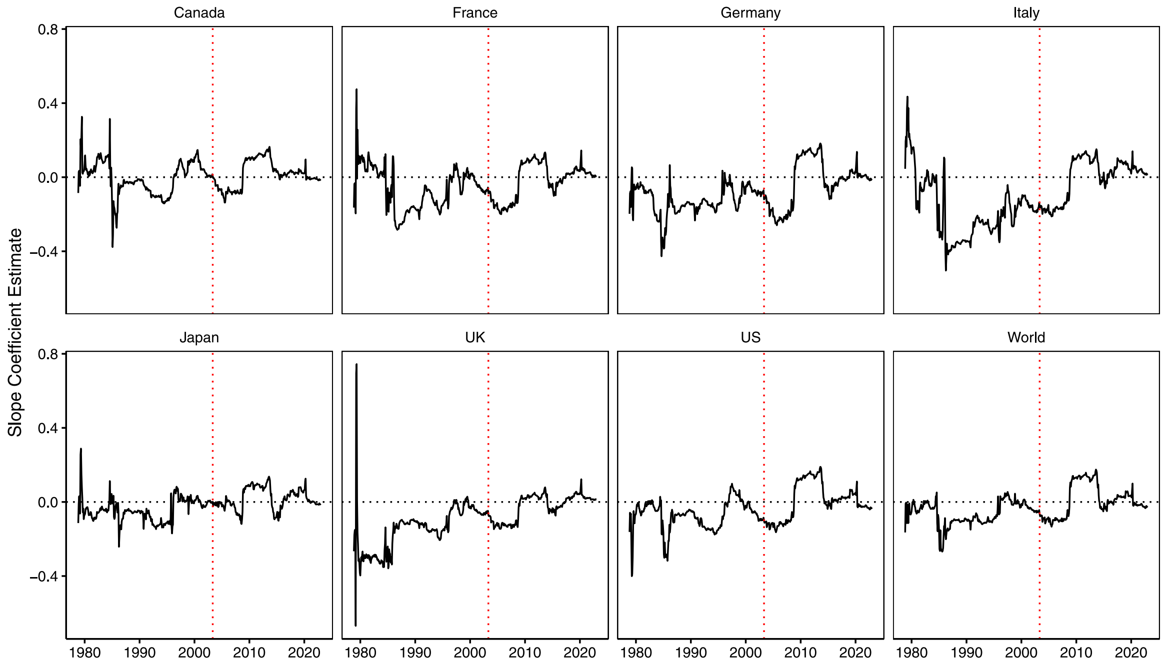

Fig. 3. Oil model slope coefficient estimates based on rolling window. The Figure plots the estimated gamma coefficient from the regression $r_t = \mu + \gamma o_{t-1} + \varepsilon_t$ based on overlapping five-year estimation windows. r_t is the current period return on a stock market index and o_{t-1} is last period's change in oil price. Stock market index returns are computed from local currency total return MSCI indices using monthly data. The dotted vertical line indicates the end of the DJM sample period.

the coefficient estimate is below zero at the start of the sample, and then gradually increases as we add successive months of data to the estimation. Towards the end of the sample, all confidence intervals include or are very close to zero. Fig. 3 plots the coefficient estimates using five-year rolling windows. The estimated slope coefficients are highly volatile and frequently changing sign.

This parameter instability further suggests that the oil model provides unreliable signals for an investment strategy. Such lack of robustness to variations in time period reduces a predictor's practical usefulness, as highlighted by Hsu et al. (2019).

5. Assessing the oil model's prediction performance

We investigate the prediction performance of the oil model using the graphical tool proposed by Goyal and Welch (2003, 2008). This graphical tool allows visual detection of any periods that abnormally contribute to the performance of a prediction model. An upward-sloping graph indicates that the oil model predicts the stock market return better than a forecast based simply on the mean market return. The Goyal–Welch graphs are thus well suited to assess the performance of the oil model during the events listed in Table 1. Goyal and Welch (2008) advice that a well-specified signal needs a reasonably good in-sample (IS) and out-of-sample (OOS) performance over the entire sample period. This requires a generally upwards drift in performance that is not confined to short or unusual sample periods.

First, we present the oil model's IS performance, which uses the full sample to determine parameter values. IS performance is measured as the cumulative difference between the sum of squared errors (SSE) using the unconditional mean of stock returns and the oil model to predict next period's stock return, i.e.

$$IS(\tau) = \sum_{t=\text{November 1973}}^{\tau} SSE(t)^{\text{Mean Full Sample}} - SSE(t)^{\text{Oil model Full Sample}}, \tag{2}$$

where τ = November 1973, … , December 2022.

Fig. 4 plots the oil model's IS performance for each of the eight stock indices. The solid line shows the IS performance for all dates, whereas the dotted line shows the performance with the extreme events removed. Removing the events substantially detracts from the predictive performance of the oil model, rendering the final data point in these line plots very close to zero in all markets. There are two main reasons for this. The first is that the Arab-Israel war and Arab embargo in 1973 (Event 1) is key to establishing the strong negative relation between lagged oil prices and stock returns. Removing that event, attenuates the negative γ in the oil model. Second, as is clear from the solid lines that show all dates, the SSE difference increases sharply during most of the events. This suggests that the predictability is not a steady one where most months contribute, but rather one where a small number of extreme events create practically all the predictability.

Next, we turn to the oil model's OOS performance, where each period's forecast is based only on data available at the time. We use 20 observations before making the initial OOS forecast in May 1975. Subsequently, both models are updated each month with an additional observation to predict next month's stock market return

$$OOS(\tau) = \sum_{t=\text{May 1975}}^{\tau} SSE(t)^{\text{Mean Prevailing Sample}} - SSE(t)^{\text{Oil model Prevailing Sample}}, \tag{3}$$

where τ = May 1975, … , December 2022. The important distinction between IS and OOS performance is that OOS uses the parameter estimates available to an investor who cannot look into the future but must use prevailing information to estimate the model.

Fig. 5 plots the OOS performance for each of the eight stock indices. The solid line shows the OOS performance across the entire sample, whereas the dotted line shows the OOS performance when we remove the extreme events. A first observation is that the OOS performance of the oil model is considerably worse than the IS performance in Fig. 4, both before and after removing the events. OOS results for the oil model are worse than the mean model as of year-end 2022 in all markets without removing any dates, except in Italy. Furthermore, OOS performance is strongly dependent on the events. When we remove these events the SSE difference in all markets end below zero. This shows that the unconditional mean model outperforms the oil model in terms of prediction error across all the stock markets.

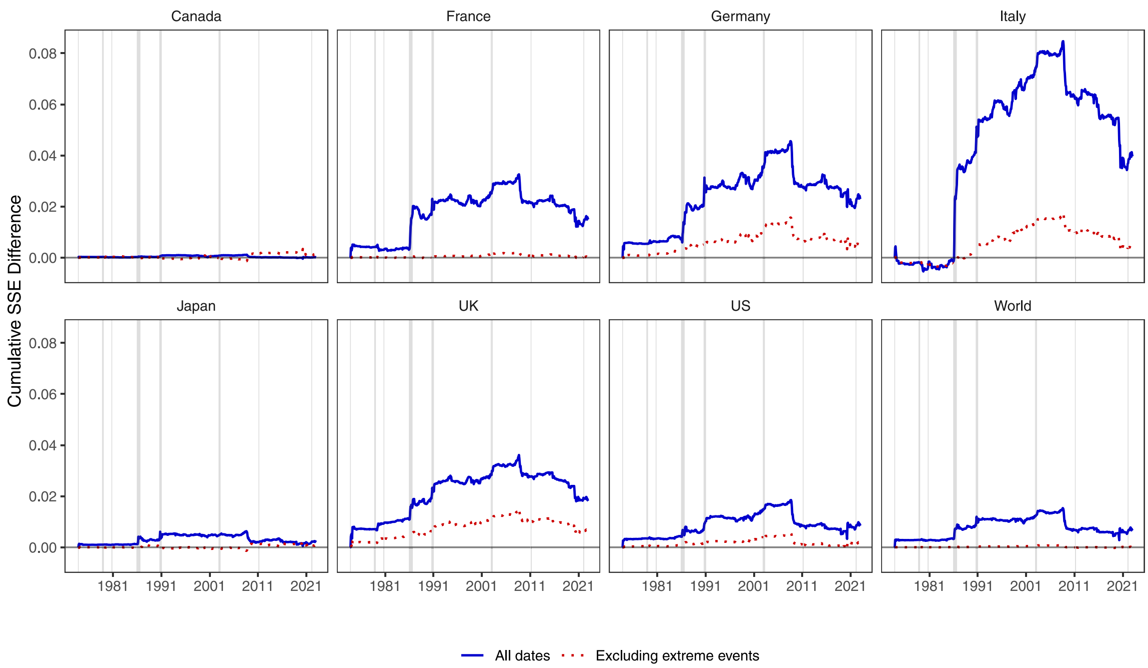

Fig. 4. Goyal–Welch In-Sample (IS) diagnostic test. The Figure plots In-Sample (IS) sum of squared errors (SSE) of the oil model relative to the simple unconditional mean stock return, as proposed by Goyal and Welch (2008)

$$IS(\tau) = \sum_{t=\text{November 1973}}^{\tau} SSE(t)^{\text{Mean Full Sample}} - SSE(t)^{\text{Oil model Full Sample}},$$

where τ = November 1973, …, December 2022. The solid line includes all dates. The dotted line excludes the events listed in Table 1, and the gray bars correspond to the same events.

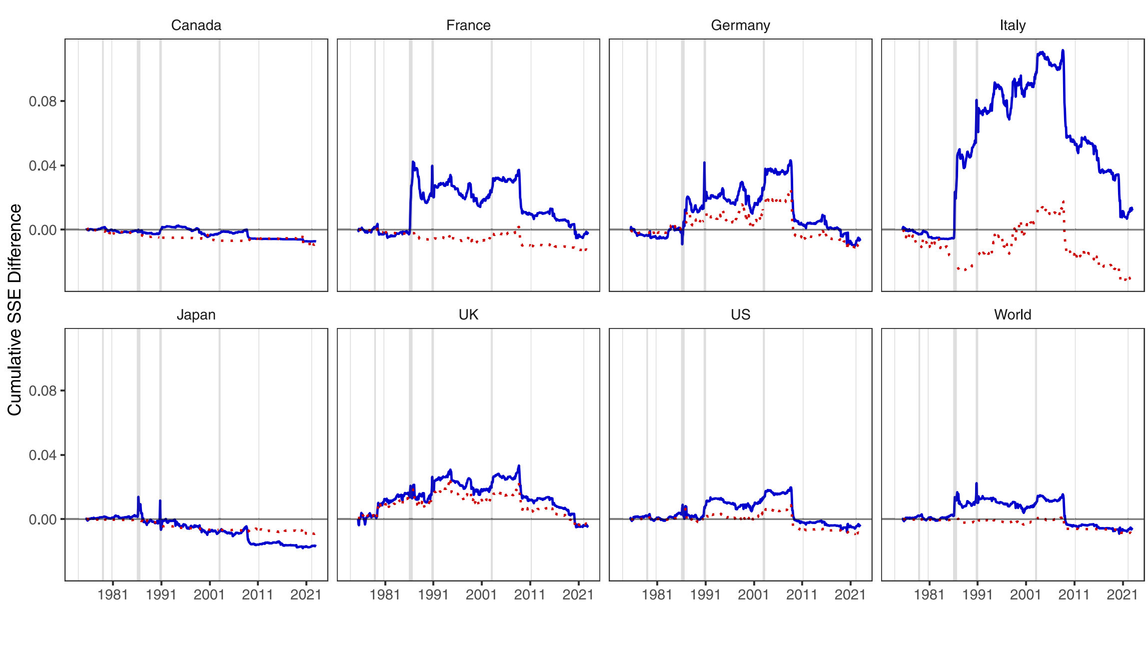

Fig. 5. Goyal–Welch Out-Of-Sample (OOS) diagnostic test. The Figure plots Out-Of-Sample (OOS) sum of squared errors (SSE) of the oil model relative to the simple unconditional mean stock return, as proposed by Goyal and Welch (2003, 2008)

$$OOS(\tau) = \sum_{t=\text{May 1975}}^{\tau} SSE(t)^{\text{Mean Prevailing Sample}} - SSE(t)^{\text{Oil model Prevailing Sample}},$$

where τ = May 1975, …, December 2022. The solid line includes all dates. The dotted line excludes the events listed in Table 1, and the gray bars correspond to the same events.

Table 5
Market-timing trading strategy results using WTI oil prices.

	Buy and hold strategy				Oil strategy								
	$\bar{r}_m$ (1)	σ_m (2)	SR_m (3)	$\bar{r}_f$ (4)	$\bar{r}_s$ (5)	σ_s (6)	SR_s (7)	r_e^{cum} (8)	α (9)	t_α (10)	β (11)	t_β (12)	λ_m (13)
Panel A: 1988:08–2003:04													
Canada	7.90	15.73	0.11	6.24	8.06	14.12	0.13	2.32	0.49	0.30	0.80	14.59	0.75
France	8.43	20.08	0.11	6.13	9.90	15.92	0.24	21.75	2.33	0.91	0.63	9.60	0.69
Germany	6.69	23.16	0.07	4.96	9.72	16.52	0.29	44.66	3.87	1.27	0.51	5.89	0.63
Italy	7.17	23.61	−0.05	8.34	14.65	16.85	0.37	110.36	6.92	2.28	0.52	7.82	0.55
Japan	−6.06	20.28	−0.38	1.65	−5.98	18.07	−0.42	1.08	−1.54	−0.78	0.79	9.99	0.88
UK	8.23	15.61	0.05	7.39	12.16	12.08	0.39	57.84	4.26	2.17	0.60	8.93	0.71
US	10.72	15.09	0.39	4.86	9.19	13.28	0.33	−22.59	−0.19	−0.11	0.77	14.48	0.78
World	5.97	15.09	0.07	4.86	6.44	13.18	0.12	6.90	0.73	0.43	0.76	12.76	0.77
Panel B: 2003:05–2022:12													
Canada	8.19	13.42	0.50	1.51	8.07	13.42	0.49	−2.24	−0.11	−1.00	1.00	926.64	1.00
France	7.58	16.29	0.42	0.71	6.99	15.03	0.42	−11.65	0.45	0.31	0.85	15.55	0.85
Germany	7.51	18.28	0.38	0.61	4.20	16.26	0.22	−65.13	−1.85	−1.05	0.79	13.83	0.78
Italy	3.43	20.02	0.11	1.24	0.92	17.14	−0.02	−49.33	−1.92	−0.96	0.73	12.82	0.66
Japan	6.39	17.53	0.36	0.04	5.36	16.65	0.32	−20.20	−0.39	−0.34	0.90	27.02	0.92
UK	6.91	12.94	0.41	1.65	6.85	12.05	0.43	−1.19	0.66	0.60	0.86	20.36	0.83
US	9.29	15.00	0.54	1.22	9.16	14.60	0.54	−2.48	0.31	0.35	0.95	41.45	0.88
World	8.36	15.53	0.46	1.22	8.11	15.06	0.46	−4.99	0.18	0.19	0.94	39.25	0.89
Panel C: 1988:08–2022:12													
Canada	8.07	14.44	0.31	3.54	8.07	13.71	0.33	0.08	0.47	0.61	0.90	32.71	0.89
France	7.94	17.99	0.27	3.04	8.24	15.40	0.34	10.10	1.61	1.12	0.73	16.17	0.78
Germany	7.16	20.48	0.23	2.48	6.56	16.37	0.25	−20.47	1.11	0.61	0.64	10.54	0.71
Italy	5.03	21.62	0.03	4.28	6.80	17.11	0.15	61.03	2.06	1.15	0.62	13.25	0.62
Japan	1.06	18.82	0.02	0.73	0.50	17.32	−0.01	−19.12	−0.51	−0.43	0.85	18.73	0.90
UK	7.48	14.13	0.24	4.11	9.13	12.07	0.42	56.65	2.58	2.30	0.72	16.31	0.78
US	9.90	15.02	0.47	2.78	9.18	14.03	0.46	−25.06	0.19	0.20	0.87	31.92	0.84
World	7.34	15.33	0.30	2.78	7.39	14.27	0.32	1.91	0.67	0.70	0.87	28.43	0.84

The Table compares the performance of DJM's market-timing strategy based on the oil model (WTI oil price) to that of a simple buy-and-hold strategy. The first oil model prediction is in August 1988, and we continue to re-estimate the model every month over the sample period. If the expected return based on the oil model is higher than the risk-free rate, then we fully invest in the stock market; if it is lower then we invest in short-term government bills. This produces a strategy return series $r_{s,t}$ for each stock market index. Columns (1) to (4) report the average annualized market return, the corresponding standard deviation, the Sharpe ratio, and the average annualized three-month bill rate, respectively. Columns (5) to (13) report results for the oil strategy. Column (5) is the average annualized return obtained by following the oil strategy; Column (6) is the volatility, Column (7) is the Sharpe ratio, Column (8) is the cumulative return from the oil strategy in excess of the market, Column (9) is the alpha from the regression $r_{s,t} - r_{f,t-1} = \alpha + \beta(r_{m,t} - r_{f,t-1}) + \varepsilon_t$; Column (10) is the t-statistic for alpha with White standard errors; Column (11) is the beta from the same regression; Column (12) is the t-statistic for beta with White standard errors; and Column (13) is the ratio of months invested in the stock market (λ_m). All returns are reported in percent.

6. Trading strategy based on oil price changes

DJM suggest using the oil model as a signal for a market-timing investment strategy. Their proposed trading strategy only invests in the stock market if the oil model predicts a stock market return that exceeds the risk-free rate. Otherwise, the strategy invests in government bills.

Following DJM, we compare the return on the oil model-based market-timing strategy with a simple buy and hold strategy. We first estimate the oil model using data from October 1973 to July 1988. These estimates and the last oil price change are then used to predict stock market returns in August 1988. Next, we re-estimate the model every month updating our model with the last month of observations. We then determine whether the predicted stock market return is higher than the current risk-free rate, i.e., if $E_{t-1}(r_t) = \mu + \gamma o_{t-1} \geq r_{f,t-1}$, where r_t denotes a stock market return and $r_{f,t-1}$ denotes the same country's short-term bill rate (observed at the end of month $t-1$).

This procedure generates a series of oil strategy returns $r_{s,t}$ for each of our eight stock market indices

$$r_{s,t} = \begin{cases} r_{m,t} & \text{if } \mu + \gamma o_{t-1} \geq r_{f,t-1} \\ r_{f,t-1} & \text{otherwise.} \end{cases} \quad (4)$$

We do not deduct trading costs, unlike DJM who deduct 0.10% based on using index futures contracts to implement this market-timing strategy. Deducting trading costs would only subtract further from the strategy's performance.

As Eq. (4) makes clear, the strategy's return deviates from the market return only when it predicts an investment below the risk-free rate. In those cases, the risk-free rate, usually a small positive return, is obtained rather than the market return. Therefore, the strategy's potential for outperformance depends on its ability to avoid market losses. In addition to considering the raw returns, we evaluate the trading strategy's performance by regressing $r_{s,t}$ on its corresponding stock market return $r_{m,t}$ using the following regression

$$r_{s,t} - r_{f,t-1} = \alpha + \beta(r_{m,t} - r_{f,t-1}) + \varepsilon_t. \quad (5)$$

Panel A of Table 5 shows trading strategy results from the DJM sample period. The average return on the oil strategy (Column 5) is higher than the average return on the buy and hold strategy (Column 1) in all markets except the US. Comparing columns (3) and (7), the oil strategy yields higher Sharpe ratios than buy and hold except in Japan and the US.[15] However, Columns (9) and (10) show that the oil strategy yields insignificant alphas for all stock markets, except Italy and UK.[16]

In contrast to our replication of their results, DJM find that the oil strategy generates significant positive alpha estimates in the majority of indices they investigate. One reason for this discrepancy could be that DJM report market beta estimates of 0.5 or less for all indices, which is far below our estimates. We would expect the market beta to be of a similar magnitude to the fraction of time that the oil strategy is invested in the stock market.

Pursuing this idea, we use the US market to offer a graphical illustration. DJM propose that we invest in the stock market only if

[15] Japan exhibits a negative average return because a large part of this estimation period is during Japan's "lost decade" (from about 1991 to 2001).

[16] We follow DJM and report t-statistics based on White standard errors from these regressions; Newey–West standard errors leave our conclusion unchanged.

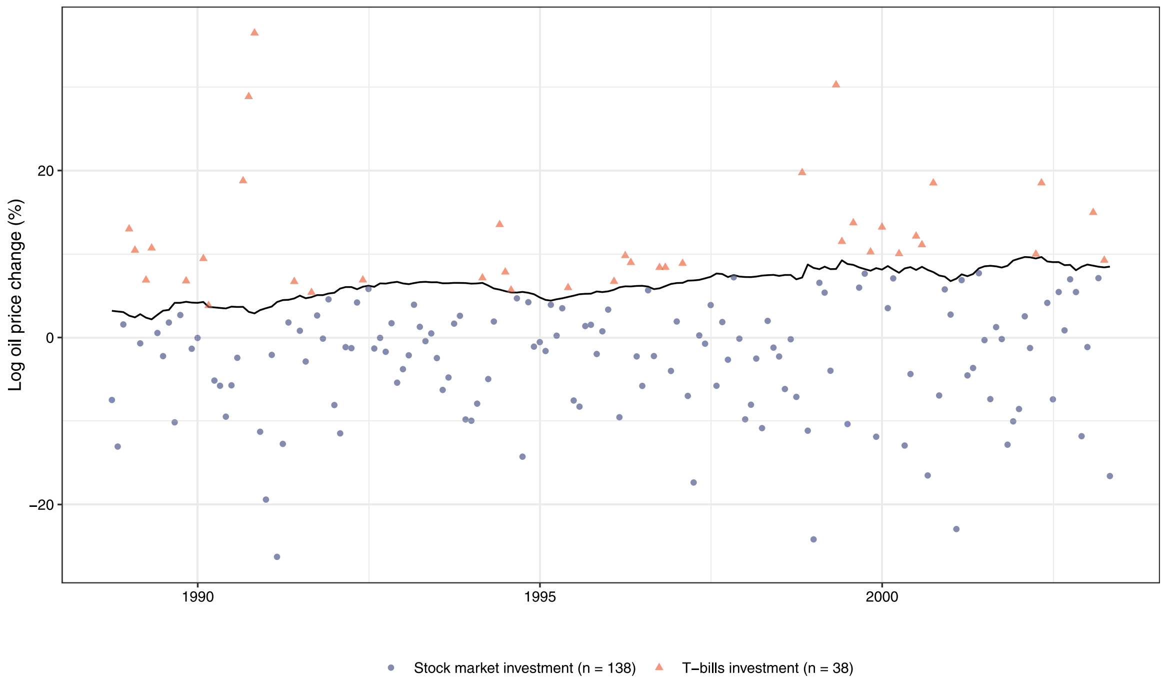

Fig. 6. Investment resulting from trading strategy in the United States. The Figure plots the investment type suggested by the oil model during the DJM sample period. The trading strategy suggests investing in the stock market only if the oil model predicts a next period stock return greater than the prevailing risk-free rate. The solid line marks the boundary for log oil price changes: Above this boundary the investment is government bills and below this boundary the investment is the stock market. The circles below the solid line show oil price changes that resulted in a stock market investment. The triangles above the line show oil price changes that resulted in an investment in government bills. The estimation period is August 1988 to April 2003.

the oil model predicts a return exceeding the prevailing risk-free rate, i.e., if $\mu_{t-1} + \gamma_{t-1} o_{t-1} \geq r_{f,t-1}$. Rearranging and skipping subscripts on coefficient estimates, the strategy invests in government bills whenever $o_{t-1} \geq (r_{f,t-1} - \mu)/\gamma$. Fig. 6 plots the case for the US market in DJM's sample period, where the solid line marks the boundary $(r_{f,t-1} - \mu)/\gamma$. In this period, the oil price change was below this boundary level in 138 of the 177 months from August 1988 to April 2003, or 78% of the months. The US oil model thereby signals a stock market investment in more than three quarters of the sample period, which is hard to reconcile with DJM's reported beta estimate for the US of 0.43. Column (13) in Table 5 reports the ratio of months invested in the stock market for all indices. For the DJM sample period (Panel A), the average ratio is 72%.

These results relate to the requirement put forth by Schwert (2003) for a significant stock market forecasting regression to be considered evidence of market inefficiency. Because expected excess return cannot be negative in a rational market with risk-averse investors, Schwert (2003) proposes that a model of market inefficiency must be able to predict negative excess return. DJM argue that their oil model meets this standard; the model includes forecasts of negative excess returns. While this holds true, our results show that the oil model predicts negative excess returns much less frequently than DJM's reported results indicate.

Panel B of Table 5 reports results for the oil model strategy from the end of the DJM sample in 2003 through to 2022. The average return on the oil strategy is now lower than the average return on the buy and hold strategy in all markets. Furthermore, none of the alphas is significantly different from zero. Panel C of Table 5 reports results for the full sample period, in which results are mixed across markets and only the UK shows a significant alpha.

DJM use Arab Light as the oil price when implementing their market-timing strategy. The Online Appendix offers results from the same estimations when using Arab Light as the oil price. The results, in terms of significant alpha, are if anything worse than those for WTI that we report in Table 5. The Online Appendix plots the excess returns from the trading strategy for each market, both when using WTI and Arab Light, and shows that the excess returns obtained by an investor differ significantly depending on the choice of oil price. However, the predictive regressions offer little guidance with respect to which oil price series an investor should use as the signal.

7. Conclusion

Oil price changes are closely followed and easily accessible. The existence of a worldwide stock market mispricing simply based on oil price changes would therefore pose a severe challenge to the view that markets are informationally efficient. DJM document that oil price changes predict stock market returns worldwide, and they argue that this predictability constitutes evidence of market inefficiency. We show that their results are largely driven by well-known oil price shocks following extreme geopolitical events, primarily four events in their sample period: the 1973 Arab-Israel war, the 1986 OPEC collapse, the 1990/91 Persian gulf war, and the 2003 invasion of Iraq. In the period after DJM's sample ends in 2003 and up to 2022, a period with few similarly disruptive geopolitical events, the forecasting relation is insignificant, and all coefficients change sign. We argue that the attenuation of the predictive relationship from oil price changes to stock returns after DJM's sample period ends in 2003 is due to a lack of sufficiently disruptive geopolitical events after 2003, rather than investors learning about a profitable oil-related market inefficiency from the academic literature. Moreover, we document that a market-timing trading strategy based on oil price changes typically generates insignificant abnormal returns compared to a simple buy and hold strategy both before and after 2003. We conclude that the oil puzzle documented by DJM is limited to a very small number of extreme geopolitical events during their sample period.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.eneco.2024.107659.

References

Alquist, R., Kilian, L., 2010. What do we learn from the price of crude oil futures? J. Appl. Econ. 25, 539–573.

Baumeister, C., Kilian, L., 2016. Forty years of oil price fluctuations: Why the price of oil may still surprise us. J. Econ. Perspect. 30, 139–160.

Caldara, D., Iacoviello, M., 2022. Measuring geopolitical risk. Amer. Econ. Rev. 112, 1194–1225.

Cederburg, S., Johnson, T.L., O'Doherty, M.S., 2023. On the economic significance of stock return predictability. Rev. Finance 27, 619–657.

Cooper, I., Priestley, R., 2009. Time-varying risk premiums and the output gap. Rev. Financ. Stud. 22, 2801–2833.

Cunado, J., Gupta, R., Lau, C.K.M., Sheng, X., 2020. Time-varying impact of geopolitical risks on oil prices. Def. Peace Econ. 31, 692–706.

Dong, X., Li, Y., Rapach, D.E., Zhou, G., 2022. Anomalies and the expected market return. J. Finance 77, 639–681.

Driesprong, G., Jacobsen, B., Maat, B., 2008. Striking oil: Another puzzle? J. Financ. Econ. 89 (2), 307–327.

Feng, G., Giglio, S., Xiu, D., 2020. Taming the factor zoo: A test of new factors. J. Finance 75, 1327–1370.

Gately, D., 1986. Lessons from the 1986 Oil Price Collapse. Brook. Pap. Econ. Act. 237–284.

Goyal, A., Welch, I., 2003. Predicting the equity premium with dividend ratios. Manage. Sci. 49, 639–654.

Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of equity premium prediction. Rev. Financ. Stud. 21 (4), 1455–1508.

Goyal, A., Welch, I., Zafirov, A., 2023. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction Ii. Swiss Finance Institute Research Paper.

Hamilton, J.D., 2003. What is an oil shock? J. Econometrics 113 (2), 363–398.

Hamilton, J.D., 2013. Historical oil shocks. In: Parker, R.E., Whaples, R. (Eds.), Routledge Handb. Major Events Econ. Hist. 239–265.

Harvey, C.R., Liu, Y., Zhu, H., 2016. And the cross-section of expected returns. Rev. Financ. Stud. 29, 5–68.

Henkel, S.J., Martin, J.S., Nadari, F., 2011. Time-varying short-horizon predictability. J. Financ. Econ. 99, 560–580.

Hong, H., Stein, J.C., 1999. A unified theory of underreaction, momentum trading and overreaction in asset markets. J. Finance 54 (6), 2143–2184.

Hou, K., Xue, C., Zhang, L., 2020. Replicating anomalies. Rev. Financ. Stud. 33, 2019–2133.

Hsu, J., Kalesnik, V., Kose, E., 2019. What is quality? Financ. Anal. J. 75, 44–61.

Jiang, H., Skoulakis, G., Xue, J., 2018. Oil and Equity Return Predictability: The Importance of Dissecting Oil Price Changes. Robert H. Smith School Research Paper No. RHS 2822061.

Jones, C.M., Kaul, G., 1996. Oil and the stock markets. J. Finance 51, 463–491.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. Amer. Econ. Rev. 99 (3), 1053–1069.

Kilian, L., Murphy, D.P., 2014. The role of inventories and speculative trading in the global market for crude oil. J. Appl. Econometrics 29, 454–478.

Kilian, L., Park, C., 2009. The impact of oil price shocks on the U.S. stock market. Internat. Econom. Rev. 50, 1267–1287.

Kim, J.H., Rahman, L.M., Shamsuddin, A., 2019. Can energy prices predict stock returns? An extreme bounds analysis. Energy Econ. 81, 822–834.

Lamont, O., 1997. Cash flow and investment: Evidence from internal capital markets. J. Finance 52 (1), 83–109.

Lochstoer, L.A., Tetlock, P., 2020. What drives anomaly returns? J. Finance 75, 1417–1455.

McLean, D.R., Pontiff, J., 2016. Does academic research destroy stock return predictability? J. Finance LXXI, 5–31.

MSCI, 2022. MSCI world index fact sheet. January 31.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55 (3), 703–708.

Newey, W.K., West, K.D., 1994. Automatic lag selection in covariance estimation. Rev. Econ. Stud. 61, 631–654.

Omar, A.M.A., Wisniewski, T.P., Nolte, S., 2017. Diversifying away the risk of war and cross-border political crisis. Energy Econ. 64, 494–510.

Rapach, D.E., Zhou, G., 2013. Forecasting stock returns. In: Elliott, G., Timmermann, A. (Eds.), Handb. Econ. Forecast. 2A, 327–383.

Schwert, G., 2003. Anomalies and market efficiency. In: Constantinides, G.M., Harris, M., Stulz, R. (Eds.), Handb. Econ. Finance 939–972.

Smyth, R., Narayan, P.K., 2018. What do we know about oil prices and stock returns? Int. Rev. Financ. Anal. 57, 148–156.

White, H., 1980. A heteroskedasticity-consistent covariance estimator and direct test for heteroskedasticity. Econometrica 48, 817–838.

Wolfers, J., Zitzewitz, E., 2009. Using markets to inform policy: The case of the Iraq War. Economica 76, 225–250.

Yergin, D., 2012. The Prize. Simon & Schuster UK Ltd.

Zhang, Q., Yang, K., Hu, Y., Jiao, J., Wang, S., 2023. Unveiling the impact of geopolitical conflict on oil prices: A case study of the Russia-Ukraine War and its channels. Energy Econ. 126.