BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel — 2020-8

Last updated: March 23, 2020

The Effect of Oil Price Shocks on Asset Markets: Evidence from Oil Inventory News

by Ron Alquist,[1] Reinhard Ellwanger[2] and Jianjian Jin[3]

[1] AQR Capital Management
Greenwich, Connecticut, United States
ronalquist@gmail.com

[2] International Economic Analysis Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9
rellwanger@bankofcanada.ca

[3] British Columbia Investment Management Corporation
Victoria, British Columbia, Canada
jianjin78@gmail.com

ISSN 1701-9397

Acknowledgements

The views expressed in this paper are those of the authors; no responsibility for them should be attributed to AQR Capital Management, the Bank of Canada or the British Columbia Investment Management Corporation. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein.

Abstract

We quantify the reaction of U.S. equity, bond futures, and exchange rate returns to oil price shocks driven by oil inventory news. Across most sectors, equity prices decrease in response to higher oil prices before the 2007/08 crisis but increase after it. Positive oil price shocks cause a depreciation of the U.S. dollar against a broad range of currencies but have only a modest effect on bond futures returns. The evidence suggests that changes in risk premia help to explain the time-varying effect of oil price shocks on U.S. equity returns.

Bank topics: Financial markets; Recent economic and financial developments
JEL codes: D83, E44, G14, G15, Q41, Q43

Résumé

Nous quantifions la réaction des cours des actions, des contrats à terme sur obligations et des rendements des taux de change sur les marchés américains à des chocs pétroliers provoqués par des annonces sur les stocks d'hydrocarbures. Dans la plupart des secteurs, les cours des actions baissent du fait d'une hausse des prix du pétrole avant la crise de 2007-2008, alors qu'ils augmentent après la crise. Les chocs positifs des prix du pétrole entraînent une dépréciation du dollar américain face à de nombreuses monnaies mais n'ont que peu d'incidence sur les rendements des contrats à terme. Les données semblent indiquer que les variations des primes de risque peuvent aider à comprendre l'effet variable dans le temps des chocs pétroliers sur les rendements des actions américaines.

Sujets : Marchés financiers; Évolution économique et financière récente
Codes JEL : D83, E44, G14, G15, Q41, Q43

Non-technical summary

Oil price fluctuations have been linked to variation in the prices of financial assets such as equities, foreign exchange rates, and interest rates. However, because oil prices and asset prices move for a variety of reasons (for example, oil prices and asset prices mutually influence each other and respond jointly to macroeconomic developments), identifying the effects of oil price fluctuations on asset prices remains a significant challenge. We address this challenge by using the information contained in weekly U.S. oil inventory news, which allows us to identify variation in oil prices induced by oil-market-specific information.

To construct weekly oil inventory news, we subtract the actual change in U.S. inventories of crude oil, gasoline, and distillate inventories, as reported by the U.S. Energy Information Administration's Weekly Petroleum Status Report, from the expected change in inventories as reported by Bloomberg. Consistent with existing evidence, we find that higher-than-expected (lower-than-expected) U.S. oil inventories lead to systematic decreases (increases) in oil prices in the minutes following the announcement. The empirical approach is based on instrumental variables estimations in which we use the three types of oil inventory news as instruments for oil futures returns. We then analyze the effect of oil price changes on asset returns during a narrow window of 15 minutes around the announcement.

Our results support existing evidence for a structural break in the relationship between oil and asset returns around September 2008. We document that before the 2007/08 crisis, higher oil prices are associated with lower equity market returns, while after the crisis, higher oil prices are associated with higher equity market returns. This pattern holds for aggregate equity market returns and is pervasive across different sectors, including those with limited direct exposure to energy prices. The estimates for bond returns follow the reverse pattern. Bond futures returns tend to increase with higher oil prices before the crisis and to decrease with higher oil prices after the crisis. While these results suggest that nominal interest rates have become increasingly aligned with oil price fluctuations, the estimates are economically small and indicate that the effects of oil price changes on nominal interest rates are limited. Finally, we show that higher oil prices are associated with a depreciation of the U.S. dollar against a broad range of currencies, including those of major oil importers (such as the euro area).

The response of interest rates to oil prices, in combination with the time-varying effect of oil prices on equity returns, suggests that oil prices have become increasingly related to equity risk premia in the post-crisis period. We also investigate whether the informational content of U.S. oil inventories about global oil supply or demand conditions changed over time, but we find little evidence supporting this claim. More generally, our results show that oil price changes associated with inventory news have, on average, a more negative effect on U.S. stock returns than general news arrivals, highlighting a particular transmission channel for oil-market-specific news.

1 Introduction

Oil price fluctuations have important implications for the terms of trade, investment, output, and other macroeconomic aggregates of both oil-importing and oil-exporting economies. Yet, even before oil price shocks are fully transmitted to the real economy, the prices of financial assets adjust to reflect market expectations about the response of macroeconomic fundamentals to such shocks. Recent empirical research has related oil price fluctuations to variation in equity market returns (Kilian and Park 2009; Gao, Hitzemann, Shaliastovich, and Xu 2017; Ready 2017), exchange rates (Akram 2004; Chen, Rogoff, and Rossi 2010; Ferraro, Rogoff, and Rossi 2015), and interest rates (Datta, Johannsen, Kwon, and Vigfusson 2018; Kilian and Zhou 2019). However, because oil prices and asset prices are jointly determined with other macroeconomic variables in general equilibrium (Frankel 2008; Akram 2009; Kilian 2009; Chen, Rogoff, and Rossi 2010; Hitzemann 2016), identifying the effects of oil price fluctuations on asset prices remains a significant challenge.

This paper uses weekly news about U.S. oil inventories to investigate and quantify the effect of oil price shocks on the returns of different financial assets and the expectations about macroeconomic outcomes that the changes in returns reflect. Changes in oil inventories are a fundamental feature of oil markets and play a central role in the intertemporal relationship between current and future supply and demand conditions (Alquist and Kilian 2010; Kilian and Lee 2014; Kilian and Murphy 2014; Alquist, Bauer, and Rios 2014). As such, news of higher-than-expected (lower-than-expected) U.S. oil inventories leads to systematic decreases (increases) in oil prices in the minutes following the announcement (see, e.g., Halova, Kurov, and Kucher 2014). As shown in figure (1), the response of oil prices to news about U.S. crude oil inventories is strong and systematic, in a way that is consistent with economic theory. We use this variation in oil prices and a comprehensive, high-frequency dataset of financial variables, including stocks, bonds, and exchange rates, to study how information about oil market fundamentals is transmitted to asset prices and the broader economy. This approach differs from existing studies, which often rely on low-frequency exclusion restrictions in the context of structural models (Kilian and Park 2009; Ready 2017; Kilian and Zhou 2019) or on indirect evidence (Gao, Hitzemann, Shaliastovich, and Xu 2017; Datta, Johannsen, Kwon, and Vigfusson 2018) to establish causality between oil prices and financial variables.

Our empirical results support existing evidence for a structural break in the relationship between oil and asset returns around September 2008 (Lombardi and Ravazzolo 2016; Aït-Sahalia and Xiu 2016; Foroni, Guérin, and Marcellino 2017; Datta, Johannsen, Kwon, and Vigfusson 2018). We document that before the 2007/08 crisis, higher oil prices are associated with lower equity market returns, while post-crisis, higher oil prices are associated with higher equity market returns. This pattern holds for aggregate equity market returns and

Figure 1: Response of oil futures returns to crude oil inventory news

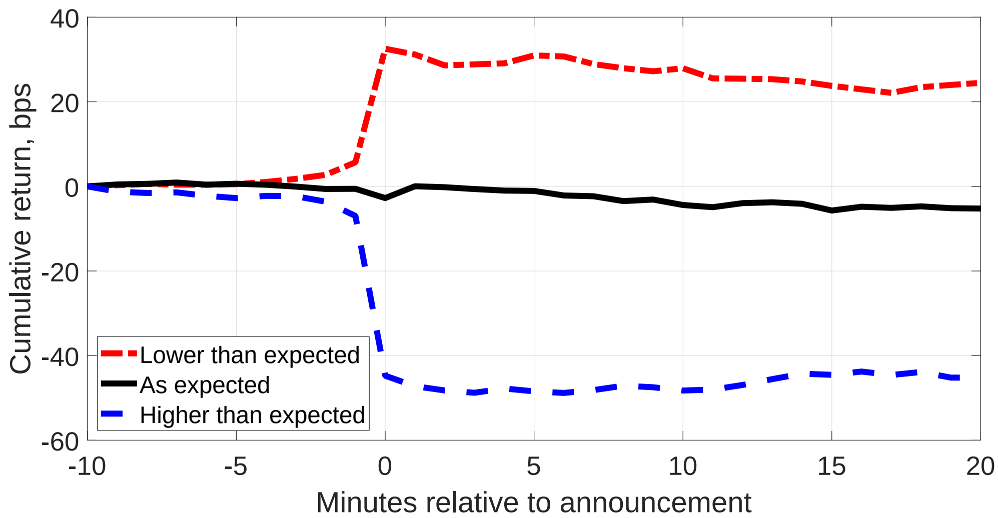

Crude oil inventory news is computed as the actual change in U.S. crude oil inventories as reported in the *Weekly Petroleum Status Report* minus the expected change according to the Bloomberg survey. *Lower than expected* stands for the average cumulative return of crude oil futures for the lowest 33% of the news, *As expected* stands for the average cumulative return of crude oil futures for the medium 33% of the news, and *Higher than expected* stands for the average cumulative return of crude oil futures for largest 33% news. *bps* stands for basis points. The sample period is 2003M10 to 2017M10.

is pervasive across different sectors, including those with limited direct exposure to energy prices, such as health care. The estimates for bond returns follow the reverse pattern. Bond futures returns tend to increase with higher oil prices pre-crisis and to decrease with higher oil prices post-crisis. While these results suggest that nominal interest rates became increasingly aligned with oil price fluctuations, the estimates are economically small and indicate that the effects of oil price changes on nominal interest rates are limited. Finally, we show that higher oil prices are associated with a depreciation of the U.S. dollar against a broad range of currencies, including those of major oil importers (such as the euro area).

Our empirical approach is based on instrumental variables (IV) estimation to capture the variation in oil prices related to the predetermined oil inventory news. We interpret the news as reflecting mostly oil-market-specific information, which induces different co-movement among oil and asset returns than more general macroeconomic news. Accordingly, the IV estimates for the equity market returns indicate a consistently lower correlation between oil returns and the equity market than ordinary least squares (OLS) regressions. This approach is consistent with the idea that general macroeconomic news induces a positive correlation between these variables. A similar pattern is obtained for exchange rates, where the depreciation of the U.S. dollar against other currencies is more pronounced for OLS regressions. These results highlight that oil-market-specific news tends to have very different effects on the relationship between oil prices and financial variables than more general macroeconomic news.[1]

News about inventories reflects the availability of oil for future consumption and the willingness of market participants to carry physical oil into the future. Thus, the oil price shocks associated with this news are different from the structural oil supply and demand shocks identified in widely used structural vector autoregressive (SVAR) models. Still, our results confirm several findings from this strand of literature, such as changes in the relationship between oil-market-specific demand shocks and U.S. stock market returns (Foroni, Guérin, and Marcellino 2017; Datta, Johannsen, Kwon, and Vigfusson 2018) and the depreciation of the U.S. dollar following higher oil prices (Kilian and Zhou 2019). This paper provides support for these findings using independent evidence and a different empirical approach.

Existing studies that examine the effect of macroeconomic news on stock markets have documented that "good" news can depress stock prices when market participants expect monetary policy to react to such news (Boyd, Hu, and Jagannathan 2005; Andersen, Bollerslev, Diebold, and Vega 2007). Following a similar argument, Datta, Johannsen, Kwon, and Vigfusson (2018) suggest that the inflation associated with higher oil prices may have been interpreted as good news during the zero lower bound episode covering much of the post-crisis sample. Although our results are broadly consistent with the idea that the importance of news about future growth and discount rates varies over the business cycle, the estimated response of bond prices to oil price shocks suggests that changes in monetary policy alone are an unlikely explanation for the time variation in the effects of oil price shocks on equity returns. Instead, changes in equity risk premia seem to play a larger role in explaining the positive effect of higher oil prices on equity returns after the crisis. More generally, these findings highlight the importance of looking at the reaction of a broader range of assets in order to understand the effect of news about oil market fundamentals on asset returns.

[1] This interpretation is also consistent with Kilian and Vega (2011), who document that between 1983 and 2008, daily oil prices did not react to U.S. macroeconomic news. Instead, oil prices did react strongly to oil inventory news, which highlights the oil-market-specific nature of such news.

2 Empirical framework

2.1 Measuring the effect of oil price shocks using oil inventory news

A central insight from recent research on the macroeconomic effects of oil price fluctuations is that oil prices are endogenous with respect to macroeconomic developments. Consider estimating β^i in

$$y_t^i = \mu^i + \beta^i r_t + \epsilon_t^i, \tag{1}$$

where y_t^i is the return on asset i, r_t is the oil futures return, and ϵ_t^i is an error term representing the effect of all other public and private news on the return on asset i. Endogeneity arises because asset prices affect oil prices, or because asset and oil prices jointly respond to news about current or expected economic conditions. Hence the oil return is correlated with the error term in equation (1), and the OLS estimates are biased and inconsistent.

Our empirical strategy is based on an IV approach that exploits variation in oil prices around the release of U.S. oil market statistics. The first stage IV regression takes the form of a standard announcement study,

$$r_t = \mu^r + \alpha^r z_t + \epsilon_t^r, \tag{2}$$

where r_t are the oil futures returns around the inventory announcement, z_t denotes the oil inventory news, and ϵ_t^r is an error term. The second stage then uses the predicted values from equation (2), $\widehat{r_t} = \widehat{\mu}^r + \widehat{\alpha}^r z_t$, to estimate β^i in equation (1).[2] Under the premise that oil market news affects asset returns through the price of oil, this IV approach identifies the causal effect of oil price shocks on asset markets and provides consistent estimates of β^i. By combining equations (2) and (1), one can see that the regression equation that relates asset return i to oil inventory news,

$$y_t^i = \bar{\mu}^i + \alpha^i z_t + \bar{\epsilon}_t^i, \tag{3}$$

can be interpreted as the reduced-form relationship of our IV model. The parameters α^r and α^i are difficult to interpret because the news does not have a natural unit. Instead, $\beta^i = \alpha^i/\alpha^r$ measures the reaction of asset returns to a 1% increase in the price of oil that is induced by the news about oil inventories.

[2]All IV estimations are performed using the standard 2SLS approach with robust standard errors.

2.2 Testing for time variation in the effect of oil price shocks on asset markets

Existing evidence suggests that the relationship between oil price fluctuations and asset returns experienced a dramatic shift around the onset of the financial crisis. This shift has been documented for the reduced-form correlation across a variety of different frequencies (see, e.g., Lombardi and Ravazzolo 2016; Aït-Sahalia and Xiu 2016), as well as in the context of structural oil market models (Foroni, Guérin, and Marcellino 2017; Datta, Johannsen, Kwon, and Vigfusson 2018). To test for a change in the effect of oil price shocks on financial variables, we interact r_t with a dummy variable d_t, which takes on the value of 1 for all observations after August 2008. Denoting the interaction term between oil futures returns and the dummy variables dr_t, we estimate the following equation:

$$y_t^i = \tilde{\mu}^i + \beta_1^i r_t + \beta_2^i dr_t + \tilde{\epsilon}_t^i, \tag{4}$$

where the endogenous variables r_t and dr_t are instrumented by the oil inventory news and their interaction with d_t.[3] The estimated marginal effects of oil price shocks on asset returns are given by $\widehat{\beta}_1^i$ for the period before 2008M8, and by $\widehat{\beta}_1^i + \widehat{\beta}_2^i$ for the post-crisis period. Standard errors for the marginal effects are obtained using the delta method. The p-values for the null hypothesis that β_2^i equals zero can be used to assess the statistical difference in the marginal effects across the different sample periods. As above, our hypothesis is that the IV estimates will be generally different from their OLS counterparts, obtained by estimating equation (4) by OLS.

3 Data description and implementation

3.1 Oil inventory news

Our data for commercial U.S. inventories of crude oil, gasoline, and distillates inventories are obtained from the U.S. Energy Information Administration (EIA) Weekly Petroleum Status Report. The report contains a set of U.S. oil market statistics for the week ending the Friday prior to the release. It is typically released on Wednesday at 10:30 a.m. Eastern Time for normal weeks, and on Thursday at 11:00 a.m. for weeks that include holidays.

Ahead of each release, market intelligence companies collect professional forecasters' and market participants' expectations about the most prominent variables in the release, namely

[3] The results are robust to the inclusion of d_t as an additional regressor. The specification in equation (4) does not include the main effect for d_t because d_t is uncorrelated with the other regressors and oil inventory news, which is consistent with theoretical predictions about high-frequency futures returns.

privately held crude oil, gasoline, and distillate inventories. This set of expectations allows us to measure the news component in the change of each type of inventory by subtracting the expected change in inventories from the actual change in inventories reported in the release:

$$\text{news}_t^{raw,i} = \text{actual}_t^i - \text{survey median}_t^i, \qquad i = \text{crude oil, gasoline, distillates,} \tag{5}$$

where *actual* refers to the reported change in inventories and *survey median* to the median expectation reported by Bloomberg. For the baseline results, we standardize the news by dividing the raw news by its sample standard deviation, $\widehat{\sigma}$:

$$\text{news}_t^i = \frac{\text{news}_t^{raw,i}}{\widehat{\sigma^i}}, \qquad i = \text{crude oil, gasoline, distillates.} \tag{6}$$

The timing of the data underlying the inventory news is important for our empirical approach to identify the effect of oil price shocks on financial variables. The approach relies on the premise that oil inventory news is exogenous with respect to the non-oil information flow at the time of the announcement. This should indeed be the case, as both the actual announcement and the expectations are determined well ahead of the release. The Weekly Petroleum Status Report refers to the weekly ending stocks for the previous week, while the survey measure is gathered and published on the day prior to the data release.

Table (1) shows the descriptive statistics for the oil inventory news during our sample period, which covers 2003M10 to 2017M10. The three inventory news measures are closely centered around zero, indicating that the forecast median is, on average, an unbiased predictor of the actual change in inventories. Crude oil inventory news correlates only weakly with news about gasoline and distillates inventories, with sample correlations of -0.04 and -0.12, respectively. Gasoline and distillate inventory news is modestly positively correlated, with a sample correlation of 0.31. Overall, the correlation among the news indicates that gasoline and distillate inventory news provides an additional source of information beyond that contained in the crude oil inventory release.

3.2 Oil futures and asset prices

To measure the effect of oil prices on financial variables, we combine a set of intraday data on West Texas intermediate (WTI) futures prices, equities, bond futures, and exchange rates from different data sources.[4] We use the nearby NYMEX WTI futures contract as a proxy for the oil spot price. Each nearby contract is rolled over into the next contract 7 trading days before its last trading day to avoid pricing issues related to any lack of

[4]The oil futures series are obtained from Portara. The equities price data are obtained from QuantQuote; the bond price data are obtained from the Chicago Mercantile Exchange; and the foreign exchange rate data are obtained from Reuters.

liquidity that might occur close to maturity. For equities, we rely on a set of exchange-traded funds (ETFs) tracking the aggregate S&P 500 (Ticker SPY) and nine different sub-sectors: Consumer Staples (XLP), Consumer Discretionary (XLY), Technology (XLK), Health Care (XLV), Industrials (XLI), Utilities (XLU), Materials (XLB), Energy (XLE), and Financial (XLF). To measure the effect of oil price shocks on interest rates, we use U.S. Treasury bond futures prices for 2-year, 5-year, and 10-year maturities. The foreign exchange rate data are measured against the U.S. dollar (USD) and consist of the currencies of net oil exporters, namely Canada (CAD) and Mexico (MXN); other commodities exporters, namely Australia (AUS), New Zealand (NZD), and South Africa (ZAR); and net oil importers, namely UK (GBP) and the euro area (EUR).

Table (2) shows the coverage of our data. The sample period for all data except the exchange rates are 2003M10 to 2017M10. The exchange rate data during the early sample period are somewhat limited, with major rates, AUD, CAD, EUR, and GBP, available from April 2006, NZD and ZAR from July 2007, and MXN from May 2009. This implies that tests for differences in the response of financial variables to oil price shocks across the different sample periods, 2003M10 to 2008M8 and 2008M9 to 2017M10, could have limited power for some exchange rate pairs and are infeasible for the case of the USD-MXN exchange rate.

3.3 Announcement and non-announcement returns

To exploit the variation in oil prices induced by news, we focus on a narrow window around the oil inventory announcements. Specifically, we compute returns for the $[-5; +10]$-minute window around the announcement, following practice in the existing literature (Gay, Simkins, and Turac 2009; Halova, Kurov, and Kucher 2014). As shown by the empirical results below, this interval allows for sufficient time for the news to be incorporated in oil futures and other asset prices, while ensuring that news about oil inventories are the main determinant of oil price changes during the event study.

A natural question is whether the IV approach provides different estimates than regular OLS estimates during "normal" times, i.e., outside of the announcement window. To test this idea, we estimate equation (4) using OLS and the same time-of-the-day returns (typically 10:25–10:40 a.m. Eastern Time) on the day prior to the inventory announcement. This accounts for potential intraday seasonality and makes our results directly comparable to previous studies that conducted a similar exercise.[5]

The original data were obtained as transactions in tick format or 1-minute intervals. We converted all time stamps to U.S. Eastern Time and deleted observations with implausibly large jumps. The tick data were aggregated to 1-minute data using the last transaction

[5]See, e.g., Halova, Kurov, and Kucher (2014).

within a given minute, classifying observations with no recorded transactions as missing. For every asset, 15-minute cumulative returns were computed by taking the last available price within a 5-minute window before the –5-minute and +10-minute mark around the announcement.[6]

The summary statistics of the returns of nearby oil futures and other asset returns are in table (2). The number of observations is smaller for the bond futures returns, which are most affected by our data-cleaning procedure. The table also compares the mean and standard variation of the variables on announcement and non-announcement days. For oil futures and energy stock returns, the return variation is significantly larger on announcement days than on non-announcement days, indicating a particularly strong effect of the oil inventory news on these variables. By contrast, the return variance is only slightly elevated on announcement days for the other financial returns.

4 Results

4.1 Is oil inventory news a strong instrument?

The first stage IV regression, displayed in table (3), shows that oil prices react strongly to the oil inventory news. Crude oil, gasoline, and distillates inventory news all enter statistically significantly and with the expected sign. The negative coefficient for these variables implies that higher-than-expected (lower-than-expected) inventories are associated with higher (lower) net supply and thus reduce (increase) the price of oil. The point estimates imply that before 2008M9, positive one-standard-deviation crude oil news decreases the price of crude oil by about 0.57%, while one-standard-deviation gasoline and distillate news decreases the price of crude by about 0.30% and 0.38%, respectively. According to the estimated coefficients associated with the interaction terms, these effects are smaller for the second sample period in the case of crude oil and distillate inventories. However, they remain statistically and economically significant, with a corresponding decrease of around 0.32% for crude oil inventory news and 0.12% for distillate inventory news.

Consistent with the weak correlations across the different types of news, these results show that oil product inventory news contains additional information for crude oil prices beyond that contained in crude oil inventory news. This evidence suggests that including the gasoline and distillate inventory news in the first-stage regression improves the prediction of the oil price returns during the announcement window and thus helps to sharpen the statistical inference around our IV estimates.

[6] As shown below in the robustness section, the results are robust to varying the size of the announcement window.

Indeed, crude oil, gasoline, and distillate inventory news variables and their interaction with the time dummy are also "strong" in a statistical sense. The F-statistics associated with the joint significance of all instruments is above 56, indicating that the instruments pass tests for weak instruments by a wide margin (see, e.g., Stock and Yogo 2005). Combined, the instruments explain about 33% of the variation in oil prices over the announcement window. Overall, these results confirm existing research showing that crude oil prices react strongly to oil inventory news (Halova, Kurov, and Kucher 2014).

4.2 The response of stock market returns to oil price shocks

Our main results are based on the IV specification that includes the time dummy for the post-2008M9 period, where we use the various oil inventory news variables as instruments for the oil price shocks and their interaction with the dummy variable.[7] Table (4) displays the regression results along with the estimated marginal effects for the aggregate U.S. stock market index and for 9 major sectors.

There is a strong difference in the reaction of aggregate U.S. stock market returns to changes in oil prices across the sample periods that mimics the changes in the unconditional correlations. Whereas aggregate stock returns decrease with higher oil prices during the pre-crisis period, they increase with higher oil prices in the post-crisis period. Both effects are economically significant: A 10% increase in oil prices is associated with a 0.8% decline (1.1% increase) in the aggregate stock market in the pre-crisis (post-crisis) period. The t-statistic for the coefficient associated with the interaction term also shows that the difference across the sample periods is statistically significant at all conventional significance levels.

This pattern is pervasive across most major sectors. For Materials, Financials, Industrials, Technology, Consumer Staples and Discretionaries, and Health Care, the marginal effects of oil price shocks on stock market returns are negative in the pre-2008M9 sample and positive thereafter. In all cases, the difference is statistically significant and of a magnitude similar to that of the change in the effect on the aggregate stock market index. The only exceptions to this pattern are the two sectors with the highest dependence on energy prices, Energy and Utilities. Unsurprisingly, we find a consistently positive effect from higher oil prices on stock market returns in the Energy sector. The effect appears to be relatively weaker during the second sample period, but it remains statistically and economically meaningful. In contrast to the other sectors, the effects of oil price shocks on returns in the Utilities sector are small during both samples and only marginally statistically significant in the post-2008M9 sample. Intuitively, companies in this sector are generally able to pass on much of the fluctuations in primary energy prices to consumers and thus have a natural

[7] See section 2.2 for a detailed description of the specification.

hedge against such fluctuations.

The results are also consistent with existing evidence on the transmission channel of oil price shocks on the U.S. economy that works through discretionary income. Research on this topic argues that higher oil prices reduce discretionary income and thus constitute demand shocks to most sectors (Lee and Ni 2002; Baumeister and Kilian 2016; Jo, Karnizova, and Reza 2019). Accordingly, we find that the sector that is most strongly exposed to discretionary income shocks, Consumer Discretionaries, displays a more adverse exposure to higher oil prices relative to other non-Energy sectors.

The contrast with the results from OLS regression on non-inventory days lends support to our empirical strategy of focusing on the variation in oil that can be traced to oil-market-specific news. Across all sectors except Energy, the IV estimates tend to be much lower than the OLS estimates, which are displayed in table (6). This finding is expected. The OLS estimates likely reflect the reaction of both stock market returns and oil price changes to general macroeconomic developments and news, and this news tends to move both variables in the same direction. The IV estimates are fundamentally different because they are driven by the news about future availability of oil. Still, one of our key findings is the marked level shift in the effect of oil price shocks on stock market returns that is pervasive across different sectors and estimation strategies.

4.3 The response of bond returns to oil price shocks

Table (5) displays the IV estimates of the effect of oil price shocks on bond futures prices. Compared to the stock market returns, the point estimates indicate a much weaker response of bond futures returns. Part of this could be due to the tendency of bond markets to be affected by liquidity shocks around sensitive announcements (Jiang, Lo, and Verdelhan 2011). Still, for all maturities, the marginal effects indicate a positive reaction of bond futures prices to higher oil prices in the first sample period. The effect is statistically significant at the 10% confidence level for the 2-year futures, at the 1% confidence level for the 5-year futures, and at the 5% confidence level for the 10-year futures. As for the stock returns, the results for the second sample period are markedly different, as highlighted by the t-statistic for the interaction term in the IV regression. During this period, 2-year and 5-year bond futures did not respond significantly, while 10-year bond futures responded negatively to higher oil prices.

Economically, however, all effects appear to be small. Before 2008M9, a 10% increase in the oil price is associated with a merely 0.6 basis point (bps) decrease in the nominal interest rate reflected in the 2-year and 5-year bond futures, and a 0.5 bps decrease in the interest rate of 10-year bond futures. Likewise, after the financial crisis, a 10% increase in the oil price is associated with a 0.7 bps increase in the nominal yield of 10-year bond futures, while the effect on the shorter maturities is smaller and statistically insignificant. The OLS estimates on non-inventory days, displayed in table (7), exhibit a similar pattern, but they are generally larger (in absolute value) than their IV counterparts. Thus, general macroeconomic news seems to be more important for the correlation between oil and bond futures returns than oil-market-specific news.

4.4 The response of exchange rate returns to oil price shocks

Table (5) displays the effect of oil price changes on the various exchange rates. Since our exchange rate data coverage for the pre-crisis period is limited, the results we obtain for this period need to be interpreted with some caution. Overall, however, the results point to a consistent pattern that higher oil prices are associated with a depreciation of the U.S. dollar. Not surprisingly, this depreciation is particularly strong against currencies of oil exporters (the Canadian dollar and the Mexican peso) and against those of other commodity-exporting countries (the Australian and New Zealand dollar and the South African rand).[8] Interestingly, the U.S. dollar also depreciates relative to the currencies of other oil-importing economies, the euro and the British pound. These results are consistent with recent evidence from a SVAR model of the global oil market and exchange rates that suggests that the U.S. exchange rate depreciates in the wake of various structural oil market shocks, at least in the short term (Kilian and Zhou 2019).

With the caveat of limited observations for the first sample period in mind, the interaction effects and marginal effects suggest that the effect of oil price changes on the various exchange rates increased during the second half of the sample. The comparison with the OLS regressions, shown in table (7), indicates that the effect of oil price shocks on exchange rates is somewhat weaker after oil-market-specific news than after more general macroeconomic news.

Overall, the results suggest an important causality that runs from oil prices to the U.S. dollar exchange rates that helps to explain the negative correlation observed between these variables during the last decades. The most common explanations for causal effects from oil prices are the terms-of-trade channel (see, e.g., Amano and Van Norden 1998a; Amano and Van Norden 1998b) or a wealth transfer (and portfolio) channel (see, e.g., Bodenstein, Erceg, and Guerrieri 2011). Both channels imply a weakening of the currency of oil-importing economies relative to those of oil exporters and thus rationalize the stronger effects observed for those currencies. The strength of such effects also depends on oil-exporting investors'

[8] For the pre-crisis period, the estimated coefficient is insignificant for NZD and negative for ZAR, likely due to the fact our dataset only comprises these exchange rates from 2007.

preferences for reinvesting their windfalls in a particular currency, which is one potential explanation for the depreciation of the U.S. dollar against currencies of other oil-importing economies, in particular the euro.

4.5 Robustness exercises

In this section, we show that our qualitative findings are robust to specific modelling choices, such as the standardization of news, the inclusion of gasoline and distillate inventory news in addition to crude oil inventory news, and the size of the announcement window.

The first robustness exercise reconsiders the baseline IV specification, which is based on three different instruments. Including several instruments has the advantage of sharpening the inference in IV regressions. This idea is supported by the first-stage regression, which highlights significant explanatory power from gasoline and distillate inventory news that is orthogonal to the information contained in crude oil inventory news. Still, the introduction of additional sources of variation could bias the IV estimates if the additional instruments are not valid.

Tables (8), (9), and (10) display the results from an alternative specification, which relies only on crude oil inventory news and its interaction with the time dummy to instrument the oil futures returns. The first two columns of these tables show the estimated marginal effects along with the benchmark estimates for the S&P 500 Index ETF, 10-year bond futures, and CAD returns.[9] For all three series, the point estimates for the marginal effects are very similar across the two specifications. The biggest difference can be observed for the estimates of the marginal effect of oil price shocks on the Canadian dollar exchange rate during the post-crisis period, but even in this case, the estimate only changes slightly from 0.09 (baseline specification) to 0.115 (alternative specification). Instead, the principal effect of dropping the gasoline and distillate inventory news is an increase in the standard errors, which reduces the statistical significance of some coefficients without changing their magnitude.[10]

The second robustness exercise scrutinizes the standardization of oil inventory news. The baseline results are computed with news that is normalized to have unit standard deviation. In practice, this normalized news is obtained from dividing the "raw" surprises by their sample (time-series) standard deviation. While this standardization is common in the announcement literature, using the sample standard deviation estimated over the entire sample period might underweight the surprise component in raw news around more tranquil

[9] We focus on these three series to conserve space. Similarly close results were obtained for the other series.

[10] The robustness with respect to this specification is also supported by over-identification tests, which do not reject the null hypothesis that all instruments are valid. The p-values associated with the test statistic according to Wooldridge (1995) are 0.14 for the S&P Index returns, 0.75 for the 10-year bond futures returns, and 0.17 for the Canadian dollar exchange rate returns.

sub-samples. To gauge the importance of this effect, we also performed robustness checks based on a different standardization, which exploits the cross-sectional dispersion inventory forecasts across different forecasters.

Specifically, we consider the alternative news, $\text{news}_t^{i,alt}$, defined as

$$\text{news}_t^{i,alt} = \frac{\text{news}_t^{raw,i}}{\widehat{\sigma_t^i}}, \qquad i = \text{crude oil, gasoline, distillates,} \tag{7}$$

where $\widehat{\sigma_t^i}$ corresponds to the cross-sectional standard deviation of the expectations of changes in inventories across different forecasters. Intuitively, this standardization increases the weight on surprises that feature a lot of forecaster agreement ex ante, while it decreases the weight on surprises that feature low forecaster agreement. In contrast to the baseline standardization using the time-series standard deviation, this form of standardization does not contain a look-ahead bias.

Although this form of standardization is very different from the benchmark news, both types of news are highly correlated. The correlations between the news standardized via equations (6) and (7) are 0.95 for crude oil and gasoline and 0.94 for distillate. Consequently, the point estimates for the IV specification that relies on the alternative news, which are displayed in column (3) of tables (8), (9), and (10), are almost identical to the benchmark estimates. The estimates for the standard errors are similarly close, suggesting that the benchmark results do not depend on a specific choice of normalization of the news.

The final exercise shows that our results are robust to the size of the announcement window, which in the benchmark estimations is 15 minutes. We consider results from two different window sizes, 5 minutes and 30 minutes, which are displayed in columns (4) and (5) of tables (8), (9), and (10). Again, the point estimates for the shorter and longer window sizes are very close to the benchmark results. Only the marginal effects for the 10-year bond futures returns for the 30-minute window appear somewhat attenuated, which is likely due to the additional noise that is introduced with the larger size of the event window. The primary effect of increasing the window size appears to be an increase in the standard errors, which renders some of the coefficients less significant than suggested by the benchmark results. By contrast, the estimated standard errors are somewhat tighter for the 5-minute window than for the 15-minute window. However, the statistical and economic conclusions remain very similar across the two specifications.

5 Interpreting the time-varying effects of oil price shocks

5.1 Oil inventory news and structural shocks

A key insight from our empirical results is that the response of stock market returns to inventory-induced oil price shocks changed significantly around the 2007/08 crisis. Kilian (2009), Kilian and Park (2009), and subsequent studies have shown that the effects of oil price changes on the macroeconomy and financial variables depend on the nature of the underlying oil price shock. Generally, changes in inventories reflect a composition of various structural forces affecting the oil market. Thus, a potential explanation for our finding is that the informational content of oil inventories in terms of the underlying structural shocks changed after the 2007/08 crisis. For example, aggregate demand shocks, which are typically associated with a positive co-movement between oil price changes and stock market returns, might have contributed more strongly to oil inventory surprises after the financial crisis.

Although no high-frequency structural model exists that would allow us to evaluate this hypothesis directly, there are reasons to believe that the changing informational content of oil inventory news in terms of its structural decomposition is not the main driver of the time variation in the oil–stock return relationship. First, Foroni, Guérin, and Marcellino (2017) and Datta, Johannsen, Kwon, and Vigfusson (2018) document similar time variations in structural models of the global oil market and stock market returns. These models explicitly control for global oil supply and aggregate demand shocks. Second, changes in global crude oil inventories appear to be predominantly driven by oil supply and other oil-market-specific demand shocks, while aggregate demand shocks only have an insignificant impact on these inventories (Kilian and Murphy 2014). In a similar fashion, surprise changes in U.S. oil inventories should mainly reflect oil supply and oil-market-specific shocks. This idea is also supported by the cross-sectional responses of stock market returns across various sectors. Existing empirical evidence suggests that energy stocks respond particulary strongly to oil demand shocks (Kilian and Park 2009; Ready 2017). By contrast, we document a markedly lower reaction of energy stocks to oil price changes related to oil inventory news over the post-crisis period. For the stocks of other sectors, the differences between the estimates from the IV and OLS approaches are roughly constant. This suggests that the time variation in the stock return response to oil price shocks results from a different reaction to the same type of news, rather than the changing structural composition of the news.

Yet another way of gauging the changing informational content of oil inventory news is via a reduced-form decomposition of changes in inventories that relies on additional data released in the Weekly Petroleum Status Report. Intuitively, any structural changes in the underlying drivers of oil inventory news should also be reflected in a changing relationship between the reduced-form variables.

The decomposition relies on the mathematical identity that the change in inventories is equal to the difference between flow production and flow consumption. While the commercial inventory numbers are considered to be the key statistics of the report, the report also contains other information on U.S. oil market variables that allows us to compute the decomposition for the case of crude oil inventories. Changes in U.S. commercial crude oil inventories are roughly equal to the sum of U.S. oil production and net imports, minus refinery intakes:[11]

$$\Delta \text{Inventories}_t^{US,Commercial} \approx \text{Production}_t^{US} + \text{Net Imports}_t^{US} - \text{Refinery Inputs}_t^{US}. \tag{8}$$

One can see by differencing equation (8) that changes in the change of commercial crude oil inventories can be attributed to a combination of changes in production, net imports, and refinery inputs.

A natural hypothesis is that changes in the structural shocks underlying inventory news should also be reflected in a changing composition of oil inventory news in terms of its reduced-form components. To investigate this hypothesis, table (11) reports the regressions of the changes in the change in U.S. crude oil inventories on changes in U.S. crude oil production, net imports, and refinery inputs for the pre-crisis and post-crises periods. The results suggest that the relationship between inventories and production, refinery inputs, and imports was relatively stable over time. The coefficients associated with changes in U.S. oil production and refinery intake increase in magnitude during the post-crisis period. However, the R-squareds associated with these variables are close to zero in both samples, suggesting that these variables have limited explanatory power for changes in crude oil inventories. By contrast, the R-squareds associated with changes in U.S. net imports of crude oil is well over 60% in both periods. This tight relationship suggests that news about U.S. crude oil inventories mainly reflects information about U.S. net imports of crude oil. However, there appears to be very limited evidence for a change in the primitive components of crude oil inventories and thus for a changing informational content of this news in terms of structural shocks.

5.2 Interpretation via the Gordon decomposition

Instead of focusing on the role of different structural shocks, research in finance has highlighted the idea that the effect of macroeconomic news on the stock market is time-varying because it is state dependent. For example, good macroeconomic news, such as lower unem-

[11] The identity holds only approximately here because it does not take into account changes in strategic petroleum reserves.

ployment, tends to have a negative effect on stock market returns during economic expansions, and a positive effect during contractions (Boyd, Hu, and Jagannathan 2005). Formally, these authors show via the Gordon decomposition that the effect of news on stock market returns can be decomposed into the effect on the three determinants of stock prices, namely on information about future interest rates, on the equity risk premium, and on economic growth and the associated corporate earnings and dividends. This decomposition allows us to investigate the time-varying response of stock market returns to oil price shocks through the underlying drivers of stock market fluctuations rather than the structural oil price shocks determining oil inventories.

We first address the relationship between oil inventory news and dividends. Similar to existing studies, we focus on U.S. economic growth as a proxy for expected dividends, which are not observable at high frequency (Boyd, Hu, and Jagannathan 2005). Table (12) displays regressions on 1- to 4-week-ahead changes in the Philadelphia Fed's Aruoba-Diebod-Scotti U.S. Business Conditions (ADS) index on changes in crude oil inventory news. The specifications also include a time dummy for the post-crisis period, which is interacted with the oil inventory news. This specification allows us to test whether a predictive relationship exists and whether it has changed over time.

The results provide no evidence that oil inventory news predicts changes in the ADS index. The coefficients associated with the news and their interaction with the time dummy are small and statistically insignificant, and the explanatory power of these variables in terms of the R-squared is negligible. These results are consistent with Law, Song, and Yaron (2018), who report that cash flow news did not contribute meaningfully to the time variation in stock return sensitivity to macroeconomic news. They are also consistent with the previous section, which argues that changes in the informational content of oil inventory news with respect to current economic conditions is likely to be small.

Andersen, Bollerslev, Diebold, and Vega (2007), Boyd, Hu, and Jagannathan (2005), and Law, Song, and Yaron (2018) argue that the monetary policy stance has important effects on the response of stock market returns to news. Thus, good news in the macroeconomic sense tends to be bad news for the stock market during expansions, when monetary policy tends to be tighter and such news increases the likelihood of higher real interest rates in the future. By contrast, the same good news tends to be good news for the stock market during recessions, when monetary policy tends to be less tight. Datta, Johannsen, Kwon, and Vigfusson (2018) have recently argued that this effect might have been particularly strong during the post-crisis period, when the U.S. policy rate was at or near the zero lower bound.

In terms of our setting, this explanation suggests that oil price shocks induced by oil inventory news had a smaller (i.e., more negative) effect on future interest rates during the post-crisis period. Because our empirical exercise covers bond futures returns, we can evaluate the effect of oil price changes on nominal rates directly. Overall, the estimates suggest that the effects of oil price shocks on current nominal bond returns are economically small. If anything, higher oil prices were related to *higher* nominal bond prices and thus to *lower* nominal interest rates in the pre-crisis period, and to *lower* nominal bond prices and thus to *higher* nominal interest rates in the post-crisis period. Thus, changes in monetary policy stance alone are an unlikely explanation for the time variation in the effect of oil price shocks on stock market returns.

By elimination, our results thus suggest that oil prices could instead have been increasingly related to equity risk premia in the post-crisis period. A potential channel through which oil prices affect equity risk premia is inflation expectations. In fact, oil prices are frequently associated with inflation expectations (Coibion and Gorodnichenko 2015). At the same time, recent research highlights the importance of time variation in the inflation risk premium for stock market returns (Boons, Duarte, Roon, and Szymanowska 2019). A natural conjecture is that during the pre-crisis period—a period of comparatively high inflation expectations—higher oil prices demanded higher inflation risk compensation and thus tended to decrease stock prices on impact. By contrast, the inflation risk premium might have reversed in the post-crisis period, when deflationary pressures prevailed.

6 Conclusion

In this paper, we use weekly crude oil inventory announcements to study the transmission of news about oil market fundamentals to asset prices. Across most sectors, equity prices responded negatively to higher oil prices before the crisis, but they exhibited a markedly positive response after it. By contrast, we find little evidence that bond prices responded meaningfully to oil price changes, although we do document a consistent depreciation of the U.S. dollar against a range of other currencies, including those of other oil-importing economies, in response to higher oil prices.

The response of interest rates to oil prices, in combination with the time-varying effect of oil prices on equity returns, suggests that oil prices were increasingly related to equity risk premia in the post-crisis period. We also investigate whether the informational content of U.S. oil inventories about global oil supply or demand conditions changed over time, but we find little evidence for this claim. More generally, our results show that oil price changes associated with inventory news have, on average, a more negative effect on U.S. stock returns than general news arrivals, highlighting this particular transmission mechanism for oil-market-specific news.

References

Aït-Sahalia, Yacine and Dacheng Xiu (2016). “Increased correlation among asset classes: Are volatility or jumps to blame, or both?” *Journal of Econometrics* 194.2, pp. 205–219.

Akram, Q Farooq (2004). “Oil prices and exchange rates: Norwegian evidence.” *The Econometrics Journal* 7.2, pp. 476–504.

Akram, Q Farooq (2009). “Commodity prices, interest rates and the dollar.” *Energy Economics* 31.6, pp. 838–851.

Alquist, Ron, Gregory H Bauer, and Antonio Diez de los Rios (2014). “What does the convenience yield curve tell us about the crude oil market?” *Bank of Canada Working Paper.*

Alquist, Ron and Lutz Kilian (2010). “What do we learn from the price of crude oil futures?” *Journal of Applied Econometrics* 25.4, pp. 539–573.

Amano, Robert A and Simon Van Norden (1998a). “Exchange rates and oil prices.” *Review of International Economics* 6.4, pp. 683–694.

Amano, Robert A and Simon Van Norden (1998b). “Oil prices and the rise and fall of the U.S. real exchange rate.” *Journal of International Money and Finance* 17.2, pp. 299–316.

Andersen, Torben G, Tim Bollerslev, Francis X Diebold, and Clara Vega (2007). “Real-time price discovery in global stock, bond and foreign exchange markets.” *Journal of International Economics* 73.2, pp. 251–277.

Baumeister, Christiane and Lutz Kilian (2016). “Lower oil prices and the US economy: Is this time different?” *Brookings Papers on Economic Activity* 2016.2, pp. 287–357.

Bodenstein, Martin, Christopher J Erceg, and Luca Guerrieri (2011). “Oil shocks and external adjustment.” *Journal of International Economics* 83.2, pp. 168–184.

Boons, Martijn, Fernando Duarte, Frans de Roon, and Marta Szymanowska (2019). “Time-varying inflation risk and stock returns.” *Journal of Financial Economics* (forthcoming).

Boyd, John H, Jian Hu, and Ravi Jagannathan (2005). “The stock market’s reaction to unemployment news: Why bad news is usually good for stocks.” *The Journal of Finance* 60.2, pp. 649–672.

Chen, Yu-Chin, Kenneth Rogoff, and Barbara Rossi (2010). “Can exchange rates forecast commodity prices?” *Quarterly Journal of Economics* 125, pp. 1145–1194.

Coibion, Olivier and Yuriy Gorodnichenko (2015). “Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation.” *American Economic Journal: Macroeconomics* 7.1, pp. 197–232.

Datta, Deepa, Benjamin Kramer Johannsen, Hannah Kwon, and Robert J Vigfusson (2018). “Oil, equities, and the zero lower bound.” *Finance and Economics Discussion Paper No. 2018-058, Federal Reserve Board.*

Ferraro, Domenico, Kenneth Rogoff, and Barbara Rossi (2015). "Can oil prices forecast exchange rates? An empirical analysis of the relationship between commodity prices and exchange rates." *Journal of International Money and Finance* 54, pp. 116–141.

Foroni, Claudia, Pierre Guérin, and Massimiliano Marcellino (2017). "Explaining the time-varying effects of oil market shocks on US stock returns." *Economics Letters* 155, pp. 84–88.

Frankel, Jeffrey A (2008). "The effect of monetary policy on real commodity prices." *In: Asset Prices and Monetary Policy.* University of Chicago Press, pp. 291–333.

Gao, Lin, Steffen Hitzemann, Ivan Shaliastovich, and Lai Xu (2017). "Oil volatility risk." *Working Paper, University of California Riverside.*

Gay, Gerald D, Betty J Simkins, and Marian Turac (2009). "Analyst forecasts and price discovery in futures markets: The case of natural gas storage." *Journal of Futures Markets* 29.5, pp. 451–477.

Halova, Marketa W, Alexander Kurov, and Oleg Kucher (2014). "Noisy inventory announcements and energy prices." *Journal of Futures Markets* 34.10, pp. 911–933.

Hitzemann, Steffen (2016). "Macroeconomic fluctuations, oil supply shocks, and equilibrium oil futures prices." *SSRN Working Paper.*

Jiang, George J, Ingrid Lo, and Adrien Verdelhan (2011). "Information shocks, liquidity shocks, jumps, and price discovery: Evidence from the U.S. Treasury market." *Journal of Financial and Quantitative Analysis* 46.2, pp. 527–551.

Jo, Soojin, Lilia Karnizova, and Abeer Reza (2019). "Industry effects of oil price shocks: A re-examination." *Energy Economics* 82, pp. 179–190.

Kilian, Lutz (2009). "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market." *The American Economic Review* 99.3, pp. 1053–1069.

Kilian, Lutz and Thomas K Lee (2014). "Quantifying the speculative component in the real price of oil: The role of global oil inventories." *Journal of International Money and Finance* 42, pp. 71–87.

Kilian, Lutz and Daniel P Murphy (2014). "The role of inventories and speculative trading in the global market for crude oil." *Journal of Applied Econometrics* 29.3, pp. 454–478.

Kilian, Lutz and Cheolbeom Park (2009). "The impact of oil price shocks on the U.S. stock market." *International Economic Review* 50.4, pp. 1267–1287.

Kilian, Lutz and Clara Vega (2011). "Do energy prices respond to US macroeconomic news? A test of the hypothesis of predetermined energy prices." *Review of Economics and Statistics* 93.2, pp. 660–671.

Kilian, Lutz and Xiaoqing Zhou (2019). "Oil prices, exchange rates and interest rates." *CEPR Discussion Paper No. DP13478.*

Law, Tzuo Hann, Dongho Song, and Amir Yaron (2018). “Fearing the Fed: How Wall Street reads Main Street.” *SSRN Working Paper*.

Lee, Kiseok and Shawn Ni (2002). “On the dynamic effects of oil price shocks: A study using industry level data.” *Journal of Monetary Economics* 49.4, pp. 823–852.

Lombardi, Marco J and Francesco Ravazzolo (2016). “On the correlation between commodity and equity returns: Implications for portfolio allocation.” *Journal of Commodity Markets* 2.1, pp. 45–57.

Ready, Robert C (2017). “Oil prices and the stock market.” *Review of Finance* 22.1, pp. 155–176.

Stock, James H. and Motohiro Yogo (2005). *Testing for weak instruments in linear IV regression*. Ed. by Donald W. K. Andrews and James H. Stock. Cambridge: Cambridge University Press.

Wooldridge, Jeffrey M (1995). *Score diagnostics for linear models estimated by two stage least squares*. Ed. by G.S. Maddala, P.C.B. Phillips, and T.N. Srinivasan. Oxford: Blackwell.

Tables

Table 1: Summary statistics for the weekly inventory news

	obs	mean	std	Correlation with Crude oil	Gasoline	Distillates
Crude oil	733	0.003	1	1		
Gasoline	733	0.029	1	-0.04	1	
Distillates	733	0.021	1	-0.12	0.31	1

The news is computed as the difference between the actual release and the survey median as reported by Bloomberg and is normalized to have unit standard deviation. *obs* stands for the number of observations, *mean* for the sample mean, and *std* for the sample standard deviation. The sample period is 2003M10 to 2017M10.

Table 2: Summary statistics for oil futures and asset returns

		Announcement days			Non-announcement days		
	start date	obs	mean	std	obs	mean	std
Oil futures							
Nearby WTI contract	01/10/2003	733	-7.8	91.4	732	0.4	39.8
US equity ETFs							
SPY (Index)	01/10/2003	733	-0.4	19.9	732	0.5	17.5
XLB (Material)	01/10/2003	730	-1.2	30.8	731	0.9	25.5
XLE (Energy)	01/10/2003	731	-0.9	54.1	731	1.0	30.1
XLF (Financial)	01/10/2003	731	0.5	28.0	732	0.1	26.9
XLI (Industrial)	01/10/2003	726	-0.5	24.0	726	0.6	20.8
XLK (Technology)	01/10/2003	732	0.2	22.8	731	1.2	21.0
XLP (Consumer Staple)	01/10/2003	725	0.3	16.1	726	1.0	14.6
XLU (Utility)	01/10/2003	731	-0.6	21.1	731	1.2	19.9
XLV (Health Care)	01/10/2003	725	-0.4	18.3	728	0.1	16.7
XLY (Consumer Discretionary)	01/10/2003	705	-0.4	25.4	715	0.5	21.6
U.S. Treasury bond futures							
2-year	01/10/2003	714	0.0	0.5	712	0.0	0.4
5-year	01/10/2003	710	0.0	1.3	714	0.0	1.0
10-year	01/10/2003	696	0.0	1.8	700	0.0	1.5
Exchange rates							
AUD	16/04/2006	601	-0.7	13.0	600	-0.1	12.1
CAD	16/04/2006	601	-1.1	12.7	600	-0.1	9.5
EUR	16/04/2006	558	-0.5	10.3	570	-0.2	9.6
GBP	16/04/2006	601	0.0	9.3	599	-0.4	8.9
MXN	06/05/2009	442	-0.5	12.7	440	0.2	11.5
NZD	27/07/2007	533	-0.9	13.7	531	0.3	12.9
ZAR	27/07/2007	525	-1.5	19.2	527	0.7	16.6

The returns are computed as the cumulative returns for a 15-minute window around the Weekly Petroleum Status Report release for announcement days and as the cumulative returns for the same 15-minute window on the previous day for the non-announcement returns. Returns are expressed in bps. *obs* stands for the number of observations, *mean* for the sample mean, and *std* for the sample standard deviation. All exchange rates are measured against the U.S. dollar. The sample period is 2003M10 to 2017M10.

Table 3: First-stage regression results for the two endogenous variables, oil futures returns and their interaction with the time dummy

	Oil futures returns	Oil futures returns × d_t
Crude Oil inventory news	-56.56***	-0.15
	(4.54)	(0.40)
Gasoline inventory news	-29.70***	-0.16
	(3.98)	(0.40)
Distillate inventory news	-37.81***	-0.15
	(5.54)	(0.47)
Crude Oil inventory news × d_t	24.86***	-31.58***
	(5.80)	(3.63)
Gasoline inventory news × d_t	2.12	-27.50***
	(5.83)	(4.26)
Distillate inventory news × d_t	25.33***	-12.35***
	(6.79)	(3.94)
Const	-6.91**	-5.44**
	(2.72)	(2.32)
Observations	733	733
R-squared	0.33	0.26
F-stat	56.20	18.65

d_t stands for a dummy variable that takes on the value of 1 for all observations after 2008M8. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 4: IV estimates of the effect of oil price shocks on U.S. equity returns

	S&P Index	Materials	Energy	Financial	Industrial	Technology	Con. Staple	Utility	Health Care	Con. Disc.
Oil futures returns	-0.081***	-0.032	0.627***	-0.220***	-0.156***	-0.170***	-0.149***	-0.019	-0.120***	-0.244***
	(0.016)	(0.026)	(0.048)	(0.034)	(0.022)	(0.022)	(0.016)	(0.021)	(0.017)	(0.032)
Oil futures returns × d_t	0.191***	0.214***	-0.128**	0.286***	0.260***	0.259***	0.176***	0.060*	0.158***	0.303***
	(0.027)	(0.039)	(0.061)	(0.043)	(0.033)	(0.032)	(0.023)	(0.031)	(0.026)	(0.041)
Constant	0.329	0.117	2.953**	0.888	0.102	0.784	0.329	-0.328	-0.245	-0.389
	(0.644)	(0.986)	(1.214)	(0.978)	(0.801)	(0.760)	(0.539)	(0.733)	(0.653)	(0.868)
Observations	733	730	731	731	726	732	725	731	725	705
R-squared	0.24	0.24	0.63	0.14	0.19	0.18	0.15	0.05	0.13	0.18
F-stat	51.75	39.47	337.40	47.37	68.25	75.66	85.99	4.09	51.68	62.26
					Marginal effects					
pre-2008M9	-0.081***	-0.032	0.627***	-0.219***	-0.156***	-0.170***	-0.149***	-0.019	-0.120***	-0.244***
	(0.016)	(0.026)	(0.048)	(0.034)	(0.022)	(0.022)	(0.016)	(0.021)	(0.017)	(0.032)
post-2008M9	0.110***	0.183***	0.499***	0.066***	0.104***	0.089***	0.028*	0.041*	0.038**	0.059**
	(0.021)	(0.030)	(0.039)	(0.025)	(0.025)	(0.023)	(0.016)	(0.023)	(0.019)	(0.026)

d_t stands for a dummy variable that takes on the value of 1 for all observations after 2008M8. *Con.* stands for Consumer and *Disc.* for Discretionary. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 5: IV estimates of the effect of oil price shocks on U.S. bond futures and exchange rate returns

	Bond futures			*Exchange rates*						
	2-year	5-year	10-year	AUS	CAD	EUR	GBP	MXN	NZD	ZAR
Oil futures returns	0.0012*	0.0030***	0.0042**	0.024*	0.037***	0.031***	0.006		0.007	-0.052**
	(0.0006)	(0.0012)	(0.0019)	(0.013)	(0.012)	(0.010)	(0.008)		(0.018)	(0.025)
Oil futures returns × d_t	-0.0017**	-0.0052***	-0.0107***	0.035*	0.053***	-0.006	0.020	0.065***	0.038*	0.109***
	(0.0009)	(0.0024)	(0.0027)	(0.018)	(0.016)	(0.017)	(0.012)	(0.014)	(0.022)	(0.031)
Constant	-0.012	-0.043	-0.061	-0.156	-0.317	-0.197	0.278	-0.193	-0.443	-0.973
	(0.017)	(0.044)	(0.066)	(0.457)	(0.402)	(0.417)	(0.368)	(0.521)	(0.537)	(0.771)
Observations	714	710	696	601	601	558	601	442	504	525
R-squared	0.04	0.05	0.06	0.23	0.39	0.07	0.06	0.25	0.18	0.16
F-stat	4.12	6.95	16.38	25.73	100.08	15.49	8.19	21.92	14.73	14.67
					Marginal effects					
pre-2008M9	0.0012*	0.0030***	0.0042**	0.024*	0.037***	0.031***	0.006		0.007	-0.052**
	(0.0006)	(0.0012)	(0.0019)	(0.013)	(0.012)	(0.010)	(0.008)		(0.018)	(0.025)
post-2008M9	0.0005	-0.0022	-0.0066***	0.059***	0.090***	0.025***	0.026***	0.065***	0.045***	0.057***
	(-0.0007)	(0.0020)	(0.0019)	(0.012)	(0.009)	(0.010)	(0.009)	(0.014)	(0.012)	(0.017)

d_t stands for a dummy variable that takes on the value of 1 for all observations after 2008M8. All exchange rates are measured against the U.S. dollar. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 6: OLS estimates of the effect of oil price shocks on U.S. equity returns

	S&P Index	Materials	Energy	Financial	Industrial	Technology	Con. Staple	Utility	Health Care	Con. Disc.
Oil futures returns	-0.045**	-0.013	0.361***	-0.093***	-0.065***	-0.126***	-0.079***	-0.04	-0.075***	-0.114***
	(-0.022)	(0.034)	(0.053)	(0.031)	(0.023)	(0.031)	(0.021)	(0.027)	(0.020)	(0.032)
Oil futures returns × d_t	0.276***	0.337***	0.137*	0.345***	0.330***	0.350***	0.207***	0.162***	0.208***	0.368***
	(0.039)	(0.056)	(0.067)	(0.058)	(0.048)	(0.046)	(0.032)	(0.040)	(0.035)	(0.050)
Constant	0.242	0.495	0.674	-0.206	0.231	0.896	0.803	1.057	-0.080	0.157
	(0.588)	(0.861)	(0.893)	(0.947)	(0.703)	(0.726)	(0.517)	(0.732)	(0.597)	(0.745)
Observations	732	731	731	732	726	731	726	731	728	715
R-squared	0.18	0.16	0.36	0.10	0.17	0.14	0.10	0.04	0.08	0.16
F-stat	27.25	27.11	95.74	17.45	23.85	30.02	21.84	9.26	17.68	27.87

d_t stands for a dummy variable that takes on the value of 1 for all observations after 2008M8. *Con.* stands for Consumer and *Disc.* for Discretionary. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 7: OLS estimates of the effect of oil price shocks on U.S. bond futures and exchange rate returns

	Bond futures			*Exchange rates*						
	2-year	5-year	10-year	AUD	CAD	EUR	GBP	MXN	NZD	ZAR
Oil futures returns	0.0019***	0.0048***	0.0097***	0.046**	0.032	0.034*	0.015		0.086**	-0.012
	(0.0006)	(0.0014)	(0.0022)	(0.019)	(0.022)	(0.020)	(0.016)		(0.038)	(0.076)
Oil futures returns × d_t	-0.0028***	-0.0089***	-0.021***	0.073**	0.082***	0.060**	0.084***	0.175***	0.020	0.167**
	(0.0009)	(0.0021)	(0.003)	(0.029)	(0.030)	(0.026)	(0.022)	(0.028)	(0.045)	(0.081)
Constant	0.010	0.034	0.061	-0.210	-0.209	-0.379	-0.534	-0.159	0.105	0.358
	(0.014)	(0.038)	(0.057)	(0.471)	(0.358)	(0.378)	(0.336)	(0.457)	(0.536)	(0.682)
Observations	712	714	700	600	600	570	599	440	531	527
R-squared	0.02	0.03	0.07	0.12	0.18	0.12	0.15	0.29	0.10	0.11
F-stat	5.56	9.14	22.80	18.52	18.02	18.55	21.20	39.47	12.55	17.38

d_t stands for a dummy variable that takes on the value of 1 for all observations after 2008M8. All exchange rates are measured against the U.S. dollar. Robust standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. The sample period is 2003M10 to 2017M10.

Table 8: Robustness exercises for the S&P Index returns

	Marginal effects				
	(1)	(2)	(3)	(4)	(5)
pre-2008M9	-0.081***	-0.089***	-0.091***	-0.068***	-0.081***
	(0.016)	(0.024)	(0.017)	(0.014)	(0.022)
post-2008M9	0.110***	0.112***	0.101***	0.084***	0.116***
	(0.021)	(0.026)	(0.021)	(0.013)	(0.031)
Observations	733	733	733	733	733
R-squared	0.24	0.15	0.23	0.26	0.19
F-stat	51.75	63.78	50.63	64.93	27.73

Column (1) presents the baseline IV estimates, column (2) presents the IV estimates using the crude oil inventory news and its interaction with the time dummy as the only instruments, column (3) presents the estimates using an alternative standardization of the news based on the cross-sectional dispersion of inventory forecasts, column (4) presents the results for returns during a 5-minute announcement window, and column (5) presents the estimates for returns during a 30-minute announcement window. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 9: Robustness exercises for the 10-year bond futures returns

	Marginal effects				
	(1)	(2)	(3)	(4)	(5)
pre-2008M9	0.004**	0.004	0.004**	0.005***	0.003
	(0.002)	(0.003)	(0.002)	(0.001)	(0.003)
post-2008M9	-0.007***	-0.008***	-0.007***	-0.010***	-0.005*
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)
Observations	696	696	696	707	682
R-squared	0.06	0.05	0.06	0.11	0.05
F-stat	16.38	11.18	15.64	37.46	5.11

Column (1) presents the baseline IV estimates, column (2) presents the IV estimates using the crude oil inventory news and its interaction with the time dummy as the only instruments, column (3) presents the estimates using an alternative standardization of the news based on the cross-sectional dispersion of inventory forecasts, column (4) presents the results for returns during a 5-minute announcement window, and column (5) presents the estimates for returns during a 30-minute announcement window. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 10: Robustness exercises for Canadian dollar–U.S. dollar exchange rate returns

	Marginal effects				
	(1)	(2)	(3)	(4)	(5)
pre-2008M9	0.037***	0.033**	0.034***	0.048***	0.032**
	(0.012)	(0.014)	(0.013)	(0.009)	(0.015)
post-2008M9	0.090***	0.115***	0.091***	0.103***	0.089***
	(0.009)	(0.014)	(0.009)	(0.008)	(0.013)
Observations	601	601	601	601	601
R-squared	0.39	0.37	0.39	0.45	0.31
F-stat	100.88	73.01	101.40	180.32	54.96

Column (1) presents the baseline IV estimates, column (2) presents the IV estimates using the crude oil inventory news and its interaction with the time dummy as the only instruments, column (3) presents the estimates using an alternative standardization of the news based on the cross-sectional dispersion of inventory forecasts, column (4) presents the results for returns during a 5-minute announcement window, and column (5) presents the estimates for returns during a 30-minute announcement window. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 11: Decomposition of weekly changes in the change in U.S. crude oil inventories in contributions from changes in U.S. crude oil production, changes in net crude oil imports, and changes in refiner net input of crude oil

	Pre-2008M9			*Post-2008M9*		
	Δ^2 Inv	Δ^2 Inv	Δ^2 Inv	Δ^2 Inv	Δ^2 Inv	Δ^2 Inv
Δ U.S. Production	0.88			2.60		
	(1.79)			(1.62)		
Δ Net Imports		5.17***			6.11***	
		(0.28)			(0.25)	
Δ Refiner Net Input			-1.76**			-2.83***
			(0.85)			(0.66)
Constant	-17.80	-6.00	-23.93	-22.95	51.01	5.79
	(255.39)	(150.17)	(252.59)	(219.24)	(133.51)	(213.28)
Observations	256	256	256	477	477	477
R-squared	0.00	0.65	0.02	0.01	0.62	0.04
F-stat	0.24	346.14	4.25	2.57	612.32	18.44

The dependent variable, Δ^2 *Inv*, stands for the weekly change in the change in U.S. commercial crude oil inventories, Δ *U.S. Production* stands for the change in U.S. crude oil production, Δ *Net Imports* for the change in U.S. net imports of crude oil, and Δ *Refiner Net Input* for the change in U.S. refiner net input of crude oil. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.

Table 12: Forecasting regressions for U.S. economic conditions

	1week	2week	3week	4week
Crude oil inventory news	-0.002	-0.007	-0.011	-0.014
	(0.007)	(0.015)	(0.021)	(0.025)
Crude oil inventory news × d_t	0.004	0.008	0.008	0.006
	(0.008)	(0.017)	(0.023)	(0.028)
d_t	0.023***	0.048***	0.073***	0.096***
	(0.008)	(0.012)	(0.021)	(0.027)
Constant	-0.014**	-0.030**	-0.045**	-0.060***
	(0.006)	(0.012)	(0.017)	(0.022)
Observations	732	731	730	729
R-squared	0.01	0.01	0.02	0.02
F-stat	3.07	3.46	3.96	4.49

The dependent variable is the 1- to 4-week-ahead forecast in the change of the Aruoba-Diebod-Scotti Business Conditions Index. d_t stands for a dummy variable that takes on the value of 1 for all observations after 2008M8. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The sample period is 2003M10 to 2017M10.