Oil and Equity Index Return Predictability: The Importance of Dissecting Oil Price Changes*

Haibo Jiang†
Tulane University

Georgios Skoulakis‡
University of British Columbia

Jinming Xue§
University of Maryland

January 12, 2018

* The authors wish to thank Steve Baker, Murray Carlson, Lorenzo Garlappi, Xuhui (Nick) Pan, Nishad Kapadia, Carolin Pflueger, Lena Pikulina, Terry Zhang, Zhaodong (Ken) Zhong, and seminar participants at the Georgia Institute of Technology, at Tulane University, at the University of Auckland, at the University of British Columbia, the 2016 Conference on Research on Economic Theory and Econometrics, the 2017 Pacific Northwest Finance Conference, the CICF 2017 conference, and the MFA 2017 conference for their comments and suggestions. Tudor Barcan, Stacy Cheng, Beth Li, and Judah Okwuobi provided excellent research assistance. † Finance Department, Freeman School of Business, Tulane University, New Orleans, LA, 70118; Email: hjiang9@tulane.edu. ‡ Finance Department, Sauder School of Business, University of British Columbia, Vancouver, BC, Canada V6T 1Z2; Email: georgios.skoulakis@sauder.ubc.ca. § Finance Department, Smith School of Business, University of Maryland, College Park, MD 20742; Email: jinming.xue@rhsmith.umd.edu.

Oil and Equity Index Return Predictability: The Importance of Dissecting Oil Price Changes

Abstract

Using data until 2015, we document that oil price changes no longer predict G7 country equity index returns, in contrast to evidence based on earlier samples. Using a structural VAR approach, we decompose oil price changes into oil supply shocks, global demand shocks, and oil-specific demand shocks. The conjecture that oil supply shocks and oil-specific demand shocks (global demand shocks) predict equity returns with a negative (positive) slope is supported by the empirical evidence over the 1986-2015 period. The results are statistically and economically significant and do not appear to be consistent with time-varying risk premia.

JEL Classification Codes: C53; G10; G12; G14; E44; Q41

Keywords: Equity return predictability; structural VAR model; oil price change decomposition; oil supply shock; global demand shock; oil-specific demand shock

1 Introduction

The impact of oil price fluctuations on equity markets and the real economy has been of great interest to academics, policy makers, and market participants alike. Oil, as the major source of energy, plays a crucial role in the modern global economy. Oil price changes could be interpreted in different ways. On one hand, an oil price increase could be considered bad news for the economy and equity markets as it increases the cost of production in a significant number of sectors and causes consumers to reduce their consumption. Following the same line of thinking, an oil price drop would have the opposite effect and would be perceived as good news. On the other hand, lower oil prices imply lower profits for the oil sector. This will likely cause oil company shares to lose value and, to some extent, drag down the aggregate market. Analogously, higher oil prices is good news for the oil sector and could affect positively the broader market. The conventional wisdom in the past was that the former effect dominates implying that an oil price hike is considered to be bad news for equity markets. A Financial Times 2008 article was titled "US stocks rally as oil prices fall".[1] Accordingly, a possible conjecture is that positive (negative) oil price changes should predict lower (higher) subsequent stock returns. In line with this conjecture, Driesprong, Jacobsen, and Maat (2008) document that, based on data until 2003, oil price changes predict Morgan Stanley Capital International (MSCI) equity index returns with a negative and statistically significant predictive slope for a large number of countries. However, the relationship between oil price movements and stock returns evolves over time.[2] Indeed, the correlation between oil price changes and subsequent equity index returns has turned positive over the last ten years. Figure 1, where we present the two scatter plots of MSCI World index return versus the one-month lagged log growth rate of West Texas Intermediate (WTI) spot price over the 1982–2003 and 2004–2015 periods, clearly illustrates the shift over time. As a result, the predictive ability of oil price change has been dramatically reduced over the sample period covering the last thirty years. This structural change is striking and begs for an explanation.

[1] See *Financial Times*, August 8, 2008 (https://www.ft.com/content/59891010-6545-11dd-a352-0000779fd18c). [2] The financial press closely follows the dynamic relationship between oil and the stock market. Two recent articles, related to this point, in the *Wall Street Journal* are titled "Oil, stocks at tightest correlation in 26 years" on January 25, 2016 (https://www.wsj.com/articles/oil-stocks-dance-the-bear-market-tango-1453722783) and "Stocks and oil prices: correlation breakdown" on April 18, 2016 (https://www.wsj.com/articles/stocks-and-oil-prices-correlation-breakdown-1461003126). It is worth pointing out that, while most of the discussion in the financial press is concerned with contemporaneous correlations, we focus on the predictive relationship between oil price changes and future stock returns.

At the most fundamental level, oil prices move when there is a misalignment between supply and demand. Understanding what causes the oil price change in the first place can be crucial for determining the potential impact of such a change on equity markets. For instance, lower oil prices due to a slowdown in global economic activity should be viewed as bad news. As Ethan Harris, a Bank of America chief economist, put it "*If you think oil prices are dropping because of the global economy is sick, then you are less likely to see lower prices as a windfall*".[3] But prices could also fall because of excess supply of oil, in which case the message would be different. To provide an explanation for the recent positive correlation between oil price changes and aggregate US market returns, Bernanke (2016) decomposes oil price change into a demand-related component and a residual. He documents that the correlations of the two components with market returns are different and states about his finding: "*That's consistent with the idea that when stock traders respond to a change in oil prices, they do so not necessarily because the oil movement is consequential in itself, but because fluctuations in oil prices serve as indicators of underlying global demand and growth*". Recent academic literature has also discussed the potential differential effects of demand and supply shocks associated with oil price fluctuations. Although oil price shocks were often associated with oil production disruptions in the 1970s and 1980s, it has been argued that the role of global demand for oil, especially from fast-growing emerging economies, should also be emphasized (Hamilton, 2003; Kilian, 2009; Kilian and Park, 2009). Furthermore, Kilian (2009) points out that oil supply shocks, global demand shocks, and other types of shocks, all of which can cause oil prices to fluctuate, should have different effects on the macroeconomy and the stock market. Building on these ideas, we suggest that oil price changes do contain useful information for predicting future stock returns in a real-time fashion, once these changes are suitably decomposed into supply and demand shocks.

We start our empirical analysis by demonstrating that the ability of oil price changes to forecast G7 country MSCI index excess returns at the monthly frequency has diminished over the extended sample period ending in 2015, using both statistical and economic significance metrics. Employing formal structural break tests, we provide evidence of a break in the predictive relationship in the third quarter of 2008 for most of these indexes. At first sight, these results might suggest that oil price changes are useless for forecasting international equity index returns.

[3] See *Fortune*, October 15, 2014 (`http://fortune.com/2014/10/15/stock-market-plunge-oil-prices/`).

However, we contend that information contained in oil price changes becomes useful once it is suitably complemented with relevant information about oil supply and global economic activity. The key observation, as made by Kilian (2009), is that oil price changes are driven by various supply and demand shocks that fundamentally play different roles. Accordingly, using a variant of the structural VAR approach of Kilian (2009), we obtain an oil price change decomposition into an oil supply shock, a global demand shock, and an oil-specific demand shock and argue that these three different shocks should have different implications for international equity markets. Subsequently, we illustrate the ability of these three shocks to predict G7 country MSCI index excess returns, using metrics of both statistical and economic significance, and the structural stability of this predictive relationship over the last thirty years.

This paper relates to a growing literature that examines the impact of oil price shocks on the real economy and equity markets. Chen, Roll, and Ross (1986), Jones and Kaul (1996), and Kilian and Park (2009), among others, examine the contemporaneous relationship between the price of oil and stock prices.[4] Kilian and Park (2009) augment the structural VAR model of Kilian (2009) by adding the US real stock return in the vector of variables and study contemporaneous relationships between shocks embedded in oil price changes and stock returns. They examine cumulative impulse responses of real stock returns to one-time shocks to oil supply, global demand, and oil-specific demand in the crude oil market. Their results, using data in the period of 1975–2006, show that an unexpected decrease in oil production has no significant effect on cumulative US real stock returns and a positive surprise to global demand (oil-specific demand) leads to a subsequent increase (decrease) in US real stock returns. In addition, several recent papers, including Driesprong, Jacobsen, and Maat (2008), Casassus and Higuera (2012), and Narayan and Gupta (2015), investigate the ability of oil price shocks to forecast equity index returns and document a negative predictive slope of oil price changes. Unlike Kilian and Park (2009), and in line with the recent finance literature and Driesprong, Jacobsen, and Maat (2008) in particular, we use a predictive regression framework to examine whether information contained in oil price changes can be used to forecast future stock returns. Although we utilize the structural VAR model by Kilian (2009), we approach the question from the perspective of an

[4] The contemporaneous relationship between the volatility of oil prices and stock returns has been studied by several papers, such as Chiang, Hughen, and Sagi (2015) and Christoffersen and Pan (2014).

investor who wishes to use the real-time information embedded in oil price changes and captured by the three aforementioned shocks to predict future MSCI equity index returns.

Our work is also related to some recent empirical research that focuses on disentangling the intrinsic shocks embedded in oil price changes. Rapaport (2014) and Ready (2016) propose to use information from the stock market to identify the underlying types of shocks in oil price changes. Rapaport (2014) argues that using the sign and magnitude of the correlation between daily oil price changes and aggregate stock market returns, excluding oil companies, allows one to identify shocks specific to the oil market and shocks that affect the overall economy. Ready (2016) uses crude oil futures returns, returns on a global equity index of oil producing firms, and innovations to the VIX index to identify demand and supply shocks. Ready (2016) focuses on the *contemporaneous* relationship between aggregate market returns and the three shocks obtained from his decomposition. In contrast, since we focus on the ability of the various shocks embedded in oil price changes to *forecast* equity index returns, we adopt the approach advanced by Kilian (2009), which utilizes more direct proxies for oil supply and global demand and does not require information from the stock market in the decomposition.

In this paper, we make a number of contributions to the literature studying the impact of oil price fluctuations on international equity returns. First, we document that the ability of oil price changes to forecast G7 country MSCI index returns has declined significantly over the last decade. In particular, using formal structural break tests, we detect a break in the predictive relationship in the third quarter of 2008 for most of the indexes under examination.

Second, using a variant of the structural VAR approach of Kilian (2009), we obtain a decomposition of oil price change into an oil supply shock, a global demand shock, and an oil-specific demand shock. To do so, we use the first principal component of the log growth rates of WTI, Dubai, and Arab Light spot prices as a comprehensive proxy for oil price change. Moreover, we employ two proxies for global real economic activity, namely a shipping cost index and global crude steel production, and use the first principal component of their log growth rates as a comprehensive proxy for global demand growth. Importantly, all the variables that we use in our empirical tests are constructed based on real-time available information.

Third, we illustrate the ability of these three shocks to predict G7 country MSCI index returns, denominated in both local currency and US dollars. In particular, the conjecture that oil supply shocks and oil-specific demand shocks (global demand shocks) predict equity returns with a negative (positive) slope is supported by the empirical evidence over the 1986–2015 period. Moreover, we detect no structural breaks in the predictive relationship between the three aforementioned shocks and G7 country MSCI equity index returns. We also demonstrate the advantage of using the oil price decomposition instead of just the oil price change, in economic terms, by the substantial and statistically significant improvement in the performance of simple mean-variance trading strategies. Specifically, for the case of the MSCI World index, the certainty equivalent return and Sharpe ratio increase from 3.88% to 7.90% and from 0.30 to 0.56, respectively.

In addition, we examine various other aspects of the predictive relationship. To address real-time data availability concerns, we construct returns with a delay of one and two weeks and show that the results are essentially identical. We further demonstrate that, as the forecasting horizon increases from one to six months, the predictive ability of the three shocks gradually diminishes. For the case of the United States, we document that the forecasting ability of the three shocks is robust in the cross section of industries; we further show that the three shocks have low correlations with the standard macroeconomic variables used to predict aggregate market returns and that the three shocks remain strong predictors in the presence of these alternative predictors. Moreover, the estimated conditional expected returns based on the three shocks exhibit high volatility and low persistence in comparison to risk premia estimates available in the literature. Finally, these three shocks do not appear to have an effect on conditional return volatility. Collectively, these results do not appear to be consistent with the notion of time-varying risk premia.

The rest of the paper proceeds as follows. In Section 2, we describe the data that we use in our empirical exercises. In Section 3, we describe the metrics of statistical and economic significance that we employ to evaluate the ability of the various quantities of interest to predict G7 country MSCI equity index returns at the monthly frequency. In Section 4, we present empirical evidence on the forecasting ability of oil price changes and how it has changed over the last decade. In

Section 5, we introduce a decomposition of oil price changes into oil supply shocks, global demand shocks, and oil-specific demand shocks and illustrate the ability of these three shocks to forecast G7 country MSCI index returns. In Section 6, we offer some concluding remarks.

2 Data

We use five different types of data: returns on international equity indexes, short-term interest rates, oil price proxies, proxies for global economic activity, and global oil production. The full sample period is from January 1982 to December 2015.

We use returns on MSCI equity indexes for the G7 countries, denominated in both local currency and US dollars.[5] We collect monthly short-term interest rates from the International Monetary Fund (IMF) and the Organisation for Economic Cooperation and Development (OECD). We use IMF Treasury bill rates when these rates are available and short-term interest rates obtained from the OECD otherwise.[6]

We use three proxies for oil price, namely the WTI spot price, the Dubai spot price, and the Arab Light spot price.[7] Note that 75% (83%) of the log growth rates of WTI (Arab Light) prices from October 1973 to September 1981 are zero. Therefore, it is problematic to use WTI and Arab Light prices before September 1981. Therefore, in our empirical analysis, we use oil price data from 1982 onwards. Following Driesprong, Jacobsen, and Maat (2008), we use nominal oil prices.

We combine the information contained in the three proxies for crude oil spot price into a single proxy using Principal Component Analysis (PCA). The single proxy, denoted by $g^{\mathtt{P}}$ where $\mathtt{P}$ stands for price, is represented by the first principal component of the log growth rates of WTI, Dubai, and Arab Light spot prices. The details of the construction of the single PCA proxy

[5] Specifically, data on MSCI indexes for the G7 countries, i.e., Canada, France, Germany, Italy, Japan, the United Kingdom, and the United States, as well as the World MSCI index are obtained from Datastream. [6] For Canada, France, Italy, Japan, and the United Kingdom, we use Treasury bill rates from the IMF. For Germany, we use Treasury bill rates from the IMF and, from September 2007, short-term interest rates from the OECD. For the United States, we use the 1-month Treasury bill rate taken from Kenneth French's website. [7] Data on the Dubai and Arab Light spot prices are obtained from Bloomberg. Data on WTI spot oil prices for the period of between January 1982 and August 2013 are obtained from the website of St. Louis Fed. Data for the period between August 2013 and December 2015 are obtained from Bloomberg.

for oil price change are given in Appendix A.1. To make the proxy g^{P} comparable to the three individual proxies, we rescale it so that its standard deviation equals 0.09 over the sample period of January 1983 to December 2015. Table 1 presents summary statistics, including correlations, for g^{P} and the log growth rates of the three oil price proxies. Over the subsample period ending in April 2003, which is the last month in the sample used in Driesprong, Jacobsen, and Maat (2008), as well as the full sample period, g^{P} is highly correlated with the log growth rates of the three individual proxies. Figure 2 also shows that the four series track each other quite closely.

We use two proxies for global economic activity to capture changes in global demand. The first proxy is a shipping cost index constructed from data on dry cargo single voyage rates and the Baltic Dry Index (BDI). Since the supply of bulk carriers is largely inelastic, fluctuations in dry bulk cargo shipping cost are thought to reflect changes in global demand for transporting raw materials such as metals, grains, and coals by sea. Therefore, shipping cost is considered to be a useful leading indicator of global economic activity. We hand-collect data on dry cargo single voyage rates from Drewry Shipping Statistics and Economics for the period between January 1982 and January 1985. Rates for seven representative routes are reported each month. We compute the monthly log growth rates of the shipping cost for each route, and then, following Kilian (2009), obtain their equally-weighted average.[8] Data on the BDI from January 1985 to December 2015 are obtained from Bloomberg.

The second proxy for global economic activity is global crude steel production. Ravazzolo and Vespignani (2015) argue that world steel production is a good indicator of global real economic activity. Steel is widely used in a number of important industries, such as energy, construction, automotive and transportation, infrastructure, packaging, and machinery. Therefore, fluctuations in world crude steel production reflect changes in global real economic activity. We obtain monthly crude steel production data for the period between January 1990 and December 2015 from the website of the World Steel Association. The reported monthly figure represents crude steel production in 66 countries and accounts for about 99% of total world crude steel production. In addition, we hand-collect monthly data for the period between January 1968 to October 1991

[8] It is, however, worth noting that there is an important difference between our proxy for global economic activity and the one constructed in Kilian (2009). Specifically, in Kilian (2009), the average growth rate is cumulated, then deflated using the US CPI and finally detrended. In that sense, the proxy in Kilian (2009) is a level variable. In contrast, our proxy is a growth rate.

from the Steel Statistical Yearbook published by the International Iron and Steel Institute. Crude steel production exhibits strong seasonality and, hence, we seasonally adjust the data, as we explain in Appendix A.1.

As in the case of oil price proxies, we use PCA to construct a single proxy for global demand growth. The single proxy, denoted by $g^{\mathtt{D}}$ where D stands for demand, is represented by the first principal component of the log growth rates of the shipping cost index and global crude steel production. The details of the construction of the single PCA proxy for global demand growth are given in Appendix A.1. The correlations between $g^{\mathtt{D}}$ and the log growth rates of the shipping cost index and global crude steel production are 0.82 and 0.74, respectively. Figure 3 shows that $g^{\mathtt{D}}$ tracks closely the two individual proxies most of time, except for a few instances in which one of the two proxies takes extreme values.

Finally, we obtain oil production data, covering the period between January 1982 and December 1991, from the website of the US Energy Information Agency.[9] In addition, we hand-collect data on the total supply of crude oil, natural gas liquids, processing gains, and global biofuels, for the period between December 1991 and December 2015, from the monthly Oil Market Report obtained from the website of the International Energy Agency. Combining data from the two sources, we construct a time series of monthly log growth rates of world crude oil production.

3 Evaluation of predictive ability

In this paper, we examine the ability of (i) oil price changes and (ii) the oil supply, global demand, and oil-specific demand shocks embedded in these changes to forecast G7 country MSCI excess returns. We do so in the context of linear predictive regressions of the following type:

$$r^{\mathsf{e}}_{t+1} = \alpha + \boldsymbol{\beta}'\mathbf{x}_t + u_{t+1}, \tag{1}$$

where r^{e}_{t+1} is the excess return on a G7 country MSCI index, $\mathbf{x}_t = [x_{1,t} \ \cdots \ x_{n,t}]'$ is a vector of predictors, $\boldsymbol{\beta} = [\beta_1 \ \cdots \ \beta_n]'$ is the vector of predictive slope coefficients, and u_{t+1} is a zero-mean

[9] Specifically, we use Table 11.1b (World Crude Oil Production: Persian Gulf Nations, Non-OPEC, and World).

random disturbance. Specifically, $\mathbf{x}_t$ is a scalar ($n = 1$) when we evaluate oil price change as a predictor, while $\mathbf{x}_t$ is a three-dimensional vector ($n = 3$) when we evaluate the three different shocks as predictors. We also consider the i.i.d. model for r^{e}_{t+1} in which case the vectors $\boldsymbol{\beta}$ and $\mathbf{x}_t$ are null and equation (1) reduces to $r^{\mathsf{e}}_{t+1} = \alpha + u_{t+1}$. We provide evidence on predictive ability in terms of statistical as well as economic significance.

The question we wish to address is whether $\mathbf{x}_t$ can forecast the MSCI index excess return r^{e}_{t+1}. Hence, we are interested in testing the null hypotheses $H_0 : \beta_i = 0$, for $i = 1, \ldots, n$. We evaluate the statistical significance of predictive ability of $\mathbf{x}_t$ using standard metrics. Specifically, we obtain two-sided p-values for the null hypotheses $H_0 : \beta_i = 0$, $i = 1, \ldots, n$ based on standard errors computed according to two well-established approaches: the method advanced by Newey and West (1987), where the optimal bandwidth is selected following the approach in Newey and West (1994), as well as the method developed by Hodrick (1992) that imposes the no-predictability condition. Finally, we also report adjusted R-squares.

To gauge the economic significance of the predictive ability of $\mathbf{x}_t$, we consider a mean-variance investor who can invest in an MSCI index and the corresponding short-term Treasury bill. The investor uses the regression model (1) to forecast MSCI index excess returns. A trading strategy is then developed based on the resulting estimates of the conditional mean and variance of excess returns. We evaluate economic significance in terms of two commonly used metrics: (i) the certainty equivalent return (CER) and (ii) the Sharpe ratio (SR) of the associated optimal portfolio returns.

Following Campbell and Thompson (2008), we assume that the risk aversion coefficient of the mean-variance investor is $\gamma = 3$. At the end of each period t, the investor uses all available data to estimate the parameters of the linear predictive regression (1). Using these parameter estimates, the investor then obtains estimates of the mean and the variance of the MSCI index excess return r^{e}_{t+1} at time t, denoted by $\widehat{\mu}_{t+1}$ and $\widehat{v}_{t+1}$, respectively. These estimates give rise to the following optimal portfolio weight on the MSCI index:

$$\omega_t = \frac{1}{\gamma} \frac{\widehat{\mu}_{t+1}}{\widehat{v}_{t+1}}. \tag{2}$$

The rest of the investor's wealth is invested in the short-term Treasury bill. We assume that the portfolio weight on the MSCI index is constrained between a minimum and maximum feasible weight, denoted by $\underline{\omega}$ and $\overline{\omega}$, respectively. The minimum weight $\underline{\omega}$ is set equal to zero so that short-selling is precluded. Following Campbell and Thompson (2008), we set the maximum weight $\overline{\omega}$ equal to 150% so that the investor is allowed to borrow up to 50% and invest the proceeds in the MSCI index. Optimal weights are determined according to equation (2) and then the realized portfolio returns are computed. Below, we describe the two metrics, CER and SR, used in our evaluation of economic significance of predictability.

The CER of the resulting optimal portfolio from period 1 to period T based on the predictive regression (1) is given by

$$\widehat{\texttt{CER}} = \widehat{\mu}_p - \frac{1}{2}\gamma\widehat{v}_p, \tag{3}$$

with the mean $\widehat{\mu}_p$ and the variance $\widehat{v}_p$ of the realized optimal portfolio simple returns being defined by

$$\widehat{\mu}_p = \frac{1}{T}\sum_{t=0}^{T-1}(r^{\mathtt{f}}_{t+1} + \omega_t r^{\mathtt{e}}_{t+1}) \quad \text{and} \quad \widehat{v}_p = \frac{1}{T}\sum_{t=0}^{T-1}\left((r^{\mathtt{f}}_{t+1} + \omega_t r^{\mathtt{e}}_{t+1}) - \widehat{\mu}_p\right)^2, \tag{4}$$

and where $r^{\mathtt{e}}_{t+1}$ and $r^{\mathtt{f}}_{t+1}$ denote the excess return on the MSCI index and the corresponding Treasury bill rate at time $t+1$, respectively.

The SR of the resulting optimal portfolio from period 1 to period T based on the predictive regression (1) is given by

$$\widehat{\texttt{SR}} = \frac{\widehat{\mu}^{\mathtt{e}}_p}{\sqrt{\widehat{v}^{\mathtt{e}}_p}}, \tag{5}$$

with the mean and the variance of the realized optimal portfolio excess returns being defined by

$$\widehat{\mu}^{\mathtt{e}}_p = \frac{1}{T}\sum_{t=0}^{T-1}\omega_t r^{\mathtt{e}}_{t+1} \quad \text{and} \quad \widehat{v}^{\mathtt{e}}_p = \frac{1}{T}\sum_{t=0}^{T-1}(\omega_t r^{\mathtt{e}}_{t+1} - \widehat{\mu}^{\mathtt{e}}_p)^2. \tag{6}$$

We express the CERs obtained from equation (3) in annualized percentages by multiplying by 1,200 and annualize the monthly SRs from equation (5) by multiplying by $\sqrt{12}$. CER represents the equivalent risk-free rate of return that a mean-variance investor would require in exchange of a risky portfolio return series, while SR measures the average portfolio excess return per unit of risk as measured by the portfolio excess return standard deviation.

If the variables in the vector $\mathbf{x}_t$ have nontrivial predictive ability, then using the predictive regression model (1) is expected to generate a higher CER and/or SR than using an i.i.d. model for the MSCI index excess returns r^{e}_{t+1}. In this context, we naturally refer to the i.i.d. model as the baseline model (Model 1) and the predictive regression model using the vector of predictors $\mathbf{x}_t$ as the augmented model (Model 2). Denote by $\texttt{CER}_j$ and $\texttt{SR}_j$ the CER and SR of Model j, for $j = 1, 2$. Even if the point estimate of the CER and/or SR generated by the augmented model is higher that its counterpart generated by the baseline model, i.e., $\widehat{\texttt{CER}}_1 < \widehat{\texttt{CER}}_2$, one might be concerned whether this is due to genuine predictive ability of $\mathbf{x}_t$ or simply to sample variability. Therefore, it is important to test the statistical significance of any differences in CER and/or SR. We follow Garlappi, Skoulakis, and Xue (2016) who develop asymptotic tests for the null hypothesis $H_0^{\mathsf{CER}} : \texttt{CER}_1 = \texttt{CER}_2$ against the one-sided alternative $H_A^{\mathsf{CER}} : \texttt{CER}_1 < \texttt{CER}_2$ and similarly the null hypothesis $H_0^{\mathsf{SR}} : \texttt{SR}_1 = \texttt{SR}_2$ against the one-sided alternative $H_A^{\mathsf{SR}} : \texttt{SR}_1 < \texttt{SR}_2$. Our purpose is to evaluate the incremental predictive ability of $\mathbf{x}_t$ compared to the i.i.d. model for MSCI index excess returns and, hence, we naturally focus on one-sided alternative hypotheses. The same framework can be used to compare the predictive ability of oil price change to that of the vector of oil supply, global demand, and oil-specific demand shocks. For this comparison, Model 1 (baseline) would correspond to the predictive regression (1) with $\mathbf{x}_t$ consisting of the oil price change while Model 2 (augmented) would correspond to the predictive regression (1) with $\mathbf{x}_t$ consisting of the three different shocks. The tests are based on standard heteroscedasticity and autocorrelation consistent (HAC) variance-covariance matrix estimators. The reported p-values are based on the Newey and West (1987) procedure with a Bartlett kernel and optimal bandwidth selected as in Newey and West (1994). Details about the computation of p-values are presented in Appendix A.2.

4 Oil price change as a predictor of MSCI index excess returns

To examine whether oil price changes can predict equity returns, we start by revisiting the evidence documented in Driesprong, Jacobsen, and Maat (2008) who consider a sample period ending in April 2003. In particular, we estimate the following standard predictive regression model:

$$r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_{t} + u^{\mathsf{P}}_{t+1}, \tag{7}$$

where r^{e}_{t+1} is the excess return on an MSCI index and the oil price change proxy g^{P}_{t} is the first principal component obtained from three oil spot price log growth rates: WTI, Dubai, and Arab Light. We construct excess returns by subtracting the particular country short-term rate from each MSCI index return in the case of local-currency denominated indexes and by subtracting the US T-bill rate from each MSCI index return in the case of US-dollar denominated indexes.[10]

We consider the MSCI indexes for the G7 countries as well as the World MSCI index, denominated both in local currencies and US dollars. The oil price change proxy we use is the first principal component obtained from the WTI, Dubai, and Arab Light spot prices as explained in Section 2. Driesprong, Jacobsen, and Maat (2008) document negative and statistically significant estimates of the predictive slope coefficient δ^{P} for a large number of countries based on a sample that ends in April 2003. We first run the predictive regression for the sample ending in April 2003 and then consider the extended sample period ending in December 2015. We first examine the statistical significance of predictability in terms of p-values and adjusted R-squares. Furthermore, we examine its economic significance by evaluating the performance of the resulting trading strategies in terms of certainty equivalent returns and Sharpe ratios.

4.1 Evidence based on data until 2003

In Table 2, we present statistical significance results for the predictive regression (7) based on MSCI index excess returns, denominated in both local currencies and US dollars. The sample period we consider starts in January 1983 and ends in April 2003. We first focus our analysis

[10] While Driesprong, Jacobsen, and Maat (2008) use log returns in their empirical analysis, it is more convenient for us to use excess returns for the purpose of assessing the economic significance of the predictive ability of oil price changes. The results for log returns, available upon request, are very similar.

on this sample period to facilitate comparison of our results with the evidence presented in Driesprong, Jacobsen, and Maat (2008) who also use a sample ending in April 2003.

For local-currency denominated returns, the point estimate of the predictive slope δ^{P} in regression (7) is negative for all six cases. The null hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in six (four) out of six cases at the 10% (5%) level of significance according to Newey and West (1987) standard errors. When Hodrick (1992) standard errors are used, $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in five (four) out of six cases at the 10% (5%) level of significance, with Japan yielding the highest p-value equal to 0.12. The adjusted R-square is higher than 2% in five out of six cases, with Canada yielding the lowest adjusted R-square equal to 1.0%.[11] Hence, even though we use our own oil price change proxy, our sample starts at a different point in time, and we use excess returns as opposed to log returns, our results confirm the evidence reported in Driesprong, Jacobsen, and Maat (2008) on the relationship between oil price changes and subsequent global equity returns for the time period extending until April 2003.

In Table 3, we present economic significance results on the ability of oil price changes to predict MSCI index excess returns in terms of the certainty equivalent return (CER) and the Sharpe ratio (SR) of the associated trading strategies, as explained in Section 3. These results reinforce the statistical significance results reported in Table 2.

Let $\mathtt{CER}_{\mathtt{IID}}$ and $\mathtt{SR}_{\mathtt{IID}}$ denote the CER and SR achieved by the trading strategy assuming that the MSCI index excess returns are i.i.d., and $\mathtt{CER}_{\mathtt{P}}$ and $\mathtt{SR}_{\mathtt{P}}$ denote the CER and SR achieved by the trading strategy using the predictive regression model (7).

For local-currency denominated returns, the augmented model using the oil price change proxy g^{P} as predictor generates significantly higher (point estimates of) CERs and SRs compared to the baseline model that assumes that MSCI index excess returns are i.i.d. across all six cases. More importantly, the null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is rejected in six (five) out of

[11] The results for US-dollar denominated returns are qualitatively similar. Specifically, the point estimate of the predictive slope δ^{P} in regression (7) is negative for all eight cases. The null hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in seven (six) out of eight cases at the 10% (5%) level of significance according to Newey and West (1987) standard errors. When Hodrick (1992) standard errors are used, $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in six (six) out of eight cases at the 10% (5%) level of significance, with Japan again yielding the highest p-value equal to 0.27. The adjusted R-square is higher than 2% in six out of eight cases, with Canada yielding the lowest adjusted R-square equal to 0.7%. Importantly, for the case of the World MSCI index, $H_0 : \delta^{\mathsf{P}} = 0$ is strongly rejected by both methods and the adjusted R-square is equal to 5.6% which is rather high for monthly returns.

six cases at the 10% (5%) level of significance, respectively, against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$. The null hypothesis $H_0^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{P}}$ is rejected against the one-sided alternative $H_A^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} < \mathtt{SR}_{\mathtt{P}}$ in all six cases at the 5% level of significance.[12] Overall, the economic significance results confirm the evidence reported in Driesprong, Jacobsen, and Maat (2008) on the ability of oil price changes to forecast international equity index returns.

4.2 Evidence based on data until 2015

In this subsection, we extend the sample period to December 2015 and run the same predictive regressions again. As in the previous subsection, we examine both the statistical and economic significance of the predictive ability of oil price changes.

In Table 4, which corresponds to Table 2, we present statistical significance results. The evidence obtained from the extended sample is quite different: the predictive ability of oil price changes has mostly disappeared.

For local-currency denominated returns, the point estimates of the predictive slope δ^{P} in regression (7) are still negative in all six cases. However, they are much smaller in absolute value. For instance, for Japan and the UK, the δ^{P} point estimates obtained over the 1983.01–2003.04 period are -0.11 and -0.11, while they fall to -0.05 and -0.06 over the 1983.01–2015.12 period, respectively. The null hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is now rejected in only three (two) out of six cases at the 10% (5%) level of significance according to Newey and West (1987) standard errors. Moreover, we observe a substantial reduction in adjusted R-squares. For instance, for Japan and the UK, the adjusted R-squares obtained over the 1983.01–2003.04 period are 2.9% and 4.1%, while they fall to 0.3% and 1.4% over the 1983.01–2015.12 period, respectively. In the case of Canada, the adjusted R-square even becomes negative.

[12] The results for US-dollar denominated returns are qualitative similar. The augmented model based on the predictive regression model (7) still generates significantly higher (point estimates of) CERs and SRs compared to the baseline i.i.d. model for MSCI index excess returns across all eight cases. Importantly, for the case of the World MSCI index, when g^{P} is used as predictor, the CER increases from $\widehat{\mathtt{CER}}_{\mathtt{IID}} = 2.26\%$ to $\widehat{\mathtt{CER}}_{\mathtt{P}} = 7.79\%$ and, similarly, the SR increases from $\widehat{\mathtt{SR}}_{\mathtt{IID}} = 0.19$ to $\widehat{\mathtt{CER}}_{\mathtt{P}} = 0.43$. The null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is rejected in five (three) out of eight cases at the 10% (5%) level of significance, respectively, against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$, with Canada yielding the highest p-value equal to 0.15. The null hypothesis $H_0^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{P}}$ is rejected against the one-sided alternative $H_A^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} < \mathtt{SR}_{\mathtt{P}}$ in six (three) out of eight cases at the 10% (5%) level of significance, respectively, with Japan yielding the highest p-value equal to 0.14.

The results for US-dollar denominated returns are even weaker. While the δ^{P} point estimates are still negative in seven out of eight cases, they are even smaller in absolute value than their counterparts obtained for local-currency denominated returns and in the case of Canada the predictive slope estimate turns out to be positive. The null hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is rejected only in the case of Italy at the 10% level of significance, regardless of whether we use Newey and West (1987) or Hodrick (1992) standard errors. Importantly, the corresponding p-values for the USA and the World MSCI indexes are 0.22 and 0.22, respectively, according to Newey and West (1987) standard errors. In addition, the adjusted R-squares are rather low: they are less than 1% in seven out of eight cases, and even negative in the case of Canada. Hence, our results show substantially weaker statistical evidence on the relationship between oil price changes and subsequent global equity excess returns over the sample extending to December 2015 compared to the sample period ending in April 2003.

Table 5, which corresponds to Table 3, reports results on the economic significance of the predictive ability of oil price changes in terms of the CER and the SR of the associated trading strategies. These results reinforce the message conveyed by Table 4 that the forecasting ability of oil price changes has essentially disappeared over the extended sample ending in December 2015.

For local-currency denominated returns, the null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is *not* rejected against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$ in four out of six cases at the 10% level of significance, with Italy and the UK being the exceptions. We obtain the same results when we test $H_0^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{P}}$ against $H_A^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} < \mathtt{SR}_{\mathtt{P}}$.

The results for US-dollar denominated returns are even weaker. The null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is *not* rejected against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$ in any case, out of eight, at the 10% level of significance. Importantly, the corresponding p-values for the USA and the World MSCI indexes are 0.37 and 0.30, respectively.

Collectively, the statistical as well as economic significance results presented in this subsection illustrate that the forecasting ability of oil price changes has been diminished over the extended sample ending in December 2015. In the next subsection, we provide further corroborating evidence by examining the stability, or lack thereof, of the predictive relation between MSCI index excess returns and past oil price changes.

4.3 Instability of the predictive slope coefficients

The empirical evidence gathered in the previous two subsections suggests that the ability of oil price changes to predict MSCI index excess returns is not stable over time. We confirm extant results on the oil price change predictive ability using data until April 2003, consistent with the evidence in Driesprong, Jacobsen, and Maat (2008), but also show that these results do not hold in the extended sample ending in December 2015. While we obtain negative and statistically significant predictive slope estimates in the early sample, these estimates become much closer to zero and lose their statistical significance in the extended sample.

As a first attempt to shed some light on these striking findings, we estimate the predictive regression model (7) over different samples using an expanding window with the first sample being 1983.01–1993.01 and the last sample being 1983.01–2015.12. Figures 4 and 5 present the predictive slope estimates along with 95% confidence intervals, based on Newey and West (1987) standard errors, over the period 1993.01–2015.12 for local-currency and US-dollar denominated MSCI index returns, respectively. The pattern evident in these graphs is rather revealing. For the majority of the cases, the predictive slope estimates are negative and frequently statistically significant until the third quarter of 2008. For many cases, however, after that point in time the estimates start increasing to zero and quickly lose their statistical significance. This effect is more noticeable for US-dollar denominated returns.

In addition to the informal analysis based on the predictive slope estimates presented in Figures 4 and 5, we also perform formal structural break tests. Specifically, we employ the methodology developed by Bai and Perron (2003) to test for multiple structural breaks in the predictive slope coefficients. We use the Bayesian Information Criterion (BIC) to select the number of breaks.

Table 6 presents Bai and Perron (2003) structural break tests in the slope coefficient for the predictive regression (7). The second column presents the BIC values assuming no break. The third and fourth columns provide the BIC values and the corresponding break dates for the case of the one-break model. The last column shows the number of breaks selected by the BIC. For local-currency denominated index returns, the test identifies the presence of one structural break in five out of six cases, with UK being the only exception. For US-dollar denominated index returns, the test identifies the presence of one structural break in seven out of eight cases, with France being the only exception. The break dates identified in most cases fall in the third quarter of 2008. However, the break dates for Italy and Japan are October 2003 and September 1990, respectively. Overall, the structural break tests provide additional evidence against the stability of the slope coefficient in the predictive relationship between MSCI index excess returns and past oil price changes.

5 The differential roles of the various shocks embedded in oil price changes

In the previous section, we confirm the finding of Driesprong, Jacobsen, and Maat (2008) that oil price changes predict international equity index returns at the monthly frequency with a negative predictive slope based on data up to April 2003. However, we also provide compelling evidence that the predictive power of oil price changes has practically disappeared over the extended sample ending in December 2015. For most of the G7 MSCI indexes, the predictive slope estimates based on expanding windows become closer to zero and turn statistically insignificant after the third quarter of 2008. Moreover, the formal econometric tests of Bai and Perron (2003) indicate the existence of a structural break in the third quarter of 2008 for the majority of the cases, especially when US-dollar denominated returns are used. The dramatic reduction in the predictive ability of oil price changes, therefore, begs for an explanation.

In this paper, we offer an explanation that emphasizes the differential roles of the various shocks embedded in oil price changes. In particular, we adapt the structural VAR framework of Kilian (2009) that provides a decomposition of oil price changes into oil supply shocks, global demand shocks, and oil-specific demand shocks. As pointed out by Kilian (2009) and Kilian and Park (2009), oil price shocks cannot be treated as strictly exogenous with respect to the global economy. In particular, they argue that oil supply shocks, global demand shocks, and oil-specific demand shocks, the combination of which leads to the observed aggregate oil price changes, should have different effects on the macroeconomy and the stock market.

Alternatively, Rapaport (2014) and Ready (2016) propose to use information from the stock market to identify the underlying types of shocks in oil price changes. Rapaport (2014) identifies shocks specific to the oil market and shocks that affect the overall economy using the sign and magnitude of the correlation between daily oil price changes and aggregate stock market returns, excluding oil companies. Ready (2016) uses crude oil futures returns, returns on a global equity index of oil producing firms, and innovations to the VIX index to identify demand and supply shocks. He documents a strong *contemporaneous* relationship between aggregate market returns and the demand/supply shocks from his decomposition based on data from 1986 to 2011. We have confirmed that his results remain strong using data until 2015. However, the shocks identified by Ready (2016) cannot forecast future stock market returns. In contrast, the focus of our paper is the *predictive* relationship between the various shocks embedded in oil price changes and equity index returns. Hence, we find the approach advanced by Kilian (2009), which utilizes more direct proxies for oil supply and global demand and does not require stock market information to obtain the decomposition, more suitable for our purposes.

Kilian and Park (2009) augment the structural VAR model of Kilian (2009) by adding the US real stock return in the vector of variables and study contemporaneous relationships between shocks embedded in oil price changes and stock returns. They examine cumulative impulse responses of real stock returns to one-time shocks to oil supply, global demand, and oil-specific demand in the crude oil market. Their results, using data in the period of 1975–2006, show that an unexpected decrease in oil production has no significant effect on cumulative US real stock returns and a positive surprise to global demand (oil-specific demand) leads to a continuous increase (decrease) in US real stock returns. In this paper, we follow the recent finance literature, Driesprong, Jacobsen, and Maat (2008) in particular, we cast the question in a predictive regression framework using short-horizon, i.e., one-month-ahead, forecasts. Although we utilize the framework proposed by Kilian (2009), we approach the question from the perspective of an investor who wishes to use the real-time information embedded in oil price changes and captured by the three aforementioned shocks to predict subsequent equity index returns.

5.1 A decomposition based on a structural VAR model

To disentangle the supply shocks, demand shocks, and oil-specific demand shocks embedded in the observed oil price changes, we employ a variant of the structural Vector Autoregressive (VAR) framework of Kilian (2009). Specifically, we consider a VAR model based on three variables capturing changes in the (i) supply of oil; (ii) demand for oil; and (iii) price of oil. The first variable, denoted by $g_t^{\mathtt{S}}$ where $\mathtt{S}$ stands for supply, is the log growth rate of world crude oil production. The second variable, denoted by $g_t^{\mathtt{D}}$ where $\mathtt{D}$ stands for demand, is the first principal component of the log growth rates of the dry bulk cargo shipping cost index and global crude steel production. It has been argued in the literature, e.g., Kilian (2009) and Ravazzolo and Vespignani (2015), among others, that fluctuations in shipping cost and global crude steel production capture changes in global economic activity growth and demand for oil. The third variable, denoted by $g_t^{\mathtt{P}}$ where $\mathtt{P}$ stands for price, is the first principal component of the log growth rates of West Texas Intermediate, Dubai, and Arab Light spot prices. We provide a detailed explanation of the data sources and construction in Section 2.

Our purpose is to employ the structural VAR model to obtain a decomposition of oil price changes into three types of shocks and use them as predictors of MSCI index returns. We do so by, first, using real-time available information and, second, constructing three variables that are stationary in a consistent way. As a result, the choice of variables in our implementation of the structural VAR model differs from that used by Kilian (2009) in two aspects. First, the index of global real economic activity constructed in Kilian (2009) is a level variable constructed by first cumulating the average growth rate of dry bulk cargo freight rates, then deflating by the US CPI, and finally linearly detrending. Our proxy for global economic activity is the first principal component the log growth rates of two indicators, the shipping cost index and world crude steel production, constructed in a real-time fashion. Second, Kilian (2009) uses log real oil prices, measured by US refiners' acquisition costs deflated by the US CPI. We use an oil price change proxy, obtained as the first principal component of the log growth rates of three proxies for crude oil spot prices. Importantly, as pointed out by Apergis and Miller (2009), the variables of global real economic activity and log real oil price used in the structural VAR model in Kilian (2009) appear to be non-stationary. On the contrary, our variables are, by construction, stationary.

Letting $\mathbf{g}_t = [g_t^{\mathtt{S}} \quad g_t^{\mathtt{D}} \quad g_t^{\mathtt{P}}]'$ denote the vector of the three variables described above, the structural VAR model is stated as:

$$\mathbf{A}_0\mathbf{g}_t = \mathbf{a} + \sum_{i=1}^{p} \mathbf{A}_i \mathbf{g}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{8}$$

where $\mathbf{A}_0$ is a 3×3 lower triangular matrix, $\mathbf{a}$ is a 3×1 vector, $\mathbf{A}_i$ is a 3×3 matrix, for $i = 1, \ldots, p$, and $\boldsymbol{\varepsilon}_t = [\varepsilon_t^{\mathtt{S}} \;\; \varepsilon_t^{\mathtt{D}} \;\; \varepsilon_t^{\mathtt{OSD}}]'$ is a vector of uncorrelated standardized shocks. The interpretation of the fundamental shocks is as follows: $\varepsilon_t^{\mathtt{S}}$ is the oil supply shock, $\varepsilon_t^{\mathtt{D}}$ is the global demand shock, and $\varepsilon_t^{\mathtt{OSD}}$ is the oil-specific demand shock. The structural innovation vectors $\boldsymbol{\varepsilon}_t$ are, by assumption, serially and cross-sectionally uncorrelated. The reduced-form VAR innovation is $\mathbf{e}_t = \mathbf{A}_0^{-1}\boldsymbol{\varepsilon}_t$, where $\mathbf{A}_0^{-1}$ is simply the Cholesky factor of the covariance matrix $\boldsymbol{\Sigma}_e = \mathrm{Var}[\mathbf{e}_t]$. Multiplying both sides of the above equation by $\mathbf{A}_0^{-1}$ yields

$$\mathbf{g}_t = \mathbf{b} + \sum_{i=1}^{p} \mathbf{B}_i \mathbf{g}_{t-i} + \mathbf{e}_t \tag{9}$$

where $\mathbf{b} = \mathbf{A}_0^{-1}\mathbf{a}$ and $\mathbf{B}_i = \mathbf{A}_0^{-1}\mathbf{A}_i$, $i = 1, \ldots, p$. The parameters $\mathbf{b}$ and $\mathbf{B}_i$ are estimated by standard OLS and the VAR order p is selected using the BIC criterion. Writing the VAR(p) system in VAR(1) form, we obtain

$$\mathbf{y}_t = \mathbf{C}\mathbf{y}_{t-1} + \mathbf{u}_t, \tag{10}$$

where $\mathbf{y}_t$ and $\mathbf{u}_t$ are $3p \times 1$ vectors defined by

$$\mathbf{y}_t = [\; \mathbf{g}_t' - \boldsymbol{\mu}_g' \quad \mathbf{g}_{t-1}' - \boldsymbol{\mu}_g' \quad \cdots \quad \mathbf{g}_{t-p+1}' - \boldsymbol{\mu}_g' \;]', \tag{11}$$

$$\mathbf{u}_t = [\; \mathbf{e}_t' \quad \mathbf{0}_3' \quad \cdots \quad \mathbf{0}_3' \;]', \tag{12}$$

$\boldsymbol{\mu}_g$ is the mean of $\mathbf{g}_t$ and $\mathbf{C}$ is a suitable $3p \times 3p$ matrix (involving the matrices $\mathbf{B}_i$, $i = 1, \ldots, p$). The Wold representation of $\mathbf{y}_t$ reads $\mathbf{y}_t = \sum_{i=0}^{\infty} \mathbf{C}^i \mathbf{u}_{t-i}$. Denoting by $\mathbf{D}_i$ the 3×3 upper-left block of the matrix $\mathbf{C}^i$ and defining the matrix $\mathbf{F}_i = \mathbf{D}_i\mathbf{A}_0^{-1}$, we can express $\mathbf{g}_t$ as $\mathbf{g}_t = \boldsymbol{\mu}_g + \sum_{i=0}^{\infty} \mathbf{F}_i \boldsymbol{\varepsilon}_{t-i}$. The third element of the vector $\mathbf{g}_t$ is the oil price change proxy denoted by $g^{\mathtt{P}}$. Hence, we obtain the following decomposition of $g^{\mathtt{P}}$ into three components

$$g_t^{\mathtt{P}} = \mu_g^{\mathtt{P}} + x_t^{\mathtt{S}} + x_t^{\mathtt{D}} + x_t^{\mathtt{OSD}}, \tag{13}$$

where $\mu_g^{\mathtt{P}}$ is the mean of the oil price change $g_t^{\mathtt{P}}$, $x_t^{\mathtt{S}} = \sum_{i=0}^{\infty}(\mathbf{F}_i)_{31}\varepsilon_{t-i}^{\mathtt{S}}$ is the oil supply shock, $x_t^{\mathtt{D}} = \sum_{i=0}^{\infty}(\mathbf{F}_i)_{32}\varepsilon_{t-i}^{\mathtt{D}}$ is the global demand shock, and $x_t^{\mathtt{OSD}} = \sum_{i=0}^{\infty}(\mathbf{F}_i)_{33}\varepsilon_{t-i}^{\mathtt{OSD}}$ is the oil-specific demand shock.

We estimate the VAR model in equation (9) using OLS and order $p = 2$, as selected according to BIC, and obtain the decomposition in equation (13) in a real-time fashion. Specifically, for each month in the period between January 1986 and December 2015, we estimate the VAR model using all available data starting in February 1982 and ending in that month. Then, we obtain the time series of three shocks in the decomposition (13), but keep the vector of the oil supply, global demand, and oil-specific demand shocks only for the last month. Our approach of the real-time decomposition reflects all revisions to historical data of crude oil and crude steel production.[13] We plot the time series of the oil supply, global demand, and oil-specific demand shocks from January 1986 to December 2015 in Figure 6. All three series are rescaled so that they have standard deviation equal to one. Supply shocks become less volatile after 2004 while global demand shocks become more volatile after 2007. The volatility of the oil-specific demand shocks appears stable across the sample period.

5.2 The predictive power of oil supply, global demand, and oil-specific demand shocks

In this section, we examine the ability of the three shocks obtained by the oil price change decomposition (13) to forecast next-month MSCI index excess returns over the sample period from January 1986 to December 2015. We do so by running the following predictive regression:

$$r_{t+1}^{\mathtt{e}} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{D}} x_t^{\mathtt{D}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_{t+1}^{\mathtt{DEC}}, \tag{14}$$

[13] We have also obtained the full-sample version of the decomposition in equation (13) using all data from February 1982 to December 2015. The two versions resulted in very similar results in the subsequent analysis, indicating that the impact of the VAR estimation error is minimal.

where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, oil-specific demand shocks obtained in the decomposition (13), respectively. We gauge the forecasting ability of the three shocks using measures of both statistical and economic significance as in the previous section. Furthermore, we offer comparisons between model (7) that uses oil price change as the sole predictor and the decomposition-based model (14).

The three shocks identified by the decomposition (13) are anticipated to have different impacts on future equity returns. Under the assumption of inelastic demand for oil in the short run, a disruption in oil production would result in an oil price increase. This would be potentially bad news for the real economy and the stock market while the corresponding shock $x^{\mathtt{S}}$ would be positive. Hence, one expects $\beta^{\mathtt{S}}$ to be negative in the predictive regression (14). Second, positive global demand shocks stimulate the global economy as a whole, although the impact might differ across countries. One, therefore, expects a positive global demand shock to be good news for equity markets. At the same time, a positive global demand shock could drive up the price of oil, which, in turn, could have a slowing down effect on certain economies. However, the overall effect should be dominated by the first direct impact and, hence, one expects a positive slope $\beta^{\mathtt{D}}$ in the predictive regression (14). Third, following the interpretation in Kilian (2009), an oil-specific demand shock is thought to capture changes in the demand for oil driven by precautionary motives. Accordingly, a positive oil-specific demand shock is thought to originate from the increased demand for oil due to uncertainty regarding future availability of oil and so it is perceived to be bad news for the global economy and the stock market. Hence, one expects $\beta^{\mathtt{OSD}}$ to be negative in the predictive regression (14).

In Table 7, we present statistical significance results for the predictive regression (14) over the 1986.01–2015.12 sample period using MSCI index excess returns, denominated in both local currencies and US dollars. To provide a direct comparison between model (7), which uses oil price change as the only predictor, and model (14), we also estimate model (7) over the same sample period. The standard errors are computed using the Newey and West (1987) method with optimal bandwidth selected as in Newey and West (1994).

As expected, given the evidence presented in Section 4, the forecasting power of oil price changes is diminished over 1986.01–2015.12 sample period. The results are very similar to the ones presented in Table 4 corresponding to the 1983.01–2015.12 sample period. In particular, the adjusted R-square for both the World and the US index is 0.5%. In stark contrast, we find strong evidence of predictability using the decomposition-based model (14).

For local-currency denominate returns, the adjusted R-squares for the Canada, France, Germany, Italy, Japan, and UK MSCI indexes are 1.5%, 2.9%, 3.1%, 9.5%, 1.9%, and 2.3%, respectively. The point estimates of the slope coefficient $\beta^{\mathtt{D}}$ are positive in all six cases and statistically significant in four (five) cases at the 5% (10%) level of significance. The point estimates of the slope coefficient $\beta^{\mathtt{OSD}}$ are negative in all six cases and statistically significant in four (five) cases at the 5% (10%) level of significance. The estimates of the slope coefficient $\beta^{\mathtt{S}}$ are negative in four out of six cases, although statistically significant only in one case.

For US-dollar denominated returns, the results are very similar. The adjusted R-squares for the Canada, France, Germany, Italy, Japan, UK, USA, and World MSCI indexes are 1.5%, 1.5%, 1.9%, 6.6%, 0.7%, 2.1%, 4.3%, and 3.6%, respectively. The $\beta^{\mathtt{D}}$ point estimates are positive in all eight cases and statistically significant in four (six) cases at the 5% (10%) level of significance. The $\beta^{\mathtt{OSD}}$ point estimates are negative in all eight cases and statistically significant in five (six) cases at the 5% (10%) level of significance. The estimates of the slope coefficient $\beta^{\mathtt{S}}$ are negative in six out of eight cases, although statistically significant only in one case.

We repeat the above analysis computing standard errors according to the Hodrick (1992) method. The results, reported in Table 8, are similar and convey the same message. Collectively, we conclude that there is strong statistical evidence supporting the usefulness of the decomposition (13) and the ability of the three associated shocks to forecast the World and G7 country MSCI index excess returns.

In addition to the evidence on statistical significance, we also provide evidence on the economic significance of the ability of the oil supply, global demand, and oil-specific demand shocks to predict G7 country MSCI index returns. We refer the model described by the decomposition-based predictive regression (14) as the augmented model and compare it to three baseline models. The first baseline model assumes that the MSCI index excess return $r^{\mathtt{e}}_{t+1}$ is i.i.d. The second baseline model is described by the predictive regression (7) that uses the oil price change $g^{\mathtt{P}}$ as predictor. The third baseline model is described by the predictive regression

$$r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{D}} + \delta^{\mathtt{D}} g^{\mathtt{D}}_t + u^{\mathtt{D}}_{t+1}, \tag{15}$$

which uses the global demand growth proxy $g^{\mathtt{D}}$ as predictor.

Let $\mathtt{CER}_{\mathtt{DEC}}$ and $\mathtt{SR}_{\mathtt{DEC}}$ denote the CER and SR achieved by the trading strategy using the decomposition-based predictive regression model (14). Analogously, we denote by $\mathtt{CER}_{\mathtt{IID}}$, $\mathtt{CER}_{\mathtt{P}}$, and $\mathtt{CER}_{\mathtt{D}}$ ($\mathtt{SR}_{\mathtt{IID}}$, $\mathtt{SR}_{\mathtt{P}}$, and $\mathtt{SR}_{\mathtt{D}}$) the CERs (SRs) achieved by the trading strategies using the i.i.d. model, the predictive regression (7), and the predictive regression (15), respectively. To gauge the predictive ability of the shocks $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$, we test the null hypotheses $H^{\mathtt{CER}}_0$: $\mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{DEC}}$, $H^{\mathtt{CER}}_0$: $\mathtt{CER}_{\mathtt{P}} = \mathtt{CER}_{\mathtt{DEC}}$, and $H^{\mathtt{CER}}_0$: $\mathtt{CER}_{\mathtt{D}} = \mathtt{CER}_{\mathtt{DEC}}$ against their one-sided alternatives. Furthermore, in a similar fashion, we test the null hypotheses $H^{\mathtt{SR}}_0$: $\mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{DEC}}$, $H^{\mathtt{SR}}_0$: $\mathtt{SR}_{\mathtt{P}} = \mathtt{SR}_{\mathtt{DEC}}$, and $H^{\mathtt{SR}}_0$: $\mathtt{SR}_{\mathtt{D}} = \mathtt{SR}_{\mathtt{DEC}}$ against their one-sided alternatives. The economic significance test results are reported in Table 9.

For local-currency denominated index returns, the decomposition-based model (14) generates CERs that are higher than their counterparts generated by the i.i.d. model in all six cases. The difference is sizable, e.g., more than 3.76%, in annualized terms, for Japan and the UK, and statistically significant in five out of six cases at the 10% level of significance. Moreover, the decomposition-based model (14) generates CERs that are higher than their counterparts generated by model (7) based on oil price change in five out of six cases, with the exception of France. In the remaining cases, the difference is higher than 1.2%, in annualized terms, and statistically significant in the case of Japan at the 10% level of significance. The decomposition-based model (14) also performs substantially better than the model (15) based on global demand growth in terms of CER. It produces CERs that are higher in all six cases and statistically significant in four out of six cases at the 10% level of significance. The SR results are in line with the CER results. The decomposition-based model (14) generates SRs that are higher than their counterparts generated by the i.i.d. model in all six cases. The increase in SR is sizable, e.g., from 0.19 to 0.53 for the UK, and the difference is statistically significant in four out of six cases at the 5% level of significance. Moreover, the decomposition-based model (14) generates SRs that are at least as high as their counterparts generated by the model (7) based on oil price change in all six cases. The differences again are sizable and statistically significant in the case of Japan at the 5% level of significance. The decomposition-based model (14) also performs substantially better than the model (15) based on global demand growth in terms of SR. It produces SRs that are higher in all six cases and statistically significant in five out of six cases at the 10% level of significance.

The results for US-dollar denominated index returns convey the same message. Importantly, in the case of the USA MSCI index, the decomposition-based model (14) generates an annualized CER equal to 9.28% compared to 6.52%, 6.08%, and 6.20% generated by the i.i.d. model, model (7), and model (15), respectively. The corresponding p-values are 0.11, 0.08, and 0.05, respectively. Even stronger results are obtained for the World MSCI index. The decomposition-based model (14) generates an annualized CER equal to 7.90% compared to 4.03%, 3.88%, and 3.30% generated by the i.i.d. model, model (7), and model (15), respectively. The difference is statistically significant in all three comparisons with p-values equal to 0.04, 0.07, and 0.03, respectively. Strong results are obtained in terms of SR as well. In the case of the USA MSCI index, the decomposition-based model (14) generates an annualized SR equal to 0.65 compared to 0.48, 0.45, and 0.46 generated by the i.i.d. model, model (7), and model (15), respectively. The corresponding p-values are 0.12, 0.09, and 0.05, respectively. For the World MSCI index, the decomposition-based model (14) generates an annualized SR equal to 0.56 compared to 0.31, 0.30, and 0.27 generated by the i.i.d. model, model (7), and model (15), respectively. The difference is statistically significant in all three comparisons with p-values equal to 0.04, 0.06, and 0.03, respectively.

So far, in this subsection, we have provided strong statistical as well as economic significance in support of the ability of the oil supply, global demand, and oil-specific demand shocks to predict the World and G7 country MSCI index excess returns. Next, we provide further corroborating evidence on the stability of this predictive relation between MSCI index excess returns and these three shocks. As in the case on the predictive regression model (7), we use the Bai and Perron (2003) methodology to test for structural breaks in the decomposition-based model (14). The results are presented in Table 10. The test does not identify breaks in all 14 cases considered, 6 local-currency and 8 US-dollar denominated MSCI indexes. Hence, the predictive regression model (14) appears to be a stable and robust specification illustrating the importance of disentangling oil price changes into oil supply, global demand, and oil-specific demand shocks.

In the next subsection, we examine various aspects of the relationship between these three shocks embedded in oil price changes and future stock returns. In particular, we provide evidence suggesting that the documented predictability does not appear to be consistent with time-varying risk premia.

5.3 Additional evidence and robustness checks

First, to alleviate any concerns regarding the real-time availability of the data required to obtain the oil price change decomposition (13), we estimate the predictive regression $r^{\mathtt{d}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{d}}_{t+1}$ is the monthly (simple) return on the World or a G7 country MSCI index constructed with a delay of one or two weeks.[14] The results, reported in Table 11, illustrate the robustness of the forecasting ability of oil supply, global demand, and oil-specific demand shocks with respect to one- or two-week delays in the construction of the World and G7 country MSCI index excess returns.

When studying return predictability, one natural question that emerges is whether the predictors under examination can forecast asset returns over long horizons. In Table 12, we present statistical evidence on whether the oil supply, global demand, and oil-specific demand shocks obtained in the decomposition (13) can predict three-month and six-month G7 country MSCI equity returns. The evidence suggests that the predictive ability of the three shocks gradually diminishes as the horizon gets longer. In particular, at the 5% level of significance, the slope corresponding to the oil-specific demand shock is statistically significant only for Italy. Moreover, the statistical significance of the slope corresponding to the global demand shock is reduced as we move from the three-month to the six-month horizon. This evidence is reinforced by the adjusted R-squares over two-, three-, four-, five-, and six-month horizons that we report in Table 13. The adjusted R-squares exhibit a declining pattern as the horizon increases and they are less than

[14] We do not use excess returns for this exercise due to lack of availability of interest rate data for the relevant time periods.

2% for eight out of 14 cases for six-month horizon. As argued by Fama and French (1989) and Driesprong, Jacobsen, and Maat (2008), among others, predictability typically associated with time-varying risk premia is long-lived and persists over long horizons. We document that this is not the case for the oil supply, global demand, and oil-specific demand shocks and, hence, we conclude that the documented predictability is not consistent with time-varying risk premia.

We also examine whether the results on the predictability of aggregate equity index returns are robust in the cross section of US industries. Specifically, we use the 17 Fama-French value-weighted industry portfolios.[15] First, we conduct Bai and Perron (2003) structural break tests for (i) the predictive regression using oil price change as the predictor, for the 1983.01–2015.12 sample period, and (ii) the predictive regression using the oil supply, global demand, and oil-specific demand shocks as predictors, for the 1986.01–2015.12 sample period. Table 15 shows that the tests identify the presence of one structural break in 14 of 17 industry portfolios when oil price change is used as the sole predictor, with the exception of Mining and Minerals, Oil and Petroleum Products, and Utilities. In contrast, the tests do not identify a break for any industry when the three shocks embedded in oil price changes are used as predictors. Second, we examine the ability of the three shocks to forecast industry portfolio excess returns. The results are presented in Table 16, where we also provide the results of the predictive regression using oil price change as the sole predictor. Given the results of the structural break tests discussed before, the results of the oil price change regression are meaningful only for the three industries that do not exhibit a break. As expected, the estimated predictive slope on oil price change is positive for the Oil and Petroleum Products industry, although not statistically significant. Overall, there is no evidence of predictability based on oil price change alone, with 15 out of 17 adjusted R-squares being less than 1%. In contrast, there is strong evidence of predictability across the various industries based on the three shocks embedded in oil price changes, according to Newey and West (1987) standard errors. The $\beta^{\mathtt{D}}$ estimates are positive for all 17 industries and statistically significant for 12 (13) industries at the 5% (10%) level. The $\beta^{\mathtt{OSD}}$ estimates are

[15] We use monthly returns on the 17 Fama-French value-weighted industry portfolios from Kenneth French's website. The abbreviations (descriptions) of the 17 industries are Food (Food), Mines (Mining and Minerals), Oil (Oil and Petroleum Products), Clths (Texiles, Apparel and Footware), Durbl (Consumer Durables), Chems (Chemicals), Cnsum (Drugs, Soap, Perfumes, Tobacco), Cnstr (Construction and Construction Materials), Steel (Steel Works Etc), FabPr (Fabricated Products), Machn (Machinery and Business Equipment), Cars (Automobiles), Trans (Transportation), Utils (Utilities), Rtail (Retail Stores), Finan (Banks, Insurance Companies, and Other Financials), and Other (Other).

negative for 15 industries and statistically significant for seven (eight) industries at the 5% (10%) level. The $\beta^{\mathtt{S}}$ estimates are negative for all 17 industries, although statistically significant only for four industries at the 10% level. Moreover, the adjusted R-square is higher than 1.5% for 12 out of 17 industries. Overall, these results demonstrate that the predictive ability of the oil supply, global demand, and oil-specific demand shocks is strong not only for the aggregate equity index, but also across different industry portfolios.

Another natural question in the context of equity return predictability is how the proposed predictors relate to macroeconomic variables that have been extensively used in the extant literature to model time-varying expected equity returns. Due to data limitations, we examine this issue only for the case of the US. Table 17 presents the contemporaneous correlations between the oil supply, global demand, and oil-specific demand shocks and four macroeconomic variables: the log dividend yield, the term spread, the default yield spread, and the one-month T-bill rate. Results are reported for both the real-time and full-sample decompositions. The correlations are rather low in magnitude with the largest (in absolute value) being the correlation between the global demand shock and the default yield equal to -0.17 (-0.16) for the real-time (full-sample) decomposition.

In addition, we examine whether the forecasting ability of the oil supply, global demand, and oil-specific demand shocks is robust to the presence of the macroeconomic predictors in the case of the US. We examine the following linear predictive regressions $r^{\mathtt{e}}_{t+1} = \gamma^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + \boldsymbol{\theta}'\mathbf{z}_t + v^{\mathtt{P}}_{t+1}$ and $r^{\mathtt{e}}_{t+1} = \gamma^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + \boldsymbol{\lambda}'\mathbf{z}_t + v^{\mathtt{DEC}}_{t+1}$ over various sample periods, where $r^{\mathtt{e}}_{t+1}$ is the USA MSCI index excess return and $\mathbf{z}_t$ is the vector of the four macroeconomic variables mentioned above. The results are reported in Table 18. According to the evidence, the inference results we have reported so far in the paper are robust to the presence of the macroeconomic variables. In particular, the forecasting ability of oil price change over the early 1982.01–2003.04 sample period is unaffected. Moreover, over the 1986–2015 sample period, the slopes of the global demand shock and the oil-specific demand shock are negative and positive, respectively, and significant at the 5% level of significance.

In our next empirical exercise, we examine the descriptive statistics of the conditional expected excess returns based on the oil supply, global demand, and oil-specific demand shocks. In particular, we focus on the mean, the standard deviation, and the first three autocorrelations. This evidence can shed light on the issue of whether the documented predictive ability of the three shocks is consistent with time-varying risk premia. For the purposes of comparison, we use two benchmarks. The first benchmark is the predicted MSCI USA index excess return based on the four macroeconomic variables discussed above in terms of descriptive statistics. We report the results, for the time period 1986.01–2015.12, in Table 19. Overall, the predicted expected excess returns based on the three shocks are more volatile and much less persistent compared to their analogues obtained from the macroeconomic variables. One might argue that the predicted excess returns based on the macroeconomic variables are just too persistent, given the nature of these macroeconomic predictors. To address this concern, in our second comparison, we use as benchmark the equity risk premium estimates obtained by Martin (2017) based on option prices over different maturities, ranging from one month to one year.[16] We report the results for the time period 1996.01–2012.01, over which the estimates from Martin (2017) are available, in Table 20. The main message from the second comparison remains the same. In particular, the second and third order autocorrelations of the predicted expected excess returns based on the three shocks are much lower than their counterparts obtained from either the predicted excess returns based on the macroeconomic predictors or the risk premium estimates of Martin (2017). Collectively, this evidence suggests that the forecasting ability of the three shocks is not consistent with time-varying risk premia, in line with the evidence of predictability diminishing over longer horizons reported above.

We conclude this section by investigating whether there is a more direct link between time variation in expected returns and changes in risk, as captured by return volatility. Such an exercise can shed more light to the question of whether the predictive ability of the oil supply, global demand, and oil-specific demand shocks is associated with changes in risk premia. To this end, we employ an augmented EGARCH(1,1) model that includes these three shocks in the volatility equation as exogenous regressors. If the variation of expected returns is to be attributed to time-varying risk premia, we would expect that any of these three shocks has the same effect on both the drift and the volatility. In the context of the EGARCH model, we would be expect the coefficient on any of these shocks to have the same sign as in the drift equation and be

[16] We thank Ian Martin for making the data available on his website.

statistically significant. The econometric specification is:

$$r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}, \tag{16}$$

$$u^{\mathtt{DEC}}_{t+1} = \sigma_t z_{t+1}, \quad z_{t+1} \sim \text{i.i.d.}(0,1), \tag{17}$$

$$\log(\sigma^2_t) = \tau_0 + \tau_1|z_t| + \tau_2 z_t + \tau_3 \log(\sigma^2_{t-1}) + \zeta^{\mathtt{S}} x^{\mathtt{S}}_t + \zeta^{\mathtt{D}} x^{\mathtt{D}}_t + \zeta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t. \tag{18}$$

As argued above, time-varying risk premia would be consistent with $\zeta^{\mathtt{S}} < 0$, $\zeta^{\mathtt{D}} > 0$, and $\zeta^{\mathtt{OSD}} < 0$. We estimate the model using monthly excess returns on the MSCI indexes for the G7 countries as well as the World MSCI index, denominated both in local currencies and US dollars, over the 1986.01–2015.12 sample period. We consider three distributions for the disturbances z_{t+1}: Normal, Student-t, and GED. The Student-t distribution was selected according to the Bayesian Information Criterion.[17] The results are presented in Table 21. In the majority of the cases, the estimates $\zeta^{\mathtt{S}} < 0$, $\zeta^{\mathtt{D}} > 0$, and $\zeta^{\mathtt{OSD}} < 0$ are statistically insignificant at conventional levels. More importantly, whenever there is significance, the sign is the opposite of what would be consistent with time-variation of risk premia, e.g., in five out of 14 instances the estimates of $\zeta^{\mathtt{D}}$ are statistically significant but negative. This evidence is inconsistent with the notion of time-varying risk premia, reinforcing the message of the evidence documented earlier.

6 Conclusion

As the modern global economy heavily depends on oil, the price of oil is widely thought to affect global real economic activity and consequently the global equity market. An oil price drop has been considered in the past to be good news as it lowers the cost of production in a significant number of sectors and allows consumers to boost their consumption. Accordingly, one could conjecture that negative (positive) oil price changes should predict higher (lower) subsequent equity returns. Driesprong, Jacobsen, and Maat (2008) document that this is indeed the case for a large number of MSCI equity indexes based on data until 2003. However, this predictive relationship has dramatically changed over the last ten years. Specifically, the correlation between the World MSCI index return and the lagged one-month log growth rate of West Texas Intermediate spot

[17] The results are very similar across all three distributional assumptions.

price has increased from -0.22 over the 1982–2003 period to 0.26 over the 2004–2015 period. As a result, the ability of oil price change to forecast future equity returns has diminished over the sample period extending to 2015. Furthermore, using the formal econometric test of Bai and Perron (2003), we detect a structural break in the predictive relationship in the third quarter of 2008 for most of the G7 country MSCI index returns.

In this paper, we suggest that oil price changes do contain useful information for forecasting subsequent equity indexes, provided that these changes are suitably disentangled into supply and demand shocks. Using a variant of the structural VAR approach of Kilian (2009), we obtain an oil price change decomposition into an oil supply shock, a global demand shock, and an oil-specific demand shock and argue that these three different types of shocks should have different effects on equity markets. The conjecture that oil supply shocks and oil-specific demand shocks (global demand shocks) predict equity returns with a negative (positive) slope is supported by the empirical evidence over the 1986-2015 sample period. Using the oil price decomposition instead of just oil price change increases the annualized certainty equivalent return and Sharpe ratio of a mean-variance trading strategy for the World MSCI index from 3.88% to 7.90% and from 0.30 to 0.56, respectively, with the differences being statistically significant. Importantly, we detect no structural breaks in the predictive relationship between equity index returns and the three shocks in any of the 14 MSCI equity indexes that we consider. These results survive in the presence of traditional macroeconomic predictors for the case of the USA MSCI index and, in general, do not appear to be consistent with time-varying risk premia.

References

Andrews, D., 1991, "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59, 817–858.

Andrews, D. W. K., and J. C. Monahan, 1992, "An Improved Heteroscedasticity and Autocorrelation Consistent Covariance Matrix Estimator," *Econometrica*, 60, 953–966.

Apergis, N., and S. M. Miller, 2009, "Do structural oil-market shocks affect stock prices?," *Energy Economics*, 31, 569–575.

Bai, J., and P. Perron, 2003, "Computation and analysis of multiple structural change models," *Journal of Applied Econometrics*, 18, 1–22.

Bernanke, B., 2016, "The Relationship between Stocks and Oil Prices," Article, The Brookings Institution.

Campbell, J., and S. Thompson, 2008, "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies*, 21, 1509–1531.

Casassus, J., and F. Higuera, 2012, "Short-horizon return predictability and oil prices," *Quantitative Finance*, 12, 1909–1934.

Chen, N., R. Roll, and S. Ross, 1986, "Economic Forces and the Stock Market," *Journal of Business*, 59, 383–403.

Chiang, I.-h. E., W. K. Hughen, and J. S. Sagi, 2015, "Estimating Oil Risk Factors Using Information from Equity and Derivatives Markets," *Journal of Finance*, 70, 769–804.

Christoffersen, P., and X. Pan, 2014, "Oil Volatility Risk and Expected Stock Returns," Working Paper, University of Toronto and Tulane University.

Driesprong, G., B. Jacobsen, and B. Maat, 2008, "Striking oil: Another puzzle?," *Journal of Financial Economics*, 89, 307–327.

Fama, E. F., and K. R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.

Garlappi, L., G. Skoulakis, and J. Xue, 2016, "Reassessing the Role of Industries in Predicting US Aggregate Market Returns," Working Paper, University of British Columbia and University of Maryland.

Hamilton, J. D., 2003, "What is an oil shock?," *Journal of Econometrics*, 113, 363–398.

Hodrick, R., 1992, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5, 357–386.

Jones, C., and G. Kaul, 1996, "Oil and the Stock Markets," *Journal of Finance*, 51, 463–491.

Kilian, L., 2009, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053–1069.

Kilian, L., and C. Park, 2009, "The impact of oil price shocks on the U.S. stock market," *International Economic Review*, 50, 1267–1287.

Martin, I., 2017, "What is the Expected Return on the Market?," *Quarterly Journal of Economics*, 132, 367–433.

Narayan, P. K., and R. Gupta, 2015, "Has oil price predicted stock returns for over a century?," *Energy Economics*, 48, 18–23.

Newey, W., and K. West, 1994, "Automatic Lag Selection in Covariance Matrix Estimation," *Review of Economic Studies*, 61, 631–653.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Rapaport, A., 2014, "Supply and Demand Shocks in the Oil Market and their Predictive Power," Working Paper, University of Chicago.

Ravazzolo, F., and J. L. Vespignani, 2015, "A New Monthly Indicator of Global Real Economic Activity," Working Paper, BI Norwegian Business School and University of Tasmania.

Ready, R. C., 2016, "Oil Prices and the Stock Market," *Review of Finance*, forthcoming.

A Appendices

A.1 Data construction

The single oil price change proxy g^{P} is constructed in a real-time fashion using PCA. Specifically, for each month t between January 1983 and December 2015, we use data on three proxies for oil price change starting in February 1982 and ending in month t. We first rescale the three log growth rates, obtained from the West Texas Intermediate, the Dubai, and the Arab Light spot prices, so they all have variance equal to one over the given sample period and then perform PCA. The first PCA corresponding to month t is kept each time and the process is repeated using expanding windows until December 2015 is reached.

To address the strong seasonality of the global crude steel production data, we use X-13ARIMA-SEATS to compute seasonally adjusted level data from which we compute log growth rates in a real-time fashion.[18] Specifically, for each month in the period between February 1982 and December 2015, we perform seasonal adjustment on the level data starting in January 1968 and ending in that month, compute the log growth rates of the seasonally adjusted level data, and finally keep the log growth rate over the last month.

The single global demand growth proxy g^{D} is also constructed in a real-time fashion using PCA. Specifically, for each month t between January 1983 and December 2015, we use data on two proxies for global economic activity starting in February 1982 and ending in month t. We first rescale the two log growth rates, obtained from the shipping cost index and the global crude steel production data, so they all have variance equal to one over the given sample period and then perform PCA. The first PCA corresponding to month t is kept each time and the process is repeated using expanding windows until December 2015 is reached.

[18] We use the X-13 Toolbox for Matlab, written by Yvan Lengwiler, to perform seasonal filtering. The source codes are retrieved from `http://www.mathworks.com/matlabcentral/fileexchange/49120-x-13-toolbox-for-seasonal-filtering/content/x13tbx/x13.m`.

A.2 Testing for equality in certainty equivalent return and Sharpe ratio

In this appendix, we follow Garlappi, Skoulakis, and Xue (2016) who provide asymptotic tests for comparing two models and the corresponding trading strategies in terms of their certainty equivalent returns (CER) and Sharpe ratios (SR). Model 1 is the baseline model while Model 2 is the augmented one. The null hypothesis of interest is that the CERs or SRs obtained by two models are equal. Denote by $\texttt{CER}_j$ and $\texttt{SR}_j$ the CER and SR of Model j, for $j = 1, 2$. The purpose of these tests is to evaluate the increment value of the augmented model in terms of performance and, therefore, we naturally consider one-sided alternative hypotheses. Specifically, we test (i) $H_0^{\mathsf{CER}} : \texttt{CER}_1 = \texttt{CER}_2$ against $H_A^{\mathsf{CER}} : \texttt{CER}_1 < \texttt{CER}_2$ and (ii) $H_0^{\mathsf{SR}} : \texttt{SR}_1 = \texttt{SR}_2$ against $H_A^{\mathsf{SR}} : \texttt{SR}_1 < \texttt{SR}_2$. Note that, in the context of our mean-variance framework, the CER of the portfolio is expressed as a function of the first two moments of simple portfolio returns while the SR of a portfolio is expressed as a function of the first two moments of the portfolio excess returns.

Let $\mathbf{r}_t = (r_{1,t}, r_{2,t})'$ denote the pair of returns on the two portfolios at time t. These returns could be either simple or excess depending on whether we focus on the CER or the SR. Denote the mean, variance and noncentral second moment of $r_{j,t}$ by μ_j, σ_j^2, and ν_j, respectively for the portfolios $j = 1, 2$. Note that $\sigma_j^2 = \nu_j - \mu_j^2$. It follows that, in the case of simple returns, the CERs for an investor with mean-variance preferences and risk aversion coefficient equal to γ are given by $\texttt{CER}_j = \mu_j - \frac{1}{2}\gamma\left(\nu_j - \mu_j^2\right)$, $j = 1, 2$. Similarly, in the case of excess returns, the SRs are given by $\texttt{SR}_j = \frac{\mu_j}{\sqrt{\nu_j - \mu_j^2}}$, $j = 1, 2$. Therefore, the relevant hypotheses can be stated using a suitable function of the parameter vector $\boldsymbol{\theta} = (\mu_1, \mu_2, \nu_1, \nu_2)'$. We estimate $\boldsymbol{\theta}$ by the sample analogue $\widehat{\boldsymbol{\theta}} = (\widehat{\mu}_1, \widehat{\mu}_2, \widehat{\nu}_1, \widehat{\nu}_2)'$, where $\widehat{\mu}_j = \frac{1}{T}\sum_{t=1}^{T} r_{j,t}$ and $\widehat{\nu}_j = \frac{1}{T}\sum_{t=1}^{T} r_{j,t}^2$, for $j = 1, 2$. Under regularity conditions, such as stationarity and ergodicity, $\widehat{\boldsymbol{\theta}}$ asymptotically follows a normal distribution described by

$$\sqrt{T}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T} \mathbf{y}_t \xrightarrow{d} N(\mathbf{0}, \boldsymbol{\Psi}), \tag{19}$$

where $\boldsymbol{\Psi}$ is the long-run variance-covariance matrix of

$$\mathbf{y}_t = \left(r_{1,t} - \mu_1, r_{2,t} - \mu_2, r_{1,t}^2 - \nu_1, r_{2,t}^2 - \nu_2\right)'. \tag{20}$$

The matrix $\boldsymbol{\Psi}$ is given by $\boldsymbol{\Psi} = \boldsymbol{\Gamma}_0 + \sum_{\ell=1}^{\infty} (\boldsymbol{\Gamma}_\ell + \boldsymbol{\Gamma}'_\ell)$, where $\boldsymbol{\Gamma}_\ell = \mathbb{E}\left[\mathbf{y}_t \mathbf{y}'_{t-\ell}\right]$, for $\ell = 0, 1, \ldots$ and is estimated by a heteroscedasticity and autocorrelation consistent (HAC) estimator of the form

$$\widehat{\boldsymbol{\Psi}} = \widehat{\boldsymbol{\Gamma}}_0 + \sum_{\ell=1}^{T} \kappa \left(\frac{\ell}{b_T}\right) (\widehat{\boldsymbol{\Gamma}}_\ell + \widehat{\boldsymbol{\Gamma}}'_\ell), \tag{21}$$

where

$$\widehat{\boldsymbol{\Gamma}}_\ell = \frac{1}{T-\ell} \sum_{t=\ell+1}^{T} \widehat{\mathbf{y}}_t \widehat{\mathbf{y}}'_{t-\ell}, \quad \widehat{\mathbf{y}}_t = \left(r_{1,t} - \widehat{\mu}_1, r_{2,t} - \widehat{\mu}_2, r_{1,t}^2 - \widehat{\nu}_1, r_{2,t}^2 - \widehat{\nu}_2\right)', \tag{22}$$

$\kappa(\cdot)$ is a kernel function, and b_T is the bandwidth. HAC estimators have been developed by several authors including Newey and West (1987), Andrews (1991), Andrews and Monahan (1992), and Newey and West (1994). We report p-values based on the Newey and West (1987) approach with the Bartlett kernel and the optimal bandwidth computed as suggested in Newey and West (1994).

Consider testing the null hypothesis $H_0 : f(\boldsymbol{\theta}) = 0$ against the alternative hypothesis $H_A : f(\boldsymbol{\theta}) < 0$, where $f(\boldsymbol{\theta})$ is a smooth real-valued function of $\boldsymbol{\theta}$. Applying the delta method, we obtain

$$\sqrt{T}\left(f(\widehat{\boldsymbol{\theta}}) - f(\boldsymbol{\theta})\right) \xrightarrow{d} N\left(0, \nabla' f(\boldsymbol{\theta}) \boldsymbol{\Psi} \nabla f(\boldsymbol{\theta})\right), \tag{23}$$

where $\nabla f(\cdot)$ is the gradient of f. For large T, the standard error of $f(\widehat{\boldsymbol{\theta}})$ is given by

$$se(f(\widehat{\boldsymbol{\theta}})) = \sqrt{\frac{1}{T} \nabla' f(\widehat{\boldsymbol{\theta}}) \widehat{\boldsymbol{\Psi}} \nabla f(\widehat{\boldsymbol{\theta}})}, \tag{24}$$

and, therefore, the corresponding t-statistic is $t(f, \widehat{\boldsymbol{\theta}}) = \frac{f(\widehat{\boldsymbol{\theta}})}{se(f(\widehat{\boldsymbol{\theta}}))}$, yielding the one-sided p-value $p(f, \widehat{\boldsymbol{\theta}}) = \Phi(t(f, \widehat{\boldsymbol{\theta}}))$, where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution.

To test for equality of CERs, we use simple returns and the function f takes the form

$$f_{\mathsf{CER}}(\boldsymbol{\theta}) = \left(\mu_1 - \frac{1}{2}\gamma\left(\nu_1 - \mu_1^2\right)\right) - \left(\mu_2 - \frac{1}{2}\gamma\left(\nu_2 - \mu_2^2\right)\right), \tag{25}$$

with gradient equal to

$$\nabla f_{\mathsf{CER}}(\boldsymbol{\theta}) = \left(1 + \gamma\mu_1, -1 - \gamma\mu_2, -\frac{1}{2}\gamma, \frac{1}{2}\gamma\right)'. \tag{26}$$

To test for equality of SRs, we use excess returns and the function f takes the form

$$f_{\mathsf{SR}}\left(\boldsymbol{\theta}\right) = \frac{\mu_1}{\sqrt{\nu_1 - \mu_1^2}} - \frac{\mu_2}{\sqrt{\nu_2 - \mu_2^2}}, \tag{27}$$

with gradient equal to

$$\nabla f_{\mathsf{SR}}(\boldsymbol{\theta}) = \left(\frac{\nu_1}{(\nu_1 - \mu_1^2)^{\frac{3}{2}}}, -\frac{\nu_2}{(\nu_2 - \mu_2^2)^{\frac{3}{2}}}, -\frac{1}{2}\frac{\mu_1}{(\nu_1 - \mu_1^2)^{\frac{3}{2}}}, \frac{1}{2}\frac{\mu_2}{(\nu_2 - \mu_2^2)^{\frac{3}{2}}}\right)'. \tag{28}$$

Table 1: **Oil price change summary statistics.** This table presents summary statistics for three oil spot price log growth rates, i.e., West Texas Intermediate (WTI), Dubai, and Arab Light, and their first principal component g^{P}. Results are presented for the 1983.01–2003.04 and the 1983.01–2015.12 sample periods. All reported numbers are in percentages.

1983.01–2003.04 Sample Period

	WTI	Dubai	Arab Light	g^{P}
Min	−39.60	−37.76	−48.51	−49.85
Max	37.71	53.68	48.73	38.16
Mean	−0.05	−0.10	−0.08	−0.26
Std. dev.	8.17	10.51	10.94	9.18
# of obs.	244	244	244	244

Correlation Matrix

	WTI	Dubai	Arab Light	g^{P}
WTI	1.00	0.76	0.72	0.90
Dubai	0.76	1.00	0.90	0.94
Arab Light	0.72	0.90	1.00	0.92
g^{P}	0.90	0.94	0.92	1.00

1983.01–2015.12 Sample Period

	WTI	Dubai	Arab Light	g^{P}
Min	−39.60	−49.71	−48.51	−49.85
Max	37.71	53.68	48.73	38.16
Mean	0.04	0.02	0.01	−0.09
Std. dev.	8.45	10.16	10.34	9.00
# of obs.	396	396	396	396

Correlation Matrix

	WTI	Dubai	Arab Light	g^{P}
WTI	1.00	0.76	0.72	0.90
Dubai	0.76	1.00	0.91	0.95
Arab Light	0.72	0.91	1.00	0.93
g^{P}	0.90	0.95	0.93	1.00

Table 2: **Oil price change as a predictor of MSCI index excess returns: statistical significance over the 1983.01–2003.04 sample period.** This table presents in-sample statistical significance results for the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where r^{e}_{t+1} is the excess return on an MSCI index and the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. We report the two-sided p-values for the null hypotheses $H_0 : \delta^{\mathsf{P}} = 0$, denoted by NW[$p$] and H[$p$], based on Newey and West (1987) standard errors, with optimal bandwidth selected as in Newey and West (1994), and Hodrick (1992) standard errors, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	δ^{P}	NW[p]	H[p]	$\bar{R}^2$
Local-currency				
Canada	−0.06	0.06*	0.08*	1.0
France	−0.15	0.00***	0.01**	4.9
Germany	−0.17	0.00***	0.01***	4.8
Italy	−0.32	0.00***	0.00***	16.4
Japan	−0.11	0.07*	0.12	2.9
UK	−0.11	0.00***	0.01***	4.1
US-dollar				
Canada	−0.06	0.09*	0.13	0.7
France	−0.13	0.00***	0.02**	3.3
Germany	−0.15	0.00***	0.02**	3.7
Italy	−0.31	0.00***	0.00***	14.3
Japan	−0.09	0.14	0.27	0.9
UK	−0.10	0.00***	0.04**	2.3
USA	−0.12	0.00***	0.00***	5.2
World	−0.12	0.00***	0.01***	5.6

Table 3: **Oil price change as a predictor of MSCI index excess returns: economic significance over the 1983.01–2003.04 sample period.** This table presents evidence on the performance of trading strategies using oil price change as a predictor in terms of two metrics: the certainty equivalent return (CER) and the Sharpe ratio (SR). The investor forms optimal mean-variance portfolios between an MSCI index and the corresponding Treasure bill and has a risk aversion coefficient $\gamma = 3$. The weight on the MSCI index is constrained between $\underline{\omega} = 0$ and $\overline{\omega} = 150\%$. We compare the baseline model (IID), according to which the MSCI index excess return r^{e}_{t+1} is i.i.d., to the augmented model (P) described by the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The portfolio strategies use the mean and variance estimates resulting from each model respectively. Reported are the CERs in annualized percentage points and the annualized SRs for both baseline and augmented models. We also report the one-sided p-values for the null hypothesis that the augmented model does not improve the CER and SR obtained from a strategy based on the baseline model. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. The results are based on an expanding window with 60 initial observations.

	CER			SR		
	P	IID	p-value	P	IID	p-value
Local-currency						
Canada	9.65	5.56	0.04**	0.49	0.07	0.02**
France	12.26	6.19	0.05**	0.61	0.28	0.04**
Germany	13.11	4.10	0.01**	0.72	0.22	0.01***
Italy	19.11	6.64	0.00***	0.82	−0.06	0.00***
Japan	−1.22	−6.66	0.06*	0.13	−0.23	0.00***
UK	11.91	5.81	0.00***	0.54	0.08	0.00***
US-dollar						
Canada	7.28	4.20	0.15	0.38	0.12	0.11
France	10.67	5.47	0.06*	0.59	0.36	0.06*
Germany	11.70	3.80	0.02**	0.64	0.23	0.01**
Italy	13.51	2.53	0.00***	0.72	0.07	0.00***
Japan	−3.83	−6.30	0.11	−0.10	−0.25	0.14
UK	8.17	4.12	0.02**	0.45	0.21	0.02**
USA	11.61	7.81	0.11	0.66	0.45	0.08*
World	7.79	2.26	0.06*	0.43	0.19	0.06*

Table 4: **Oil price change as a predictor of MSCI index excess returns: statistical significance over the 1983.01–2015.12 sample period.** This table presents in-sample statistical significance results for the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where r^{e}_{t+1} is the excess return on an MSCI index and the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. We report the two-sided p-values for the null hypotheses $H_0 : \delta^{\mathsf{P}} = 0$, denoted by NW[$p$] and H[$p$], based on Newey and West (1987) standard errors, with optimal bandwidth selected as in Newey and West (1994), and Hodrick (1992) standard errors, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	δ^{P}	NW[p]	H[p]	$\bar{R}^2$
Local-currency				
Canada	−0.01	0.86	0.82	−0.2
France	−0.08	0.08*	0.06*	1.4
Germany	−0.08	0.11	0.10*	1.0
Italy	−0.19	0.01***	0.00***	6.2
Japan	−0.05	0.29	0.34	0.3
UK	−0.06	0.03**	0.03**	1.4
US-dollar				
Canada	0.01	0.78	0.72	−0.2
France	−0.06	0.31	0.22	0.4
Germany	−0.06	0.24	0.24	0.4
Italy	−0.17	0.01**	0.00***	4.1
Japan	−0.04	0.34	0.42	0.1
UK	−0.03	0.49	0.40	0.0
USA	−0.05	0.22	0.14	0.6
World	−0.04	0.22	0.20	0.6

Table 5: **Oil price change as a predictor of MSCI index excess returns: economic significance over the 1983.01–2015.12 sample period.** This table presents evidence on the performance of trading strategies using oil price change as a predictor in terms of two metrics: the certainty equivalent return (CER) and the Sharpe ratio (SR). The investor forms optimal mean-variance portfolios between an MSCI index and the corresponding Treasure bill and has a risk aversion coefficient $\gamma = 3$. The weight on the MSCI index is constrained between $\underline{\omega} = 0$ and $\overline{\omega} = 150\%$. We compare the baseline model (IID), according to which the MSCI index excess return r^{e}_{t+1} is i.i.d., to the augmented model (P) described by the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The portfolio strategies use the mean and variance estimates resulting from each model respectively. Reported are the CERs in annualized percentage points and the annualized SRs for both baseline and augmented models. We also report the one-sided p-values for the null hypothesis that the augmented model does not improve the CER and SR obtained from a strategy based on the baseline model. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. The results are based on an expanding window with 60 initial observations.

	CER			SR		
	P	IID	p-value	P	IID	p-value
Local-currency						
Canada	5.59	5.10	0.39	0.29	0.23	0.37
France	7.08	5.24	0.23	0.44	0.31	0.22
Germany	7.21	4.59	0.18	0.48	0.30	0.15
Italy	9.30	4.50	0.06*	0.49	−0.01	0.00***
Japan	−2.24	−2.83	0.43	0.06	−0.08	0.19
UK	7.93	5.21	0.06*	0.42	0.18	0.04**
US-dollar						
Canada	3.25	3.85	0.60	0.23	0.25	0.55
France	4.78	4.32	0.43	0.38	0.35	0.41
Germany	4.73	4.09	0.42	0.38	0.30	0.30
Italy	3.97	1.86	0.29	0.37	0.09	0.04**
Japan	−2.19	−2.27	0.48	−0.07	−0.13	0.25
UK	3.91	3.69	0.46	0.29	0.27	0.40
USA	7.74	7.06	0.37	0.52	0.48	0.36
World	4.70	3.54	0.30	0.34	0.29	0.36

Table 6: **Bai-Perron structural break tests: oil price changes for the 1983.01–2015.12 sample period.** This table presents Bai and Perron (2003) structural break tests for the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where r^{e}_{t+1} is excess return on an MSCI index and the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. Bold numbers indicate the lowest value assumed by the Bayesian Information Criterion (BIC).

	No break model	One break model		# of breaks
	BIC	BIC	Break date	selected by BIC
Local-currency				
Canada	−6.281	**−6.283**	2008.08	1
France	−5.789	**−5.790**	2008.08	1
Germany	−5.588	**−5.592**	2008.07	1
Italy	−5.514	**−5.554**	2003.10	1
Japan	−5.760	**−5.763**	1990.09	1
UK	**−6.232**	−6.228	2008.08	0
US-dollar				
Canada	−5.810	**−5.815**	2008.07	1
France	**−5.601**	−5.598	2008.07	0
Germany	−5.419	**−5.421**	2008.07	1
Italy	−5.301	**−5.335**	2003.10	1
Japan	−5.570	**−5.575**	1990.09	1
UK	−5.929	**−5.933**	2008.07	1
USA	−6.290	**−6.318**	2008.08	1
World	−6.283	**−6.308**	2008.07	1

Table 7: **Oil supply, global demand, and oil-specific demand shocks as predictors of MSCI index excess returns: statistical significance results, based on Newey and West (1987) standard errors, over the 1986.01–2015.12 sample period.** The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{D}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on Newey and West (1987) standard errors with optimal bandwidth selected as in Newey and West (1994). The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$.

	Oil price change			Oil price change decomposition						
	$\delta^{\mathtt{P}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{D}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local-currency										
Canada	−0.00	0.91	−0.3	−0.18	0.60	0.27	0.01***	−0.03	0.39	1.5
France	−0.08	0.09*	1.5	−0.43	0.39	0.23	0.06*	−0.10	0.01**	2.9
Germany	−0.08	0.12	1.0	−0.50	0.39	0.31	0.02**	−0.11	0.01**	3.1
Italy	−0.18	0.01**	6.2	−0.75	0.07*	0.32	0.03**	−0.22	0.00***	9.5
Japan	−0.05	0.32	0.3	0.21	0.56	0.31	0.00***	−0.08	0.06*	1.9
UK	−0.06	0.03**	1.4	0.17	0.70	0.16	0.13	−0.09	0.00***	2.3
US-dollar										
Canada	0.02	0.74	−0.2	−0.26	0.56	0.35	0.01***	−0.01	0.82	1.5
France	−0.06	0.30	0.4	−0.57	0.21	0.23	0.24	−0.07	0.08*	1.5
Germany	−0.06	0.26	0.3	−0.66	0.20	0.30	0.10*	−0.08	0.04**	1.9
Italy	−0.16	0.02**	4.1	−0.91	0.07*	0.32	0.12	−0.20	0.00***	6.6
Japan	−0.04	0.37	0.1	0.61	0.18	0.19	0.06*	−0.07	0.11	0.7
UK	−0.03	0.49	0.0	0.01	0.98	0.32	0.05**	−0.06	0.04**	2.1
USA	−0.04	0.27	0.5	−0.36	0.36	0.32	0.00***	−0.07	0.01**	4.3
World	−0.04	0.26	0.5	−0.19	0.58	0.31	0.01**	−0.07	0.02**	3.6

Table 8: **Oil supply, global demand, and oil-specific demand shocks as predictors of MSCI index excess returns: statistical significance results, based on Hodrick (1992) standard errors, over the 1986.01–2015.12 sample period.** The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{D}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by H[p], based on Hodrick (1992) standard errors. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$.

	Oil price change			Oil price change decomposition						
	$\delta^{\mathtt{P}}$	H[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	H[p]	$\beta^{\mathtt{D}}$	H[p]	$\beta^{\mathtt{OSD}}$	H[p]	$\bar{R}^2$
Local-currency										
Canada	−0.00	0.87	−0.3	−0.18	0.64	0.27	0.02**	−0.03	0.38	1.5
France	−0.08	0.07*	1.5	−0.43	0.40	0.23	0.10*	−0.10	0.03**	2.9
Germany	−0.08	0.11	1.0	−0.50	0.41	0.31	0.07*	−0.11	0.02**	3.1
Italy	−0.18	0.00***	6.2	−0.75	0.15	0.32	0.03**	−0.22	0.00***	9.5
Japan	−0.05	0.36	0.3	0.21	0.67	0.31	0.03**	−0.08	0.13	1.9
UK	−0.06	0.04**	1.4	0.17	0.70	0.16	0.14	−0.09	0.01***	2.3
US-dollar										
Canada	0.02	0.67	−0.2	−0.26	0.56	0.35	0.03**	−0.01	0.82	1.5
France	−0.06	0.22	0.4	−0.57	0.30	0.23	0.26	−0.07	0.11	1.5
Germany	−0.06	0.26	0.3	−0.66	0.30	0.30	0.23	−0.08	0.09*	1.9
Italy	−0.16	0.01***	4.1	−0.91	0.11	0.32	0.11	−0.20	0.00***	6.6
Japan	−0.04	0.44	0.1	0.61	0.28	0.19	0.17	−0.07	0.21	0.7
UK	−0.03	0.40	0.0	0.01	0.98	0.32	0.03**	−0.06	0.10*	2.1
USA	−0.04	0.17	0.5	−0.36	0.39	0.32	0.03**	−0.07	0.03**	4.3
World	−0.04	0.23	0.5	−0.19	0.63	0.31	0.03**	−0.07	0.04**	3.6

Table 9: **Oil supply, global demand, and oil-specific demand shocks as predictors of MSCI index excess returns: economic significance over the 1986.01–2015.12 sample period.** This table presents evidence on the performance of the trading strategies using oil supply, global demand, and oil-specific demand shocks as predictors in terms of two metrics: the certainty equivalent return (CER) and the Sharpe ratio (SR). The investor forms optimal mean-variance portfolios between an MSCI index and the corresponding Treasure bill and has a risk aversion coefficient $\gamma = 3$. The weight on the MSCI index is constrained between $\underline{\omega} = 0$ and $\overline{\omega} = 150\%$. We compare the augmented model (DEC) described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13) to three baseline models. The first baseline model (IID) assumes that the MSCI index excess return $r^{\mathtt{e}}_{t+1}$ is i.i.d. The second baseline model (P) is described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The third baseline model (D) is described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{D}} + \delta^{\mathtt{D}} g^{\mathtt{D}}_t + u^{\mathtt{D}}_{t+1}$, where the global demand growth proxy $g^{\mathtt{D}}_t$ is the first principal component obtained from the log growth rates of shipping cost index and global crude steel production. The portfolio strategies use the mean and variance estimates resulting from each model respectively. Reported are the CERs in annualized percentage points and the annualized SRs for the augmented and the two baseline models. We also report the one-sided p-values for the null hypothesis that the augmented model does not improve the CER and SR obtained from a strategy based on each of the two baseline models. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. The results are based on an expanding window with 60 initial observations.

	CER							SR						
	DEC	IID	p-value	P	p-value	D	p-value	DEC	IID	p-value	P	p-value	D	p-value
Local-currency														
Canada	6.68	4.40	0.09*	4.90	0.17	5.62	0.26	0.45	0.25	0.10	0.31	0.22	0.37	0.27
France	6.03	4.24	0.24	6.10	0.51	2.93	0.11	0.41	0.23	0.14	0.41	0.49	0.18	0.08*
Germany	7.44	3.58	0.08*	5.77	0.23	3.78	0.09*	0.51	0.17	0.03**	0.40	0.18	0.24	0.06*
Italy	10.70	4.32	0.03**	8.73	0.12	3.04	0.01***	0.61	−0.06	0.00***	0.50	0.13	0.00	0.00***
Japan	3.56	−0.20	0.05**	−1.67	0.07*	−1.37	0.01**	0.42	−0.11	0.00***	0.02	0.02**	−0.14	0.00***
UK	8.52	4.74	0.04**	7.32	0.25	3.94	0.03**	0.53	0.19	0.02**	0.45	0.28	0.13	0.02**
US-dollar														
Canada	4.06	3.31	0.37	1.96	0.25	4.12	0.51	0.34	0.26	0.30	0.19	0.23	0.34	0.48
France	0.90	3.46	0.81	2.48	0.76	1.82	0.63	0.22	0.26	0.64	0.25	0.63	0.20	0.47
Germany	3.50	3.27	0.47	2.37	0.37	3.29	0.47	0.33	0.21	0.20	0.25	0.27	0.26	0.33
Italy	5.30	2.27	0.19	3.20	0.27	0.48	0.10*	0.44	0.05	0.01***	0.34	0.23	0.05	0.02**
Japan	4.11	1.46	0.04**	0.60	0.02**	−0.69	0.01***	0.30	−0.11	0.00***	−0.05	0.00***	−0.26	0.00***
UK	5.33	3.38	0.19	2.74	0.21	2.67	0.12	0.41	0.25	0.12	0.23	0.18	0.22	0.08*
USA	9.28	6.52	0.11	6.08	0.08*	6.20	0.05**	0.65	0.48	0.12	0.45	0.09*	0.46	0.05**
World	7.90	4.03	0.04**	3.88	0.07*	3.30	0.03**	0.56	0.31	0.04**	0.30	0.06*	0.27	0.03**

Table 10: **Bai-Perron structural break tests: oil supply, global demand, and oil-specific demand shocks as predictors of MSCI index returns for the 1986.01–2015.12 sample period.** This table presents Bai and Perron (2003) structural break tests for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, oil-specific demand shocks obtained in the oil price change decomposition (13), respectively. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. Bold numbers indicate the lowest value assumed by the Bayesian Information Criterion (BIC).

	No break model	One break model		# of breaks
	BIC	BIC	Break date	selected by BIC
Local-currency				
Canada	**−6.298**	−6.257	1994.02	0
France	**−5.794**	−5.774	1995.03	0
Germany	**−5.598**	−5.573	1994.06	0
Italy	**−5.537**	−5.534	1994.06	0
Japan	**−5.737**	−5.703	1990.09	0
UK	**−6.240**	−6.197	2008.08	0
US-dollar				
Canada	**−5.803**	−5.765	2008.07	0
France	**−5.597**	−5.561	1999.02	0
Germany	**−5.436**	−5.407	1999.03	0
Italy	**−5.313**	−5.296	1990.09	0
Japan	**−5.551**	−5.531	1990.09	0
UK	**−5.976**	−5.944	2010.01	0
USA	**−6.299**	−6.270	2008.08	0
World	**−6.276**	−6.249	2008.07	0

Table 11: **Oil supply, global demand, and oil-specific demand shocks as predictors of delayed monthly MSCI index returns: statistical significance over the 1986.01–2015.12 sample period.** This table presents results for the predictive regression $r^{\mathtt{d}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{d}}_{t+1}$ is the delayed monthly return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{D}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on Newey and West (1987) standard errors with optimal bandwidth selected as in Newey and West (1994). Results based on a one-week and two-week delay are presented in the left and right panel, respectively. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	One-week delay							Two-week delay						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{D}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{D}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local-currency														
Canada	−0.21	0.55	0.29	0.00***	−0.04	0.22	1.7	−0.14	0.68	0.30	0.00***	−0.03	0.35	1.7
France	0.01	0.98	0.27	0.06*	−0.12	0.00***	2.7	−0.16	0.69	0.26	0.01**	−0.09	0.02**	2.2
Germany	−0.18	0.76	0.35	0.05*	−0.13	0.00***	3.2	−0.50	0.32	0.37	0.00***	−0.11	0.00***	3.7
Italy	−0.36	0.44	0.47	0.00***	−0.21	0.00***	7.7	−0.30	0.56	0.50	0.00***	−0.19	0.00***	7.0
Japan	−0.09	0.81	0.37	0.00***	−0.08	0.02**	2.5	−0.13	0.76	0.40	0.00***	−0.08	0.02**	2.6
UK	0.39	0.50	0.24	0.02**	−0.09	0.00***	2.4	0.27	0.56	0.24	0.02**	−0.07	0.02**	1.8
US-dollar														
Canada	−0.29	0.52	0.40	0.01***	−0.02	0.68	1.8	−0.28	0.51	0.41	0.00***	−0.01	0.83	2.3
France	0.03	0.95	0.32	0.02**	−0.08	0.01***	1.3	0.05	0.91	0.35	0.01***	−0.06	0.09*	1.6
Germany	−0.16	0.75	0.40	0.01***	−0.10	0.00***	2.2	−0.28	0.55	0.46	0.00***	−0.09	0.01***	3.3
Italy	−0.35	0.47	0.51	0.00***	−0.19	0.00***	5.5	−0.02	0.97	0.59	0.00***	−0.18	0.00***	5.6
Japan	0.19	0.71	0.27	0.02**	−0.08	0.02**	1.1	0.20	0.70	0.37	0.00***	−0.09	0.01***	2.0
UK	0.33	0.54	0.39	0.00***	−0.07	0.04**	2.4	0.40	0.36	0.41	0.00***	−0.06	0.04**	3.2
USA	−0.04	0.92	0.34	0.00***	−0.07	0.01***	3.2	0.09	0.79	0.34	0.00***	−0.03	0.23	2.6
World	−0.06	0.86	0.37	0.00***	−0.07	0.00***	3.8	0.03	0.91	0.38	0.00***	−0.06	0.02**	4.0

Table 12: **Oil supply, global demand, and oil-specific demand shocks as predictors of long-horizon MSCI index returns: statistical significance over the 1986.01–2015.12 sample period.** This table presents results for the predictive regression $r_{t,t+h} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{D}} x_t^{\mathtt{D}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_{t+h}^{\mathtt{DEC}}$, where $r_{t,t+h}$ is the h-month return on an MSCI index and $x_t^{\mathtt{S}}$, $x_t^{\mathtt{D}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{D}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on Newey and West (1987) standard errors with optimal bandwidth selected as in Newey and West (1994). Results based on three-month and six-month returns are presented in the left and right panel, respectively. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Three-month return							Six-month return						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{D}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{D}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local-currency														
Canada	−0.79	0.21	0.67	0.00***	0.01	0.80	3.4	−1.17	0.24	0.45	0.11	0.03	0.70	0.5
France	−1.18	0.12	0.76	0.05**	−0.11	0.11	4.0	−1.58	0.25	0.72	0.07*	−0.19	0.06*	2.3
Germany	−1.53	0.09*	0.92	0.01***	−0.09	0.19	4.6	−1.29	0.40	0.72	0.08*	−0.17	0.08*	1.4
Italy	−2.22	0.01***	1.04	0.00***	−0.29	0.00***	10.9	−1.97	0.13	0.92	0.02**	−0.36	0.00***	5.2
Japan	−0.97	0.26	0.63	0.01***	−0.11	0.15	2.3	−1.98	0.14	0.34	0.26	−0.25	0.06*	2.0
UK	−0.19	0.77	0.48	0.06*	−0.04	0.30	1.4	−0.09	0.93	0.59	0.06*	−0.05	0.53	0.7
US-dollar														
Canada	−0.80	0.31	0.89	0.01***	0.05	0.52	3.7	−1.29	0.28	0.56	0.22	0.06	0.52	0.4
France	−1.46	0.11	0.92	0.02**	−0.01	0.88	4.6	−2.10	0.14	0.72	0.10	−0.13	0.24	2.1
Germany	−1.84	0.05**	1.10	0.00***	0.00	0.99	5.7	−1.75	0.19	0.72	0.11	−0.13	0.15	1.3
Italy	−2.45	0.01**	1.21	0.00***	−0.20	0.04**	9.0	−2.77	0.03**	0.91	0.06*	−0.33	0.03**	4.9
Japan	−0.32	0.79	0.71	0.00***	−0.10	0.27	1.7	−1.71	0.31	0.66	0.03**	−0.26	0.19	2.1
UK	−0.25	0.75	0.95	0.00***	−0.01	0.89	5.1	−0.67	0.53	0.87	0.02**	−0.03	0.73	1.7
USA	−0.38	0.57	0.79	0.01***	−0.02	0.62	4.9	−1.02	0.27	0.72	0.02**	−0.03	0.71	1.9
World	−0.66	0.33	0.89	0.00***	−0.05	0.33	6.5	−1.25	0.22	0.78	0.02**	−0.12	0.17	3.0

Table 13: **Oil supply, global demand, and oil-specific demand shocks as predictors of long-horizon MSCI index returns: statistical significance over the 1986.01–2015.12 sample period.** This table presents adjusted R^2 for the predictive regression $r_{t,t+h} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{D}} x_t^{\mathtt{D}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_{t+h}^{\mathtt{DEC}}$, where $r_{t,t+h}$ is the h-month return on an MSCI index and $x_t^{\mathtt{S}}$, $x_t^{\mathtt{D}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). The adjusted R^2, stated in percentage, based on two-month to six-month returns are presented. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively.

	Two-month	Three-month	Four-month	Five-month	Six-month
Local-currency					
Canada	3.4	3.4	3.5	1.7	0.5
France	3.9	4.0	4.8	3.4	2.3
Germany	5.1	4.6	4.2	2.6	1.4
Italy	11.1	10.9	9.9	6.7	5.2
Japan	2.7	2.3	2.5	2.1	2.0
UK	2.6	1.4	3.0	2.0	0.7
US-dollar					
Canada	3.8	3.7	3.5	1.6	0.4
France	3.1	4.6	4.6	2.9	2.1
Germany	4.7	5.7	4.4	2.4	1.3
Italy	8.6	9.0	8.0	5.7	4.9
Japan	1.8	1.7	2.4	2.3	2.1
UK	5.3	5.1	5.5	3.6	1.7
USA	5.5	4.9	5.6	3.6	1.9
World	6.4	6.5	6.9	4.7	3.0

Table 14: **Oil supply, global demand, and oil-specific demand shocks as predictors of long-horizon MSCI index returns: statistical significance over the 1986.01–2015.12 sample period.** This table presents adjusted R^2 for the predictive regression $r_{t,t+h} = \alpha^{\texttt{DEC}} + \beta^{\texttt{S}} x_t^{\texttt{S}} + \beta^{\texttt{D}} x_t^{\texttt{D}} + \beta^{\texttt{OSD}} x_t^{\texttt{OSD}} + u_{t+h}^{\texttt{DEC}}$, where $r_{t,t+h}$ is the h-month return on an MSCI index and $x_t^{\texttt{S}}$, $x_t^{\texttt{D}}$, and $x_t^{\texttt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). The adjusted R^2, stated in percentage, based on one-month to six-month returns are presented. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively.

	One-month	Two-month	Three-month	Four-month	Five-month	Six-month
Oil price change decomposition						
Local-currency						
Canada	1.6	3.4	3.4	3.5	1.7	0.5
France	2.9	3.9	4.0	4.8	3.4	2.3
Germany	3.0	5.1	4.6	4.2	2.6	1.4
Italy	9.6	11.1	10.9	9.9	6.7	5.2
Japan	1.9	2.7	2.3	2.5	2.1	2.0
UK	2.2	2.6	1.4	3.0	2.0	0.7
US-dollar						
Canada	1.5	3.8	3.7	3.5	1.6	0.4
France	1.4	3.1	4.6	4.6	2.9	2.1
Germany	1.9	4.7	5.7	4.4	2.4	1.3
Italy	6.6	8.6	9.0	8.0	5.7	4.9
Japan	0.6	1.8	1.7	2.4	2.3	2.1
UK	2.1	5.3	5.1	5.5	3.6	1.7
USA	4.3	5.5	4.9	5.6	3.6	1.9
World	3.7	6.4	6.5	6.9	4.7	3.0
Oil price change						
	One-month	Two-month	Three-month	Four-month	Five-month	Six-month
Local-currency						
Canada	1.0	0.0	-0.5	-0.5	-0.5	-0.3
France	5.6	3.7	1.4	1.7	1.4	0.8
Germany	5.3	2.7	0.2	0.4	0.5	0.2
Italy	17.8	12.5	9.5	8.1	5.2	4.1
Japan	3.0	4.1	1.8	2.1	2.1	1.7
UK	4.4	1.6	0.1	0.6	-0.1	-0.4
US-dollar						
Canada	0.7	-0.1	-0.5	-0.5	-0.5	-0.3
France	3.9	2.3	-0.2	0.3	0.8	0.9
Germany	4.2	2.3	-0.5	-0.2	0.3	0.6
Italy	15.3	11.0	6.5	6.8	5.8	5.7
Japan	1.0	2.5	0.4	0.8	1.2	1.4
UK	2.9	1.4	-0.4	-0.4	-0.5	-0.5
USA	5.3	1.5	-0.1	0.3	-0.3	-0.4
World	5.8	4.1	0.8	1.5	1.2	1.0

Table 15: **Bai-Perron structural break tests for excess returns on Fama-French 17 value-weighted industry portfolios: oil price change as the predictor of industry portfolio excess returns for the 1983.01–2015.12 sample period, oil supply, global demand, and oil-specific demand shocks as predictors of industry portfolio excess returns for the 1986.01–2015.12 sample period.** The left panel presents Bai and Perron (2003) structural break tests for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents Bai and Perron (2003) structural break tests for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, oil-specific demand shocks obtained in the oil price change decomposition (13), respectively. The number of breaks is selected by the Bayesian Information Criterion (BIC).

	Oil price change		Oil price change decomposition
	# of breaks selected by BIC	Break date	# of breaks selected by BIC
Food	1	1991.12	0
Mines	0	–	0
Oil	0	–	0
Clths	1	2008.08	0
Durbl	1	2008.06	0
Chems	1	2008.07	0
Cnsum	1	2000.05	0
Cnstr	1	2008.08	0
Steel	1	2008.07	0
FabPr	1	2008.08	0
Machn	1	2008.08	0
Cars	1	2008.08	0
Trans	1	2008.06	0
Utils	0	–	0
Rtail	1	2008.08	0
Finan	1	2008.06	0
Other	1	2008.08	0

Table 16: **Oil supply, global demand, and oil-specific demand shocks as predictors of excess returns on Fama-French 17 value-weighted industry portfolios: statistical significance over the 1986.01–2015.12 sample period.** The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{D}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on Newey and West (1987) standard errors with optimal bandwidth selected as in Newey and West (1994). *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$.

	Oil price change			Oil price change decomposition						
	$\delta^{\mathtt{P}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{D}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Food	−0.028	0.45	0.1	−0.36	0.39	0.15	0.11	−0.04	0.23	0.9
Mines	0.018	0.74	−0.2	−0.69	0.28	0.34	0.06*	0.01	0.92	0.6
Oil	0.004	0.90	−0.3	−0.04	0.93	0.12	0.30	−0.00	0.91	−0.6
Clths	−0.077	0.22	1.0	−1.13	0.05*	0.52	0.00***	−0.12	0.01**	7.5
Durbl	−0.025	0.66	−0.1	−0.39	0.44	0.53	0.00***	−0.07	0.11	4.9
Chems	0.018	0.76	−0.2	−0.59	0.22	0.56	0.00***	−0.02	0.72	4.7
Cnsum	−0.042	0.29	0.5	−0.37	0.37	0.13	0.15	−0.05	0.16	1.0
Cnstr	−0.085	0.09*	1.5	−0.78	0.15	0.34	0.02**	−0.11	0.00***	4.7
Steel	−0.001	0.98	−0.3	−0.56	0.38	0.59	0.01**	−0.04	0.51	2.5
FabPr	−0.020	0.65	−0.2	−0.39	0.41	0.30	0.01***	−0.04	0.34	1.5
Machn	−0.064	0.32	0.4	−0.24	0.70	0.53	0.00***	−0.12	0.02**	3.9
Cars	−0.058	0.32	0.4	−0.85	0.09*	0.43	0.01**	−0.09	0.04**	3.6
Trans	−0.028	0.50	−0.0	−0.31	0.52	0.43	0.00***	−0.07	0.08*	4.1
Utils	0.005	0.87	−0.3	−0.55	0.08*	0.07	0.33	0.01	0.76	0.4
Rtail	−0.094	0.01***	2.6	−0.40	0.40	0.26	0.03**	−0.12	0.00***	5.1
Finan	−0.041	0.51	0.2	−0.80	0.09*	0.54	0.00***	−0.08	0.04**	7.0
Other	−0.055	0.22	0.7	−0.24	0.59	0.37	0.01***	−0.09	0.01***	4.2

Table 17: **Correlations between oil supply shocks, global demand shocks, oil-specific demand shocks, and several US macroeconomic variables.** This table presents the correlation matrix for oil supply shock ($x_t^{\mathtt{S}}$), global demand shock ($x_t^{\mathtt{D}}$), oil-specific demand shock ($x_t^{\mathtt{OSD}}$), the log dividend yield (dy), the term spread (tms), the default yield spread (dfy), and the one-month T-bill rate (tbl). Results are presented for the 1986.01–2015.12 sample period. The top (bottom) panel shows results of the three shocks obtained by the oil price change real-time (full-sample) decomposition.

Real-time decomposition

	$x^{\mathtt{S}}$	$x^{\mathtt{D}}$	$x^{\mathtt{OSD}}$	dy	tms	dfy	tbl
$x^{\mathtt{S}}$	1	-0.09	0.17	-0.04	0.01	-0.01	-0.02
$x^{\mathtt{D}}$	-0.09	1	0.08	-0.03	0.05	-0.17	0.04
$x^{\mathtt{OSD}}$	0.17	0.08	1	-0.08	0.00	-0.08	0.04
dy	-0.04	-0.03	-0.08	1	0.20	0.32	0.31
tms	0.01	0.05	0.00	0.20	1	0.27	-0.62
dfy	-0.01	-0.17	-0.08	0.32	0.27	1	-0.27
tbl	-0.02	0.04	0.04	0.31	-0.62	-0.27	1

Full-sample decomposition

	$x^{\mathtt{S}}$	$x^{\mathtt{D}}$	$x^{\mathtt{OSD}}$	dy	tms	dfy	tbl
$x^{\mathtt{S}}$	1	-0.08	0.02	0.01	-0.01	0.04	0.02
$x^{\mathtt{D}}$	-0.08	1	0.04	-0.08	0.06	-0.16	0.00
$x^{\mathtt{OSD}}$	0.02	0.04	1	-0.08	0.00	-0.07	0.03
dy	0.01	-0.08	-0.08	1	0.20	0.32	0.31
tms	-0.01	0.06	0.00	0.20	1	0.27	-0.62
dfy	0.04	-0.16	-0.07	0.32	0.27	1	-0.27
tbl	0.02	0.00	0.03	0.31	-0.62	-0.27	1

Table 18: **Robustness checks of the USA MSCI index excess return predictive regressions: the role of macroeconomic variables.** This table presents results for the predictive regressions $r^{\mathtt{e}}_{t+1} = \gamma^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + \boldsymbol{\theta}'\mathbf{z}_t + v^{\mathtt{P}}_{t+1}$ and $r^{\mathtt{e}}_{t+1} = \gamma^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + \boldsymbol{\lambda}'\mathbf{z}_t + v^{\mathtt{DEC}}_{t+1}$ over various sample periods. The variables in these regressions are: (i) $r^{\mathtt{e}}_{t+1}$ is excess return on the USA MSCI index, (ii) the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light, (iii) $x^{\mathtt{S}}_t$, $x^{\mathtt{D}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (13), and (iv) $\mathbf{z}_t$ is the vector of macroeconomic variables including the log dividend yield (dy), the term spread (tms), the default yield spread (dfy), and the one-month T-bill rate (tbl). The numbers in parentheses represent two-sided p-values for the null hypotheses that the slope coefficients are zero, based on Newey and West (1987) standard errors with optimal bandwidth selected as in Newey and West (1994). *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented in percentage, is denoted by $\bar{R}^2$.

Sample Period	1983.01–2003.04	1983.01–2015.12	1986.01–2015.12	1986.01–2015.12
$g^{\mathtt{P}}_t$	−0.11***	−0.04	−0.04	
	(0.00)	(0.28)	(0.25)	
$x^{\mathtt{S}}_t$				−0.34
				(0.38)
$x^{\mathtt{D}}_t$				0.34***
				(0.00)
$x^{\mathtt{OSD}}_t$				−0.06**
				(0.03)
dy_t	0.03**	0.03***	0.03***	0.03***
	(0.04)	(0.00)	(0.00)	(0.00)
tms_t	−0.78	−0.43*	−0.49*	−0.57**
	(0.14)	(0.08)	(0.08)	(0.04)
dfy_t	−0.32	−1.16	−1.36	−0.99
	(0.74)	(0.25)	(0.15)	(0.17)
tbl_t	−0.45	−0.29**	−0.33**	−0.35**
	(0.14)	(0.02)	(0.05)	(0.03)
$\bar{R}^2$	6.3	2.0	2.0	5.7

Table 19: **Descriptive statistics for MSCI index predicted excess returns: 1986.01–2015.12.** This table presents the mean, the standard deviation (STD), and the first three autocorrelations (AC(k), $k = 1, 2, 3$) of MSCI index predicted excess returns using two sets of predictors: (i) four macroeconomic variables, i.e., the log dividend yield, the term spread, the default yield spread, and the one-month T–bill rate and (ii) the oil supply, global demand, and oil-specific demand shocks obtained from the oil price decomposition. The top panel contains the results for the MSCI USA index based on the macroeconomic variables. The middle (bottom) panel contains the results for MSCI local–currency (US–dollar) denominated indexes based on the oil price decomposition.

	Mean	STD	AC(1)	AC(2)	AC(3)
Macroeconomic variables					
USA	0.65	0.74	0.97	0.93	0.89
Oil price decomposition					
Local-currency					
Canada	0.40	0.67	0.56	0.14	-0.09
France	0.50	1.08	0.32	-0.02	-0.11
Germany	0.48	1.23	0.36	-0.01	-0.12
Italy	0.20	2.12	0.28	-0.03	-0.10
Japan	0.30	0.95	0.39	0.06	-0.09
UK	0.35	0.79	0.27	0.00	-0.09
US-dollar					
Canada	0.57	0.85	0.61	0.19	-0.07
France	0.66	0.93	0.38	-0.03	-0.13
Germany	0.59	1.10	0.41	-0.02	-0.13
Italy	0.43	1.99	0.30	-0.04	-0.10
Japan	0.27	0.77	0.31	-0.02	-0.08
UK	0.56	0.87	0.46	0.09	-0.10
USA	0.65	0.99	0.44	0.04	-0.12
World	0.53	0.94	0.43	0.06	-0.12

Table 20: **Descriptive statistics for MSCI index predicted excess returns: 1996.01–2012.01.** This table presents the mean, the standard deviation (STD), and the first three autocorrelations (AC(k), $k = 1, 2, 3$) of MSCI index predicted excess returns using two sets of predictors: (i) four macroeconomic variables, i.e., the log dividend yield, the term spread, the default yield spread, and the one-month T–bill rate and (ii) the oil supply, global demand, and oil-specific demand shocks obtained from the oil price decomposition. For comparison purposes, the top panel reports the descriptive statistics for the equity premium estimates of Martin (2017), obtained from option prices, for one-month (M1), two-month (M2), three-month (M3), six-month (M6), and one-year (M12) maturities. The second panel contains the results for the MSCI USA index based on the macroeconomic variables. The third (fourth) panel contains the results for MSCI local–currency (US–dollar) denominated indexes based on the oil price decomposition.

	Mean	STD	AC(1)	AC(2)	AC(3)
Martin (2017)'s equity premium					
M1	0.41	1.15	0.78	0.55	0.45
M2	0.41	1.05	0.82	0.61	0.51
M3	0.41	0.97	0.85	0.67	0.57
M6	0.41	0.82	0.88	0.73	0.64
M12	0.39	0.69	0.90	0.80	0.72
Macroeconomic variables					
USA	0.42	1.01	0.95	0.87	0.81
Oil price decomposition					
Local-currency					
Canada	0.62	1.22	0.66	0.19	-0.08
France	0.48	1.50	0.52	0.00	-0.17
Germany	0.53	1.95	0.48	-0.06	-0.19
Italy	0.25	1.77	0.43	-0.04	-0.16
Japan	-0.16	1.04	0.66	0.22	-0.03
UK	0.23	0.91	0.40	-0.02	-0.12
US-dollar					
Canada	0.88	1.65	0.68	0.23	-0.06
France	0.49	1.49	0.55	0.05	-0.15
Germany	0.54	1.90	0.50	-0.02	-0.18
Italy	0.39	1.77	0.47	0.01	-0.14
Japan	-0.21	0.86	0.54	0.04	-0.03
UK	0.37	1.36	0.63	0.17	-0.08
USA	0.42	1.54	0.55	0.06	-0.14
World	0.32	1.39	0.58	0.11	-0.11

Table 21: **Effect of oil supply, global demand, and oil-specific demand shocks on conditional return volatility: evidence from an EGARCH(1,1) model.** This table presents results of an augmented EGARCH(1,1) model that includes the three shocks in the volatility equation as exogenous regressors. The econometric specification is: $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{D}} x^{\mathtt{D}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, with $u^{\mathtt{DEC}}_{t+1} = \sigma_t z_{t+1}$, $z_{t+1} \sim$ i.i.d.$(0,1)$, and $\log(\sigma^2_t) = \tau_0 + \tau_1|z_t| + \tau_2 z_t + \tau_3 \log(\sigma^2_{t-1}) + \zeta^{\mathtt{S}} x^{\mathtt{S}}_t + \zeta^{\mathtt{D}} x^{\mathtt{D}}_t + \zeta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t$. The reported results are based on the Student-t distribution for the disturbances z_{t+1}, selected among the Normal, Student-t, and GED distributions according to the Bayesian Information Criterion. The model is estimated using monthly excess returns on the MSCI indexes for the G7 countries as well as the World MSCI index over the 1986.01–2015.12 sample period. The top and bottom panels contain results for local-currency and US-dollar denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	τ_1	p-value	τ_2	p-value	τ_3	p-value	$\zeta^{\mathtt{S}}$	p-value	$\zeta^{\mathtt{D}}$	p-value	$\zeta^{\mathtt{OSD}}$	p-value
Local-currency												
Canada	0.16	0.03**	0.07	0.11	0.97	0.00***	−2.98	0.59	−3.16	0.04**	−0.03	0.96
France	0.33	0.00***	−0.20	0.00***	0.81	0.00***	3.05	0.75	0.07	0.98	−0.04	0.95
Germany	0.28	0.01**	−0.09	0.15	0.87	0.00***	6.73	0.45	−2.21	0.36	0.18	0.81
Italy	0.13	0.11	−0.07	0.04**	0.96	0.00***	7.24	0.24	−2.46	0.15	−0.02	0.97
Japan	0.00	0.97	−0.04	0.43	−0.79	0.00***	6.13	0.70	−8.56	0.05*	−0.02	0.98
UK	0.28	0.03**	−0.17	0.05**	0.78	0.00***	0.19	0.98	1.02	0.69	0.14	0.86
US-dollar												
Canada	0.13	0.04**	0.08	0.09*	0.96	0.00***	−4.59	0.38	−3.49	0.02**	−0.40	0.49
France	0.14	0.11	−0.21	0.00***	0.89	0.00***	5.37	0.45	−0.60	0.71	−0.27	0.58
Germany	0.21	0.04**	−0.08	0.17	0.92	0.00***	5.43	0.46	−1.85	0.32	−0.20	0.73
Italy	−0.07	0.03**	−0.02	0.54	0.98	0.00***	11.79	0.00***	−4.52	0.00***	0.26	0.52
Japan	0.00	0.95	−0.07	0.00***	0.99	0.00***	−4.64	0.30	−2.43	0.04**	0.29	0.33
UK	0.14	0.12	−0.11	0.06*	0.94	0.00***	−1.03	0.88	−0.20	0.91	−0.23	0.70
USA	0.25	0.03**	−0.18	0.03**	0.86	0.00***	−3.60	0.66	−0.37	0.86	−0.29	0.66
World	0.16	0.12	−0.18	0.02**	0.85	0.00***	4.40	0.61	−0.63	0.75	−0.22	0.73

Figure 1: **Scatter plots of US-dollar denominated MSCI World index return versus the one-month lagged log growth rate of West Texas Intermediate (WTI) spot price.** In this figure, we present the scatter plots of the US-dollar denominated MSCI World index return versus the one-month lagged log growth rate of WTI spot price over the 1982.01–2003.12 and 2004.01–2015.12 sample periods. The solid lines represent the fitted least-squares regression lines. The correlation between the MSCI World index return and the one-month lagged log growth rate of WTI spot price is -0.22, 0.26, and -0.04 over the 1982.01–2003.12, 2004.01–2015.12, and 1982.01–2015.12 sample periods, respectively.

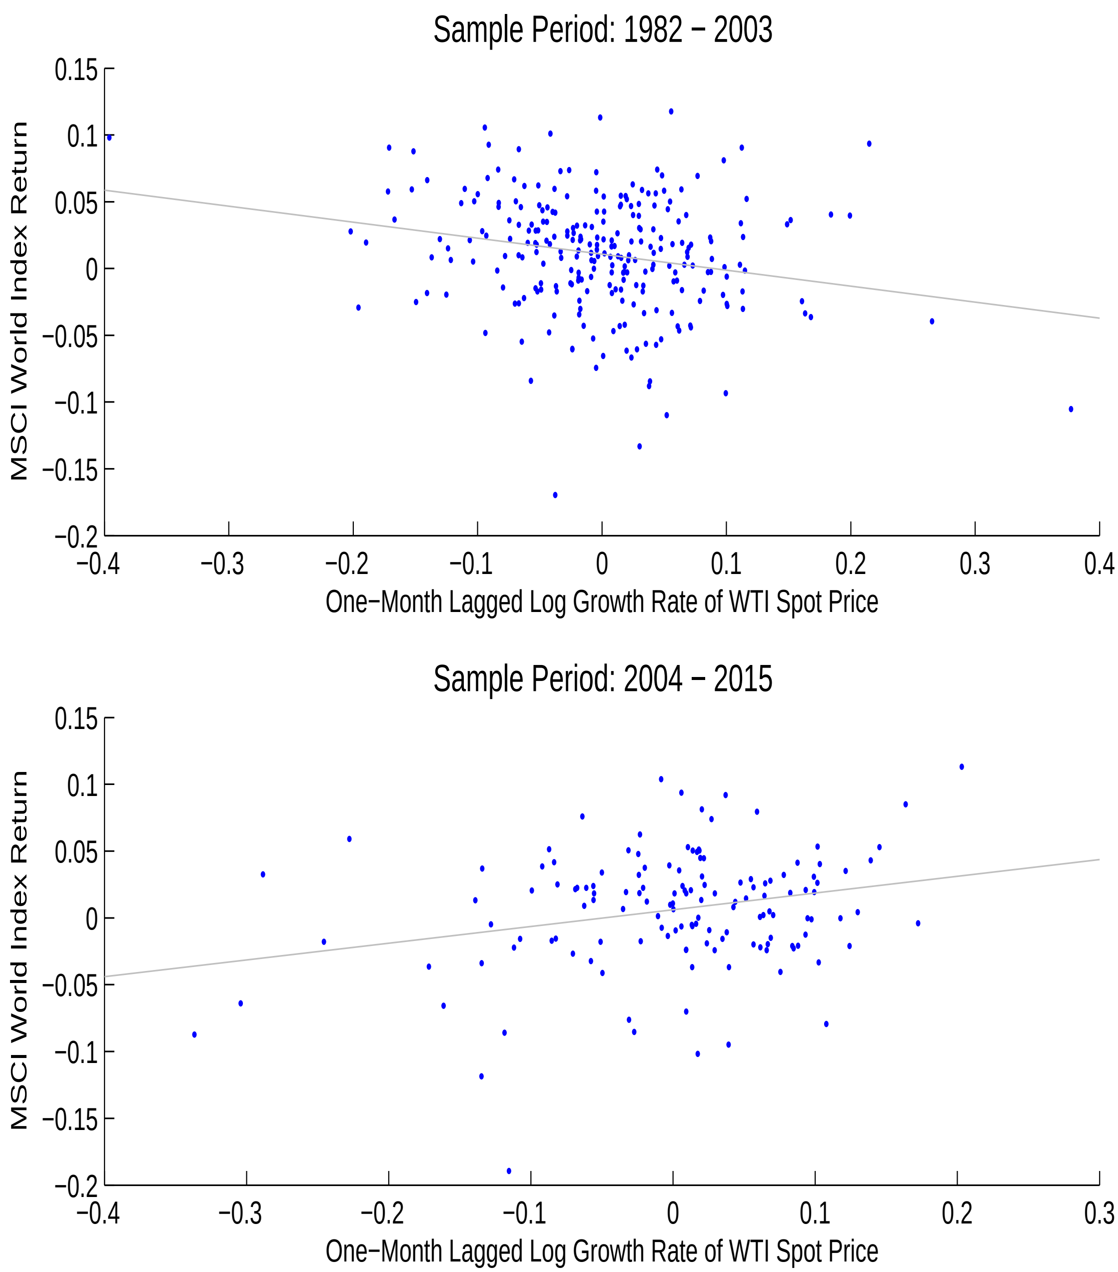

Figure 2: **Time series of oil price change proxies.** In this figure, we plot the log growth rates of three oil spot price proxies, i.e., WTI, Dubai, and Arab Light, along with their first principal component, g^{P}, over the 1983.01–2015.12 sample period. All series are rescaled so that they have a standard deviation equal to 0.09.

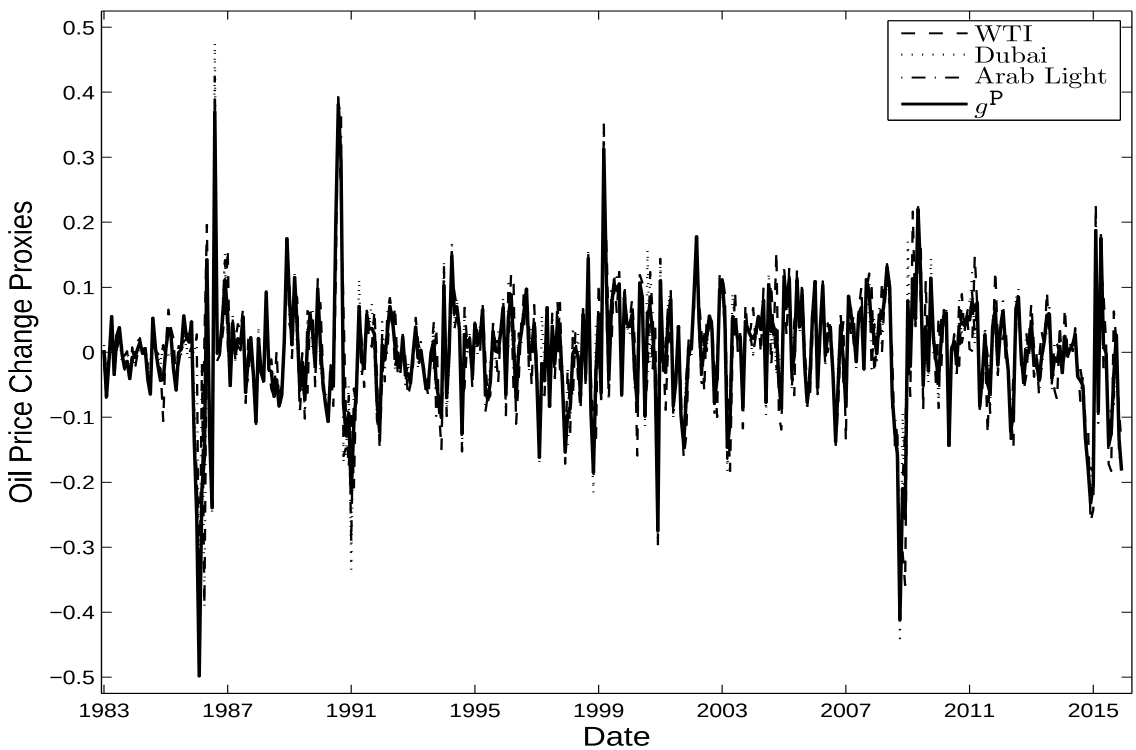

Figure 3: **Time series of global demand growth proxies.** In this figure, we plot the log growth rates of the shipping cost index and the seasonally-adjusted crude steel production, along with their first principal component, $g^{\mathtt{D}}$, over the 1983.01–2015.12 sample period. All series are rescaled so that they have a standard deviation equal to one.

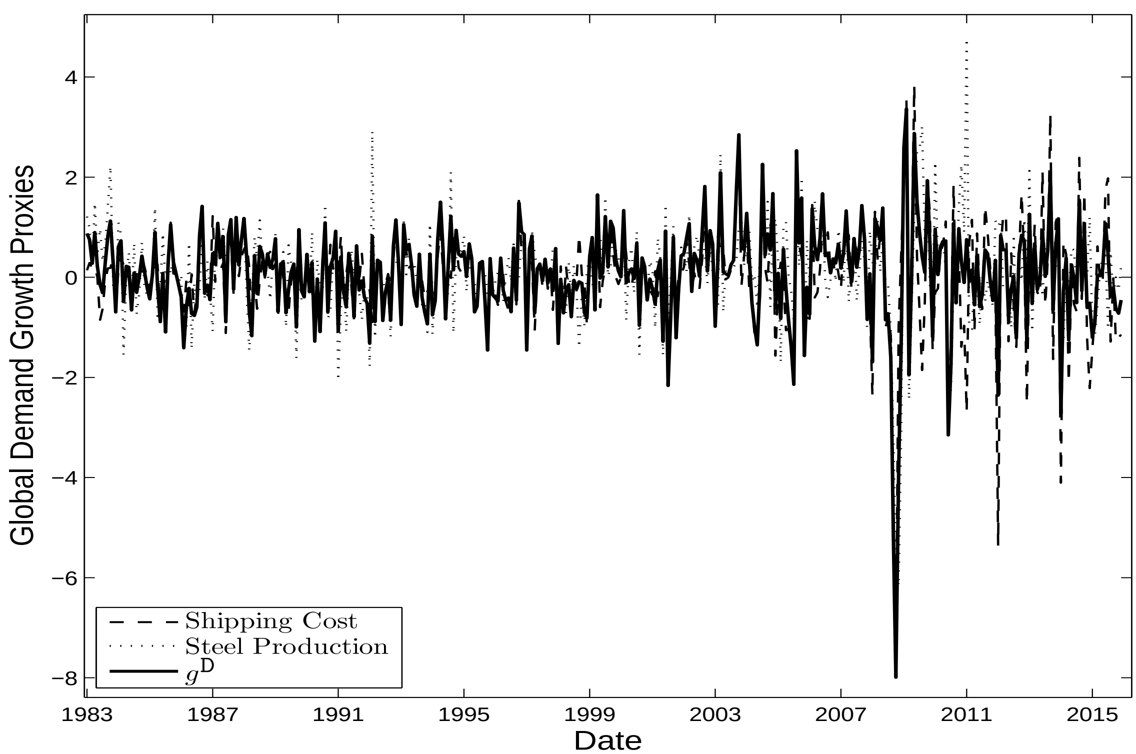

Figure 4: **Oil price change as a predictor of local-currency denominated MSCI index excess returns: slope estimates over expanding sample periods.** In this figure, we plot the time series of the slope estimates, along with the corresponding 95% confidence intervals based on Newey and West (1987) standard errors, from the predictive regression model (7) over different samples using an expanding window with the first sample being 1983.01–1993.01 and the last sample being 1983.01–2015.12.

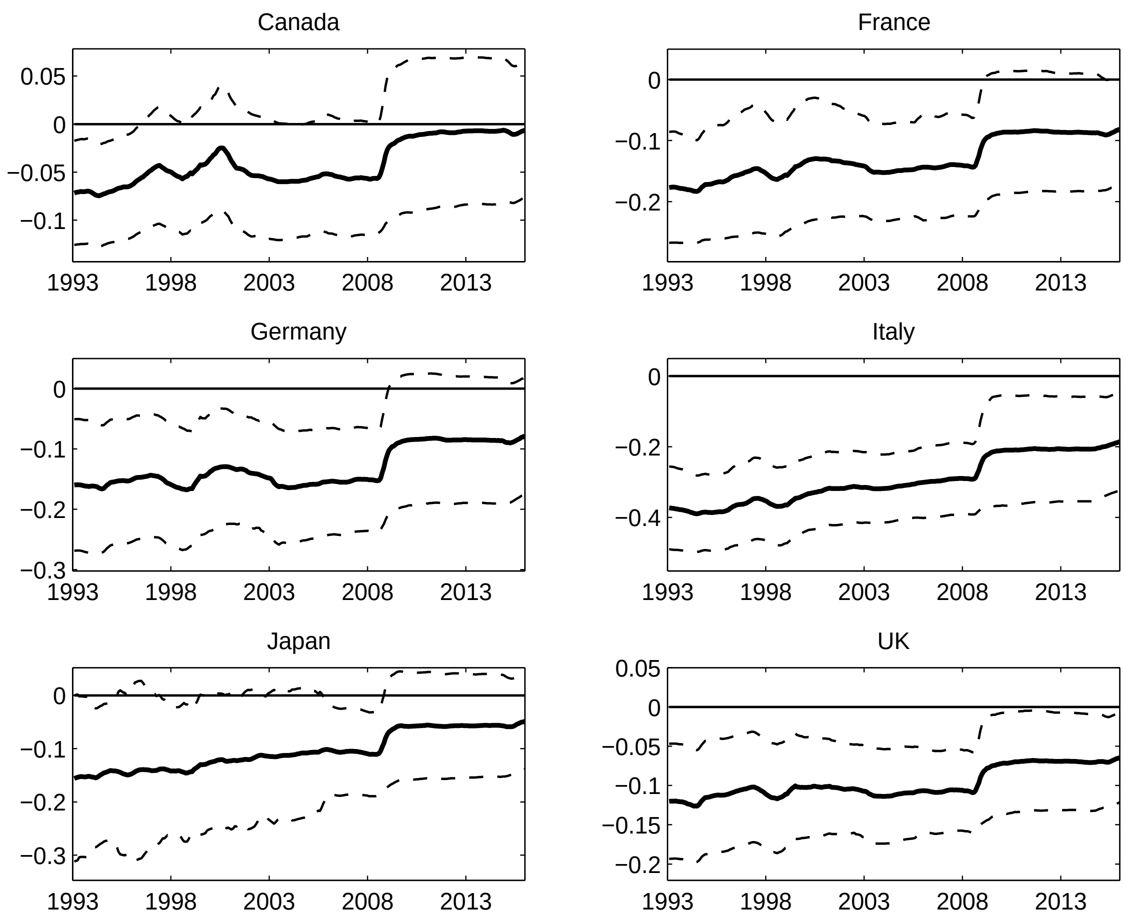

Figure 5: **Oil price change as a predictor of US-dollar denominated MSCI index excess returns: slope estimates over expanding sample periods.** In this figure, we plot the time series of the slope estimates, along with the corresponding 95% confidence intervals based on Newey and West (1987) standard errors, from the predictive regression model (7) over different samples using an expanding window with the first sample being 1983.01–1993.01 and the last sample being 1983.01–2015.12.

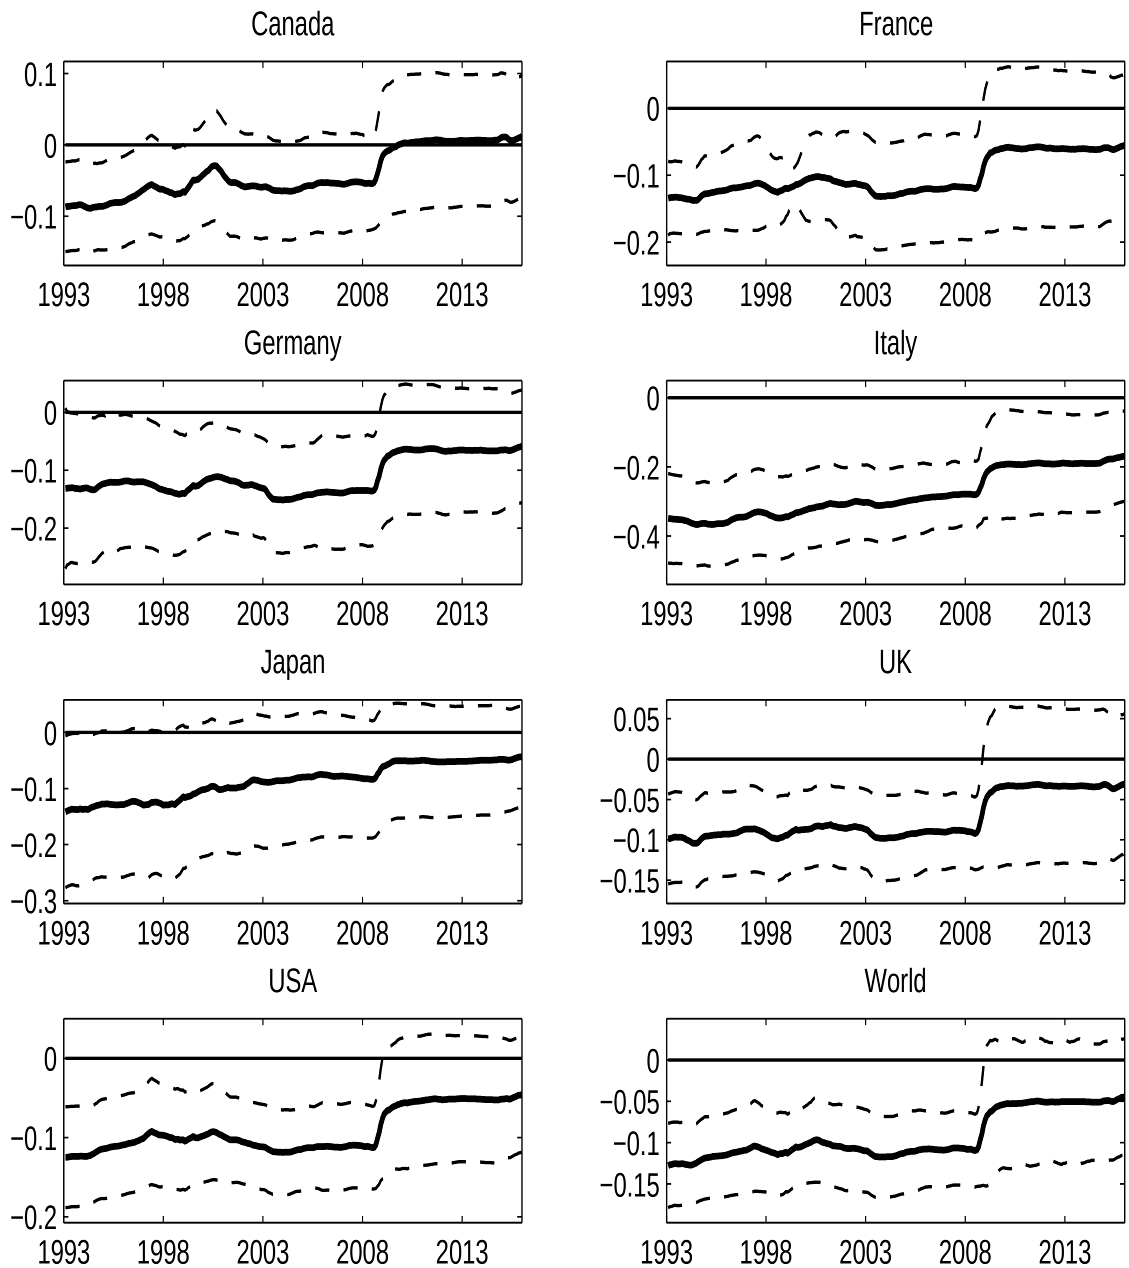

Figure 6: **Time series of oil supply, global demand, and oil-specific demand shocks.** In this figure, we plot the time series of the oil supply, global demand, and oil-specific demand shocks obtained using the decomposition in equation (13) over the 1986.01–2015.12 sample period. The shocks are obtained in a real-time fashion as described in subsection 5.1. All three series are rescaled so that they have standard deviation equal to one.