The Correlation of Oil and Equity Prices: The Role of Zero Lower Bound*

Deepa Datta, Benjamin K. Johannsen, Hannah Kwon, and Robert J. Vigfusson†

July 15, 2016

Abstract

Since 2008, oil and equity prices have moved together much more than they did previously. We show that this increased comovement is due in part to increased responsiveness of both oil prices and equity prices to macroeconomic news. Before 2008, there is little evidence that oil prices were responsive to macroeconomic news. Since 2008, reflecting the low interest rate environment and the relatively constrained monetary policy during this period (i.e. the zero lower bound), oil and equity prices have become more responsive to macroeconomic news. This finding suggests that different rules apply at the zero lower bound, implying the potential for large fiscal multipliers at the zero lower bound.

JEL Classifications: F31, F41, E30, E01, C81

1 Introduction

We document that the rules regarding oil and equity prices changed dramatically in late 2008. Oil and equity price movements became highly correlated, whereas earlier they were typically uncorrelated. Also in contrast to historical experience, oil and equity prices became responsive to macroeconomic news surprises, such as unanticipated changes in nonfarm payrolls. We provide both empirical evidence and theoretical support to show that this change in the rules results from nominal interest rates being constrained by the zero lower bound

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

†Comments and suggestions can be directed to robert.j.vigfusson@frb.gov.

(ZLB). In our theoretical work, we show that an environment in which the ZLB can cause such a change in the rules is also an environment in which a binding ZLB implies a larger government fiscal multipliers. As such, our empirical evidence is supportive of increased government spending in times when the ZLB is binding.

As can be seen in Figure 1, the correlation between oil and equity returns increased sharply in 2008. Between 1983 and 2008, the correlation fluctuated around zero, only turning sharply negative in response to events such as the 1990/1991 Gulf War. However, the correlation rose drastically in late 2008, reaching as high as 0.67 in 2010 and averaging around 0.5 until late 2013. Thereafter, the correlation has moved lower.

We also show that, in contrast to historical experience, oil and equity returns became responsive to macroeconomic news surprises, such as unanticipated changes in nonfarm payrolls. For example, Kilian and Vega (2011) report that oil prices do not have statistically significant responses to macroeconomic news surprises over the period from 1983 to 2008. Although, using data from 1957 to 2000, Boyd et al. (2005) claimed that equity prices responded positively to bad news in expansions and negatively to bad news in recessions, our results differ in that the increased responsiveness of equity prices post-2008 has outlasted the recession and instead seems to be related to the low level of interest rates.

We provide empirical evidence that this change in the rules results from the ZLB constraining nominal interest rates. Related to our work is Swanson and Williams (2014), who find that longer-term interest rates become less responsive to macroeconomic news surprises after 2008. They argue that the lack of a response is an effect of the zero lower bound on nominal interest rates. As such, one contribution of our work is showing that the ZLB affects not only interest rates by making them less responsive to surprises, but also other asset prices, including oil and equities, by making them more responsive. One important methodological contribution of our paper relative to Swanson and Williams is that, beyond to reporting results for time-varying responsiveness as was done by Swanson and Williams, we estimate and test directly the hypothesis that the responsiveness varies with monetary policy conditions, as measured by an interest rate implied by a modified Taylor Rule. Furthermore, we also test alternative hypotheses that attribute the change in responsiveness to the financialization of oil markets or increased uncertainty in the crisis era, and show that the evidence in favor the ZLB is stronger.

Our empirical evidence would be less compelling if we could not explain why the ZLB causes interest rates to have become *less* responsive to macroeconomic news surprises and oil and equity returns to have become *more* responsive. However, the explanation is intuitive and arises from the well-known result that, as interest rates become unresponsive to demand shocks, these shocks have larger effects on economic activity (Christiano et al., 2011). Increased responsiveness of oil and equity returns is a corollary of this result. To verify this claim, we provide a more formal analysis with a New-Keynesian model that is augmented to include oil. Using our New-Keynesian model, we show that, consistent with results reported in Bodenstein et al. (2013), the effects of oil supply shocks change when the ZLB binds because the monetary authority does not respond to changes in inflation at the ZLB. By contrast, away from the ZLB, changes in inflation lead to more than a one-for-one change in the nominal rate. Consequently, movements in inflation have different effects on the real interest rate at the ZLB and away from it. For example, at the ZLB, increases in inflation *decrease* real interest rates, and thus increase output and consumption. Furthermore, in our model we also show that oil and equity prices become more correlated at the ZLB in response to demand shocks.

Our empirical evidence provides important insights into the effects of being at the ZLB. There has been much debate using theoretical models about whether different rules apply at the zero lower bound. Our paper provides strong empirical evidence that the rules are different at the ZLB. As a way to test the effects of the ZLB, using the response of oil and equity returns to macroeconomic news surprises has several advantages over other methods. First, the identification is transparent. Our explanatory variables are based on surprises U.S. macroeconomic news announcements, which occur on fixed and pre-announced days. Second, these announcements are frequent. Between 2009 and 2015, we have over 800 observations for days on which announcements were made. Finally, our surprises are likely more informative than the shocks identified in related work. A full discussion of Wieland (2015) is in section 6.2, but we note that our shocks are actually likely more informative, as for example, a one standard deviation surprise in nonfarm payrolls has a similar oil price effect in one day as an oil supply shock in Wieland's work has in six months.

2 Data

Table 1 presents a summary of our main variables of interest. Our measure for the price of oil is the closing value, in dollars per barrel, of the front-month futures contract for West Texas Intermediate (WTI) crude oil for delivery in Cushing, Oklahoma obtained from NYMEX.[1] For equities, we primarily use the Fama-French constructed measure of the value-weighted

[1]The series reports the official daily closing prices at 2:30pm in the New York Mercantile Exchange. In contrast, Kilian and Vega (2011) use the daily spot price for WTI crude oil for delivery (freight on board) in Cushing, Oklahoma, as reported by the U.S. Energy Information Administration (EIA). Analyses using the EIA series, or the nearby futures price for Brent crude oil obtained from Bloomberg, generate similar results. Of these, we prefer the WTI nearby futures price as its more precise timing allows us to better relate it to the macroeconomic announcements. In supplementary analysis, we also use the WTI far futures price, which we define as the price of the furthest available December contract.

daily return on all NYSE, AMEX, and NASDAQ stocks. These oil and equity price series are in the first two panels of figure 1. To ensure that our findings are robust to excluding energy firms from the equity index, we also use returns on 12 industry-specific equity indexes from the Fama-French data library, as well as returns on the S&P 500 Ex-Energy index obtained from Bloomberg (Ticker: SPXXEGP). Summary statistics on our alternative oil and equity series can be found in Table 2.

We also conduct some analyses using metals prices, inflation compensation, and interest rates as the dependent variables. For metals returns, we use log-differences in the metals spot index constructed by the Commodities Research Bureau. Our measure of long-run inflation compensation is constructed as the spread between the yields for standard ten to five year-ahead Treasury bonds less the same spread between the yields on ten to five year Treasury inflation-protected securities (aka TIPS). We also use near-term inflation compensation, which is constructed using two to one year-ahead Treasure bonds and TIPS. As is discussed in Gürkaynak et al. (2010), inflation compensation "provides information about agents' inflation expectations, but its interpretation is complicated by inflation risk premia and the differential liquidity premia between TIPS and nominal securities. Finally, our primary measure of interest rates is the market yield on U.S. Treasury securities at a constant maturity of 2 years. We also include market yields at 1-year and 10-year constant maturity for comparison.

To generate our "daily" returns series, we first drop days that have missing values for any of our primary variables of interest: WTI futures and physical spot prices, equity prices, metals prices, interest rates, and inflation compensation when available. Next, we calculate returns and differences using these consecutive closing prices, thereby ensuring that the daily returns are calculated over the same period for each variable. We construct daily oil and metals returns as the log-difference (times 100) in the consecutive closing prices, and compare these returns to the daily percent return on equities ($Equity_t$). These returns are compared to the simple daily change in inflation compensation and interest rates. With the exception of inflation compensation, all of these variables are available for the entirety of our sample period, which covers April 6, 1983 through December 31, 2015. The data on inflation compensation begin on January 5, 1999. Additional summary statistics on these variables can be found in Table 1.

We measure macroeconomic news using the same approach that has been well-established in the empirical literature such as Beechey and Wright (2009) and Kilian and Vega (2011). It is important to note that news about macroeconomic announcements is not what macroeconomists would call a "news shock". A Beaudry/Portier-style news shock, as in Barsky et al. (2014) is information about the *future* state of the world. In contrast, our macroeconomic news announcements provide information about the *current* state of the world. We use survey results from Action Economics as the expected U.S. macroeconomic fundamentals. Macroeconomics news is defined as the difference between the announced realization of the macroeconomic aggregates and the survey expectations. We focus on the variables that Swanson and Williams (2014) use in their analysis of asset price movements during the zero lower bound period: capacity utilization, consumer confidence, core CPI, GDP (advance), initial claims, ISM manufacturing, leading indicators, new home sales, nonfarm payrolls, core PPI, retail sales excluding autos, and the unemployment rate. Although a number of these surprise series are available starting in 1980, most of our analysis extends from January 1992 through December 2015, which is the period over which all twelve of our macroeconomic news variables are available.

Since the units of measurement differ across the news indicators, we follow the common practice in this literature and normalize the surprise component of the each news announcement by the sample the standard deviation. This allows the units to be comparable across all announcements. Therefore, for each indicator j at time t the surprise component S_{jt} is

$$S_{it} = \frac{(A_{jt} - E_{jt})}{\sigma_j}$$

where A_{jt} denotes the released value of indicator j and E_{jt} refers to the market's expectation of indicator j prior to the announcement. To calculate σ_j, which is the standard deviation of the surprise component $(A_{jt} - E_{jt})$, we use the entire sample period available for each surprise. Following Beechey and Wright (2009), we flip the sign for unemployment and initial jobless claims announcements, so that positive surprises represent stronger-than-expected growth. Summary statistics for the surprise component of each announcement, $(A_{jt} - E_{jt})$, can be found in Panel C of Table 2.

In supplementary analysis, we also consider the variables' responses to oil-specific news related to the weekly data release on U.S. crude oil inventories. We measure the surprise as the difference between announced realizations of the weekly change in U.S. crude oil inventories and the ex-ante survey expectations. As with the macroeconomic news, we divide the surprise component of the inventories news by its full sample standard deviation. Survey expectations are obtained from Bloomberg for the period June 2003 to December 2015. To augment this relatively short sample analysis, we also approximate this oil news series using just the weekly changes in crude oil inventories beginning in November 1988, which is also standardized using the full sample standard deviation of the changes. Finally, our analysis of oil-specific news also conditions on federal funds rates surprises, as federal funds rate announcements have frequently been made on the same day as the weekly inventories data release. Our measure of funds rates surprises is the standardized difference between the announced federal funds rate and its futures market equivalent.

As discussed in Beechey and Wright (2009), the response of asset prices to news events occurs very rapidly, often completely adjusting within 15 minutes of the announcement. However, as was also noted in Beechey and Wright, although intradaily regressions provide more efficient estimates of the reactions to news announcements, the daily estimators also were consistent. It would seem reasonable to expect a similar result for oil prices. In addition, by using daily data, our results are most comparable to those reported in Kilian and Vega (2011).[2] Using high-frequency data, Rosa (2014) reports statistically significant results for the responses of oil prices to macroeconomic news over the 1999-2011 sample. However, he does not consider the role of time-variation, which we emphasize here, and which may explain the difference between the results reported in Rosa and those in Kilian and Vega.

3 Documenting the Increased Correlation Between Oil Prices and Equity Prices

3.1 Rolling Window Regressions

As shown in panel C of Figure 1, there has been a marked increase in the correlation between daily oil and equity returns. Using a regression to summarize comovement rather than to indicate causality, we depict the time variation in this relationship in Figure 1, which plots rolling window regression coefficients estimated over a rolling sample of one year for the model

$$Oil_t = \alpha + \beta Equity_t + \varepsilon_t. \tag{1}$$

As with the rolling correlation, we can see a sharp increase in the rolling coefficient estimates in September 2008. Using a standard Chow test, we find a statistically significant break date of September 27, 2008. Table 3 reports the slope coefficients on equity returns for the model in equation 1 estimated over three sample periods: the full sample, pre-break, and post-break. As shown in Table 3, the coefficient is slightly negative for the pre-break sample, but is large, positive, and significant for the post-break sample. The coefficient of 0.75 in the post-break sample implies that during this period, a daily return of 1 percent on the equity index is associated with an oil price increase of about 0.75 percent. We find similar results when using our alternative measures of oil prices, including the physical spot price for WTI and the nearby futures price for Brent crude oil. The results when using the

[2] Studies using higher frequency prices include Halova (2012), which looks at how oil and natural gas respond to news about oil and natural gas inventories.

WTI far futures series are qualitatively similar but quantitatively smaller, which is in line with the lower variation in futures prices, as reported in Table 2.

To test whether this changed relationship applies more broadly, we also perform the Chow test for the regression of metals returns and inflation compensation on equity returns. Figure 2 shows the time variation in the metals-equity rolling coefficient estimates. Applying the Chow test to the regression of metals on equity returns implies a statistically significant break date of October 4, 2008. As with oil, Table 3 shows that the slope coefficient on equity returns is essentially zero for the pre-break sample, but is much larger and statistically significant for the post-break sample. For both 5-year inflation compensation, we find a larger, significant coefficient for the post-break sample. A one-standard deviation increase in the daily equity return (1.1 percentage points) is associated with an increase in inflation compensation of 1.6 basis points, or almost one-third of the standard deviation for this variable.

Next, to ensure that the increased correlation between oil and equity prices is not being driven by fluctuations in the energy component of the S&P 500, we separately regress oil on each of the twelve Fama-French industry portfolios, determined by SIC codes. The results of the related Chow tests are presented in Panel B of Table 3. In the pre-break sample, returns in all of the non-energy related sectors are negatively associated with oil prices. Only the energy sector shows a positive, statistically significant relationship before the break in 2008. In contrast, post-break, all of the sectors display a positive and statistically significant relationship similar to that of the energy sector. These results confirm that our finding of an increased correlation between equity prices and oil prices is not being driven exclusively by equity prices for energy producers. Instead, the increased correlation between oil prices and equity prices is broad-based.

3.2 Covariance Decomposition

The increased correlation between oil and equity returns can also be represented by Figure 3, which shows the scatter plots of the standardized series for oil and equity returns generated by subtracting the sample mean and dividing by the sample standard deviation for each series. It is clear from the figure that the post-break period shows a stronger positive association.

To investigate this change further, we define the set of days P as the days on which the standardized returns have the same sign (and a positive product), and the set of days N, on which the standardized returns have opposite signs (and a negative product). In the figures, points in the upper-right and lower-left quadrants represent positive products days. Using the break date of September 27, 2008 identified previously, this decomposition allows us to measure the change in the *frequency* of these positive product days before and after the break

($\frac{P}{T}$), as well as the average *magnitude* of the products on those days ($\frac{1}{P}\sum_{t\in P}(Oil_t \times Equity_t)$):

$$
\begin{aligned}
Covariance &= \frac{1}{T}\sum_{t\in T}(Oil_t \times Equity_t) \\
Covariance &= \underbrace{\left(\frac{P}{T}\right)}_{\substack{\text{Positive}\\\text{Products}\\\text{Frequency}}} \underbrace{\left(\frac{1}{P}\sum_{t\in P}(Oil_t \times Equity_t)\right)}_{\substack{\text{Positive}\\\text{Products}\\\text{Avg. Magnitude}}} + \underbrace{\left(\frac{N}{T}\right)}_{\substack{\text{Negative}\\\text{Products}\\\text{Frequency}}} \underbrace{\left(\frac{1}{N}\sum_{t\in N}(Oil_t \times Equity_t)\right)}_{\substack{\text{Negative}\\\text{Products}\\\text{Avg. Magnitude}}}.
\end{aligned}
$$

We apply this decomposition to our pre-break and post-break samples in panel A of Table 4. We find that the increase in the covariance from -0.07 before the break to 0.42 after the break was driven by an increase in both the frequency of positive product days and the average magnitude of the product on those days. In the figure, the increase in the magnitude is represented by the larger average distance from the origin for the points in the positive product quadrants. Additionally, the average magnitude on negative product days fell in absolute value, which also contributes to the increase in the covariance, and is represented in the figure by the smaller distance from the origin for the points in the negative product quadrants.

A similar decomposition can be applied to the variances of the oil and equity returns series. Holding the classifications (and frequencies) constant from our covariance analysis, we examine the change in the average magnitude of the squared return on positive and negative product days. For oil, we can express this decomposition as:

$$
\begin{aligned}
Var(Oil_t) &= \frac{1}{T}\sum_{t\in T}(Oil_t)^2 \\
&= \underbrace{\left(\frac{P}{T}\right)}_{\substack{\text{Positive}\\\text{Product Days}\\\text{Frequency}}} \underbrace{\left(\frac{1}{P}\sum_{t\in P}(Oil_t)^2\right)}_{\substack{\text{Positive}\\\text{Product Days}\\\text{Avg. Magnitude}}} + \underbrace{\left(\frac{N}{T}\right)}_{\substack{\text{Negative}\\\text{Product Days}\\\text{Frequency}}} \underbrace{\left(\frac{1}{N}\sum_{t\in N}(Oil_t)^2\right)}_{\substack{\text{Negative}\\\text{Product Days}\\\text{Avg. Magnitude}}}.
\end{aligned}
$$

In panels B and C of Table 4 we can see that for positive product days, the average magnitude of the squared daily oil and equity returns both increased after the break. Additionally, the average magnitudes on negative product days both fell in absolute value after the break.

The covariance and variance decompositions show that days when oil and equity prices move together have become more important since 2008. One possible explanation for this change is that there may have been a shift from a supply-driven relationship to a demand-driven one. For example, exogenous disruptions to oil supply generally result in an increase in oil prices and a fall in the general health of the economy, and hence a negative relationship between oil prices and non-oil producer equity returns. In contrast, a demand-driven increase in oil prices is generally associated with an improving economy, resulting in a positive relationship between oil prices and equity prices. In recent years, the increase in the oil-equity correlation suggests that oil price movements have become more demand-driven. This greater role for demand shocks could be driven by an increase in the frequency or magnitude of demand shocks relative to supply shocks, or alternatively, an increase in the sensitivity of oil and equity prices to these demand shocks. In the next section, we test for this increased sensitivity by testing for a structural break and time variation in the responsiveness of oil and equity prices to macroeconomic news.

4 Estimating the Response to Macroeconomic News

4.1 Structural Break

Having found an increased correlation between oil and equity returns, we now test whether that correlation is associated with increased sensitivity of oil and equity prices to macroeconomic news. Our estimation procedure is similar to those found in earlier papers, such as Kilian and Vega (2011). For $Y_t \in \{Oil_t, Equity_t\}$, we estimate the effect of news using the model

$$Y_t = \alpha + \boldsymbol{\beta} \boldsymbol{S_t} + \varepsilon_t, \tag{2}$$

where $\boldsymbol{S_t}$ refers to the vector of standardized macroeconomic news surprises on day t. Each element β_j of $\boldsymbol{\beta}$ measures the response of oil or equity returns to a one standard deviation surprise for the corresponding announcement S_j. By estimating the daily returns around the time of the announcement, we attempt to isolate the immediate reaction of asset prices to the news announcement as much as possible. As discussed earlier, this strategy has already been applied successfully to numerous financial assets in the literature, including in Andersen et al. (2003) and Kilian and Vega (2011).

The regression estimates are based only on the sample 1992 to 2015, using data for only those days on which at least one news announcement was made.[3] By excluding days with no announcements, we likely minimize the variance of ε_t, resulting in the most precise estimates for the β_j.[4] Table 5 reports the β_j and t-statistics using robust standard errors for the

[3]Note that Chow test and covariance decomposition results in the previous sections still hold when limiting the sample to the period starting in 1992.

[4]The regression sample includes all days with at least one announcement and with available data for our dependent variables of interest. For each day in our regression sample, we set $S_{jt} = 0$ for those variables without an announcement on that day. In order to prevent these 0's from biasing the β_j, the S_{jt} are demeaned

regressions of oil and equity prices on our 12 surprise variables. When estimated over the full sample, the generally small coefficient estimates and lack of statistical significance indicate that both oil and equity prices are not responsive to macroeconomic news. Further analyses also in Table 5 show that metals, like oil and equities, are also generally not responsive to macroeconomic news. In contrast, we find more support for a significant response of long-run inflation compensation to a number of the news announcements.

Next, we explore whether the responsiveness to macroeconomic news changed after our estimated break date of September 2008 by interacting the surprise variable with a dummy for the post-crisis period. In order to prevent the analysis by being clouded by the observations around the most volatile days of the crisis, we divide this post-crisis sample, and incorporate two dummy variables in the model. The first is our dummy variable of interest, D_t^{post}, which represents the dummy for dates after November 1, 2008. The second, D_t^{crisis}, represents the dummy for the dates between September 1, 2008 and October 31, 2008.

$$Y_t = \alpha_0 + \boldsymbol{\beta_1 S_t} + \boldsymbol{\beta_2 S_t} D_t^{post} + \beta_3 D_t^{post} + \boldsymbol{\psi_1 S_t} D_t^{crisis} + \psi_2 D_t^{crisis} + \varepsilon_t \tag{3}$$

Tables 6 and 7 report results from the oil and equity regressions, respectively. For both oil and equities, we find that the β_{1j} are generally not statistically significant, pointing to little effect of macroeconomic news before the break date. To examine whether the effect changed after the break, we examine the β_{2j} coefficients. Additionally, we report for each surprise j the estimated effect of the surprise after the break, which is given by $\beta_{1j} + \beta_{2j}$. Finally, we report the F-statistic and p-value associated with the test of whether this post-break effect is significantly different from zero.

Overall, the responses of oil and equity prices to macroeconomic news have changed dramatically after 2008. The responses to almost all of the surprises reported in Tables 6 and 7 have changed sign. A few other surprises go from having little to no effect before the break to a positive effect after the break. For example, before the break date, the announcement of a higher-than-expected nonfarm payrolls number would have had little to no effect on oil prices and would have had a negative effect on equity prices, which likely reflects the prospect of central bank action to cool over-heating of the economy. In contrast, after the break, the announcement of higher-than-expected nonfarm payrolls has a positive effect on oil and equity prices, which likely reflects that the announcement was taken an indicator of stronger than expected demand during this period. One might claim that we are making too strong of a claim because the post-break impact response of equity prices is not statistically significant at conventional p-values and, more broadly, less than half of the

(using the regression sample of S_{jt}) before inclusion in the regression.

dummy variables for post-2008 are individually statistically significant. However, we note that, under the null hypothesis of no structural break, the result that of having almost all of the coefficients switching signs would itself be an incredibly rare event.

4.2 Time Variation

Andersen et al. (2007) report time variation in the responses of equity prices to macroeconomic news before 2008. Likewise, Swanson and Williams (2014) report that the response of interest rates to macroeconomic news varied before 2008. As such, in this section we consider whether we also find similar time variation in the responses of equity price returns and oil prices to the macroeconomic surprises. One way to measure the time variation in the sample is to observe time variation in our parameters of interest when estimating each of our news regressions over a rolling sample. Our alternative, adapted from Swanson and Williams, allows us to incorporate the information from all twelve surprises into one model. For each of our dependent variables of interest, we measure the time variation in the response to the surprises by estimating monthly rolling regressions of the form:

$$Y_t = \alpha^\tau + \delta^\tau \hat{\boldsymbol{\beta}} \boldsymbol{S_t} + \varepsilon_t^\tau, \tag{4}$$

where each regression is estimated over a 1-year rolling window. In contrast to Swanson and Williams, which estimates δ^τ and β^τ jointly as a nonlinear least squares problem, we define $\hat{\boldsymbol{\beta}}$ as a vector of fixed parameters estimated over the subsample 2009 to 2012 using the regression $Y_t = \alpha + \boldsymbol{\beta} \boldsymbol{S_t} + \varepsilon_t$. Table 8 reports these $\hat{\beta}_j$ for oil, equities, metals, and inflation compensation, and Figures 4 and 5 plot the estimates of δ^τ. For oil and equities, δ^τ are near zero in the early part of the sample, but move sharply positive beginning with the rolling sample that incorporates data after our identified break date of late 2008. Unsurprisingly, we find a similar pattern in the δ^τ for our alternative measures of the oil price, including the far futures and physical spot prices for WTI crude oil and the nearby futures price for Brent crude oil, as shown in the top panel of Figure 5. Furthermore, we show in the bottom panel of Figure 5 that the δ^τ for metals as well as 1-year and 5-year inflation compensation also follow a similar pattern.

The timing of the sharp increase is also consistent with the finding in Swanson and Williams (2014) that the interest rate becomes less sensitive to macroeconomic news around the Zero Lower Bound (ZLB) era. We replicate the Swanson and Williams result using the model in equation 4. As reported in Table 8, we estimate the $\hat{\beta}_j$ using the period 1992 to 2000, as the period 2009 to 2012 betas are (unsurprisingly) near zero and would not be suitable for the estimation of δ^τ in equation 4. The top panel of Figure 6 plots the estimates of our δ^τ for the 2-year Treasury yields along with analogous results for 1- and 10-year Treasury yields. As found by Swanson and Williams, we can see that interest rates became less responsive to macroeconomic news in the ZLB era. The results also support the finding that the shorter maturity yields were less responsive to macroeconomic news than the longer maturity yields, based on expectations for how long the ZLB period would last. Additionally, we find that both the 1- and 2-year yields become more responsive to news at the very end of the sample (in 2015), likely as market participants expect the end of the ZLB era. Finally, the bottom panel of Figure 6 plots the estimates of the rolling δ^τ for the 2-year yield along with oil and equities, and shows clearly that just as the interest rate sensitivity to news was declining, the sensitivities of oil and equities were increasing. In addition, we see the decline in oil and equity price sensitivity towards the end of the sample, just around the time that interest rate sensitivity increased.

4.3 Kernel Regression

The timing of the variation in the response to macroeconomic news announcements points to the ZLB as one likely driving factor. In this section, we test this hypothesis more directly using kernel regression. Rather than estimating time-varying coefficients, we now estimate coefficients that vary with an underlying, or controlling variable. As a first pass, we estimate how the relationship between oil and equities changes based on a ZLB metric, defined as the prediction for the federal funds rate using the modified Taylor rule as in Bernanke (April 28, 2015). This ZLB metric is intended to capture the target federal funds rate implied by the current state of the economy, without censoring due to the ZLB.

The left panel of Figure 7 depicts this implied interest rate along with the actual federal funds rate through the first quarter of 2015. To measure inflation, p, the modified rule uses the deflator for core personal consumption expenditures (PCE), which excludes food and energy prices. For the output gap, y, it uses estimates prepared by Federal Reserve staff for FOMC meetings through 2009, and then estimates produced and published by the Congressional Budget Office through 2015. It uses only real-time data, and places a weight of 1 on the output gap: $r = p + y + 0.5(p - 2) + 2$.

In general, for a given value of the implied rate, Z_k, we estimate the coefficients $\gamma(Z_k)$ using the model

$$\gamma(Z_k) = \arg\min_{\gamma} \sum_t \phi\left(\frac{Z_t - Z_k}{h}\right)(Y_t - X_t\gamma)^2, \qquad (5)$$

where ϕ is the pdf for a standard normal distribution, and h is a bandwidth chosen by the cross-validation method. This estimation places more weight on the observed (y_t, x_t) when

Z_t is close to Z_k.

Building on earlier results for the relationship between oil prices and equity returns, we first estimate $\mathbf{\Gamma}(Z_k) \equiv \{\alpha(Z_k), \beta(Z_k)\}$ by minimizing

$$\sum_t \phi\left(\frac{Z_t - Z_k}{h}\right)(Oil_t - \alpha - \beta Equity_t)^2 \tag{6}$$

with respect to (α, β). The right panel of Figure 7 plots our estimate of $\beta(Z_k)$, and provides further evidence that oil and equities have stronger comovement (i.e. $\beta(Z_k)$ is larger) when interest rates are low, and in particular, when the implied rate is negative.

To determine the statistical significance of these results, we construct an F-test of the unrestricted model in (6) against the restricted model, in which the coefficients α and β do not vary with Z_t. The restricted model is estimated using a standard regression of Oil_t on $Equity_t$, and is equivalent to a kernel regression in which $Z_t = 1$ for all t. We compare this F-statistic to a distribution of F^{sim} generated using a wild bootstrap procedure.[5] As reported in the first row of Table 9, our results reject the null hypothesis with a p-value of less than 0.01.

Second, we again test the hypothesis that this stronger comovement is coincident with higher sensitivity of oil and equity prices to macroeconomic news announcements. Analogously to equation 2, we estimate for oil, equities, and interest rates the coefficients $\mathbf{\Gamma}(Z_k) \equiv \{\alpha(Z_k), \boldsymbol{\beta}(Z_k)\}$ by minimizing

$$\sum_t \phi\left(\frac{Z_t - Z_k}{h}\right)(Y_t - \alpha - \boldsymbol{\beta}\boldsymbol{S_t})^2 \tag{7}$$

with respect to $\{\alpha, \boldsymbol{\beta}\}$. We measure the average sensitivity to macroeconomic news announcements over the range of our implied rate by taking an average of the $\hat{\beta}_j(Z_k)$, weighted by the frequency ω_j of announcement j in the estimation, $\overline{\boldsymbol{\beta}}(Z_k) = \frac{1}{12}\sum_{j=1}^{j=12}\omega_j\hat{\beta}_j\,(Z_k)$. Figure 8 underscores the higher sensitivity of oil and equities to macroeconomic news announcements during periods with lower implied rates, and the higher sensitivity of interest rates to macroeconomic news announcements during periods with higher implied rates.

Using the same wild bootstrap procedure as before, we conduct for each dependent variable an F-test of the unrestricted model in (7) against the restricted model, in which the

[5] To generate this distribution, we run 1000 simulations. For each simulation i, we use the restricted model estimates for $\hat{\alpha}(Z_k)$, $\hat{\beta}(Z_k)$, and $\hat{\varepsilon}_t$ to generate: $Y_{it}^{sim} = \hat{\alpha}(Z_t) - \hat{\beta}(Z_t)Equity_t + \nu_{it} * \hat{\varepsilon}_t$. Note that the Y_{it}^{sim} leave the $Equity_t$ and Z_t variables fixed, thereby preserving any existing serial correlation, and then scale up and down the residuals $\hat{\varepsilon}_t$ by $\nu_{it} \sim N(0,1)$, thereby preserving heteroscedasticity. Using these Y_{it}^{sim}, we estimate $\hat{\alpha}_i(Z_k)$ and $\hat{\beta}_i(Z_k)$ under the restricted and unrestricted models, and generate a distribution of F-statistics.

coefficients do not vary with Z_t. As reported in the first line of Panel B in Table 9, we are able to reject the null hypothesis for all three dependent variables.

We can also use kernel regression to test whether oil, equity prices, and interest rates show different sensitivities to good and bad news. For example, during the ZLB period, we might expect the magnitude of the negative response of oil and equity prices to negative macroeconomic surprises to be larger than the magnitude of the positive responses to positive surprises. However, we find no evidence of asymmetry (results not shown).

4.4 Alternative Hypotheses

The previous section provides strong evidence that oil and equities have stronger comovement and are more sensitive to macroeconomic news surprises when interest rates are low, and in particular, when the implied interest rate is negative. We now turn to testing alternative hypotheses for these findings. In particular, we test whether the observed relationships are associated with two other market conditions that were roughly coincident with the ZLB period: elevated market uncertainty and increasing financialization of the oil market.

Panel A of Figure 9 plots the 90-day rolling average of the VIX, which is a measure of options-implied stock market volatility. According to this measure, market uncertainty began rising in 2007, spiked sharply in 2008 at the height of the financial crisis, and remained elevated for a few years after that. According to the common folk wisdom that all correlations go to one in a crisis, the increased uncertainty could be an alternative driver of the elevated oil-equity correlation during the ZLB period.

A second alternative hypothesis is that with increased financialization of oil markets, the greater overlap between oil market and other financial market participants resulted in greater sensitivity of the financial markets for oil to general market conditions. According to this theory, the oil market would react much more strongly to events that earlier would have moved only equity markets. We capture this trend by measuring the 90-day rolling average of the open interest across all futures contracts for oil on NYMEX, as depicted in panel B of Figure 9.

To test these alternative hypotheses, we reestimate the kernel regression of oil on equities using the VIX and then open interest as the controlling variables. We also estimate for each of our three dependent variables the model in (7), using the alternative controlling variables. Lastly, we consider models using two controlling variables, and test the null hypothesis that a model including the implied interest rate along with the VIX or open interest is equivalent to a model including just one controlling variable. Table 9 summarizes the hypotheses being tested and the p-values that result from the wild bootstrap procedure for each test.

In panel A of Table 9, we can see that the inclusion of each controlling variable in turn is informative for the regression of oil on equities. As expected, Figure 10 shows that the beta of oil on equities increases with the VIX and with open interest. We also find in Table 9 that the inclusion of any two controlling variables for the regression of oil on equities is also significantly different from the restricted model that includes just one controlling variable. In summary, we find that all three of our controlling variables are informative for explaining the relationship between oil and equities.

For the regressions of our three dependent variables on the macroeconomic news surprises, however, we find that in general, only the implied rate is informative. The implied rate is informative against the alternative of a restricted model with no controlling variable, and the unrestricted models incorporating the implied rate along with either the VIX or open interest are also significantly different from the restricted models excluding the implied rate.

In contrast, we find that the VIX and open interest are generally not informative when compared to a model with no controlling variable, nor when compared to a model that incorporates the implied rate. Additionally, Figure 11 shows that the responsiveness to news varies quite a bit with both the VIX and open interest, in non-monotonic ways. In sum, we find that the implied rate is the most informative of our three potential controlling variables, providing most support to our theory that the zero lower bound is driving the results, instead of either of our alternative hypotheses. Given these results, we now turn to a model that provides a framework under which the oil-equity relationship varies with the prevailing interest rate and shows particular variation in the ZLB period.

5 A DSGE Model with Oil

In the previous section, we presented the dramatic differences in the behavior of oil and equity prices in the pre- and post-2008 samples. Using kernel regressions, we related these differences to the implied rate, and showed that the ZLB is a key difference between the two subsamples. The New-Keynesian model has become the benchmark for policy analysis at the ZLB after Eggertsson and Woodford (2003) and Eggertsson (2004). To study the effects of oil supply shocks and demand shocks on oil and equity prices at the ZLB, we use a small-scale, closed-economy, New-Keynesian model augmented with oil similar to the model in Bodenstein et al. (2013). As in the previous section, we find that oil and equity prices exhibit dramatically different behavior in normal times versus under the ZLB.

5.1 Households

The model economy is populated by a large number of identical households. The households value consumption of final goods, C_t, consumption of oil, C_t^O, labor supplied by the household, L_t, and real bond holdings, B_t/P_t, according to a utility function given by

$$E_t \sum_{m=0}^{\infty} \beta^m \left[u\left(C_{t+m}, C_{t+m}^O, L_{t+m}, \mu_{C^O,t+m}\right) + \eta_{t+m} \Xi_{t+m} v\left(\frac{B_{t+m}}{P_{t+m}}\right)\right].$$

Here, η_t is an exogenous and stochastic process that shifts household preferences for risk-free real bond holdings, $\mu_{C^O,t+1}$ is an exogenous and stochastic process that shifts the household preference for oil consumption, and the discount factor, β, satisfies $0 < \beta < 1$. We include end-of-period risk-free real bond holdings in the utility function as in Fisher (2015) in order to capture the spread between risky and risk-free assets. We couple the bonds in the utility function with the preference shifter, η_t, to allow the spread to change over time. We assume that the function v has the property that for the steady state level of net bond holdings, $v\left(\frac{B_t}{P_t}\right) = 0$. Moreover,we define $\Xi_{t+m} = u_1\left(C_{t+m}, C_{t+m}^O, L_{t+m}, \mu_{C^O,t+m}\right)$, which allows us to isolate the spread as identically η_t. The household faces a per-period budget constraint given by

$$\frac{B_t}{P_t} + \frac{P_t^O}{P_t} C_t^O + C_t + \frac{P_t^K}{P_t} K_{t+1} \leq \frac{W_t}{P_t} L_t + R_{t-1} \frac{B_{t-1}}{P_t} + \frac{R_t^K}{P_t} K_t + T_t + \frac{P_t^K}{P_t} K_t,$$

where P_t^O is the nominal price of oil, K_t are capital holdings of the household, P_t^K is the nominal price of capital, W_t is the nominal wage, R_t is the nominal interest rate, R_t^K is the nominal rental rate of capital, and T_t are real lump-sum profits, taxes, and transfers. We denote the price of final consumption goods by P_t.

The household's first-order optimality condition for risk-free bond holdings is given by

$$1 = \eta_t + \beta E_t \left[\frac{u_1\left(C_{t+1}, C_{t+1}^O, L_{t+1}, \frac{B_{t+1}}{P_{t+1}}, \eta_{t+1}\right)}{u_1\left(C_t, C_t^O, L_t, \frac{B_t}{P_t}, \eta_t\right)} \frac{R_t}{\pi_{t+1}} \right],$$

where $\pi_t \equiv \frac{P_t}{P_{t-1}}$. Changes in η_t represent changes in the spread between the risk-free one-period bond and other assets, and are meant to capture changing preferences for risk-free bond holdings, because of, for example, a flight to safety. Thus, η_t plays an analogous role to the spread shock in Smets and Wouters (2007). We normalize real bonds to be in zero net supply. In our model experiments in section 5.7, we use shocks to η_t to cause the ZLB to bind.

We specify the function u so that

$$u\left(C_t, C_t^O, L_t, \mu_{C^O,t}\right) = \log\left[\left((1-\omega_C)^{\frac{\rho_C}{1+\rho_C}} C_t^{\frac{1}{1+\rho_C}} + (\omega_C)^{\frac{\rho_C}{1+\rho_C}} \left(\frac{C_t^O}{\mu_{C^O,t}}\right)^{\frac{1}{1+\rho_C}}\right)^{1+\rho_C}\right] - \frac{\chi}{1+\phi} L_t^{1+\phi}.$$

These preferences are similar to those used in Bodenstein et al. (2013) in that households value consumption of oil in their utility flow and in that consumption of oil is an imperfect substitute for retail-goods consumption. Here, $\mu_{C^O,t}$ is a preference shifter that affects the utility flow of oil consumption. We interpret shocks to $\mu_{C^O,t}$ as oil demand shocks.

5.2 Retailers

Retailers are perfectly competitive and produce final output, Y_t, using intermediate inputs, $Y_t(j)$. Retailers aggregate these intermediate inputs using a production technology given by

$$Y_t = \left(\int_0^1 Y_t(j)^{\frac{\epsilon-1}{\epsilon}} dj\right)^{\frac{\epsilon}{\epsilon-1}},$$

where $\varepsilon > 1$. Profit maximization implies that retailers have a demand curve, given by

$$Y_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon} Y_t$$

for each intermediate input.

5.3 Intermediate Goods Producers

There is a unit-measure of intermediate goods producers who face demand curves from the retailers. They all solve the same profit maximization problem, in which they choose $P_t(j)$ to maximize

$$E_t \sum_{m=0}^{\infty} M_{t,t+m}\left[\frac{(1+\tau)P_{t+m}(j) - MC_{t+m}}{P_{t+m}}\left(\frac{P_{t+m}(j)}{P_{t+m}}\right)^{-\epsilon} Y_{t+m} - \frac{\Phi}{2}\left(\frac{P_{t+m}(j)}{P_{t+m-1}(j)} - 1\right)^2\right],$$

where $M_{t,t+m}$ measures the household's time t valuation of real profit flows in period $t+m$, MC_t is the real marginal cost which retailers take as given, and τ is a subsidy meant to offset steady state distortions due to monopoly power. The term $\frac{\phi}{2}\left(\frac{P_{t+m}(j)}{P_{t+m-1}(j)} - 1\right)^2$ represents the cost of price adjustment, which is introduced in a similar way to Rotemberg (1982). An alternative approach to introducing a nominal rigidity is to use the Calvo model of price adjustment, as in Christiano et al. (2005). We do not use the Calvo model because it complicates the solution methodology by introducing an additional state variable and dynamic equation.

Oil inputs, ${V^O}_t(j)$, are combined with capital and labor to produce intermediate goods according to the constant-returns-to-scale production technology

$$Y_t(j) = \left((1-\omega_Y)^{\frac{\rho_Y}{1+\rho_Y}} V_t(j)^{\frac{1}{1+\rho_Y}} + (\omega_Y)^{\frac{\rho_Y}{1+\rho_Y}} \left(V_t^O(j)\right)^{\frac{1}{1+\rho_Y}} \right)^{1+\rho_Y},$$

where

$$V_t(j) = K_t(j)^{\alpha} L_t(j)^{1-\alpha}.$$

The parameter ρ_Y controls the substitutability between oil and non-oil inputs. Importantly, our production function exhibits constant returns to scale, so that firms take marginal cost as independent of their pricing decision. The first-order optimality condition of intermediate goods producers is given by

$$(\pi_t - 1)\,\pi_t = \frac{\epsilon}{\Phi}\left(\frac{MC_t}{P_t} - 1\right) Y_t + E_t M_{t,t+1}\,(\pi_{t+1} - 1)\,\pi_{t+1}.$$

Because firms are identical, in equilibrium they all set $P_t\,(j) = P_t$ and all rent the same amount of capital, hire the same amount of labor, and use the same amount of oil.

5.4 The Oil Market

In each period oil supply, O_t, is exogenously determined and follows a stochastic process so that

$$\log\left(\frac{O_t}{O}\right) = \rho_O \log\left(\frac{O_{t-1}}{O}\right) + \sigma_O \epsilon_t^O,$$

where O is the steady state value of oil supply and ϵ_t^O is a standard normal random variable. Market clearing in the oil market requires that the quantity of oil demanded equal the exogenous oil supply:

$$O_t = V_t^O + C_t^O.$$

This is a closed economy version of the oil market in Bodenstein et al. (2013).

5.5 Government Policy and the Aggregate Resource Constraint

The fiscal authority purchases government consumption, G_t, in each period. We assume that G_t evolves exogenously so that

$$\log\left(\frac{G_t}{G}\right) = \rho_G \log\left(\frac{G_{t-1}}{G}\right) + \sigma_G \epsilon_t^G.$$

In each period, lump sum taxes are set to satisfy the government budget constraint, period-by-period, with $B_t = 0$. The monetary authority sets the nominal interest rate according to a Taylor rule, which is constrained by the ZLB,

$$R_t = \max\left\{1, \frac{1-\eta}{\beta} + \theta_\pi (\pi_t - 1)\right\}.$$

Note that in our specification of the Taylor rule, the monetary authority responds only to final-goods prices, not to oil prices. Finally, the aggregate resource constraint for final output is given by

$$Y_t = C_t + G_t + \frac{\Phi}{2}(\pi_t - 1)^2.$$

5.6 Parameter Values

We assume that the household time discount factor is $\beta = 0.985$ and that in steady state the preference parameter $\eta = 0.01$. The parameter η imply a steady-state 4 percent premium on risky assets as comparted to risk-free assets on an annualized basis, and the parameter β implies a steady-state risk-free real interest rate of 2 percent. We set $\epsilon = 7$, which is well within the range considered in Altig et al. (2011). For price adjustment costs, we set $\phi = 200$, which implies that a model with Calvo price adjustment friction in which firms update their price with probability of about 0.15 would have an identical linearization as our model. The parameter in the monetary policy rule is set so that the Taylor principle is satisfied, $\theta_\pi = 1.5$. We set $\rho_Y = \rho_C = -2$ to be in-line with estimates reported in Bodenstein et al. (2013), which implies an elasticity of substitution of $\frac{1}{2}$ for the different inputs to production and in the consumption bundle. We set $\omega_Y = 0.03$, $\omega_C = 0.02$, and we parameterize χ so that steady state labor supply is 1. The choice of $\omega_Y = 0.03$, $\omega_C = 0.02$ is similar to Bodenstein et al. (2013), and implies that oil is about 4.2 percent of output and also implies that firms use more oil as an input to production that households consume.

For each of the processes η_t, $\mu_{C,t}$, O_t, and G_t, we assume that the stochastic process governing its evolution is given by

$$\log(x_t/x) = \rho_x \log(x_{t-1}/x) + \sigma_x \epsilon_t^x$$

where x represents the steady state value and ϵ_t^x is a standard normal random variable that is independent of all other processes and over time. For out baseline calibration, we set $\rho_O = \rho_{\mu_{CO}} = 0.5$, and $\rho_\eta = \rho_G = 0.8$. Additionally, we set $\sigma_O = 0.03$, $\sigma_G = 0.01$, $\sigma_{\mu_{CO}} = 0.03$ and $\sigma_\eta = 0.05$. We assume that the capital stock is fixed ($K_t = \overline{K}$) and that the steady state capital-labor ratio is 15. We allow households to trade shares of the capital stock, but assume that there is no investment decision by households or firms. For the production function, $\alpha = 0.3$. Finally, we set $\phi = 1$.

5.7 Solving the Model

We solve the model using a policy iteration methodology similar to the solution strategy introduced by Bizer and Judd (1989) and Coleman (1991). The solution algorithm reads the equilibrium conditions as a mapping from equilibrium function at time $t+1$ to the equilibrium functions at time t. To solve for an equilibrium, we conjecture equilibrium functions for time $t+1$ and then calculate the equilibrium functions at time t using the equilibrium conditions. If the conjectured functions are equal to the calculated time t functions, then we have an equilibrium. If not, we use the calculated functions as new conjectured functions and proceed until convergence, at which point we have an equilibrium.

To accommodate the ZLB, we follow Gust et al. (2016) and parameterize two versions of the equilibrium functions: one that is operative with $R_t = 1$ and one that is operative with $R_t = \frac{1-\eta}{\beta} + \theta_\pi (\pi_t - 1)$. Expectations are calculated so that if $\frac{1-\eta}{\beta} + \theta_\pi (\pi_t - 1) < 1$, then the former functions are operative. Otherwise, the latter functions determine time t prices and quantities. That is, we use the endogenous interest rate to determine which function is used to calculate the values of the equilibrium variables in our model. In this way, we allow the equilibrium functions for every variable to have a kink at the ZLB.

The equilibrium functions are projected onto Smolyak polynomials as in Judd et al. (2014). We allow up to fourth-order terms in the projection. By solving the model nonlinearly, we avoid the pitfalls associated with linear approximations at the ZLB, which have been documented in Christiano and Eichenbaum (2012) and Braun et al. (2015). By employing a projection onto polynomials, we are implicitly assuming that the nonlinearities in the model (other than the ZLB) are smooth functions of the states. Gust et al. (2016) report that this assumption holds relatively well, even in their large-scale model.

5.8 Shocks at the ZLB

To make the ZLB bind, we consider the following experiment. In period 1, $\log(\eta_t/\eta)$ is set 6 standard deviations higher than its non-stochastic steady state. We interpret this as an increase in the spread between risky and risk-free assets. All other variables are set to their ergodic means. We then compute the impulse response to a shock in period 1. That is, we compute the average path for endogenous variables in each subsequent period and compare it to the average path for endogenous variables in each subsequent period conditional on an additional shock in period 1. We then compare these impulse response functions to impulse response functions computed assuming that $\log(\eta_t/\eta)$ takes its mean value in period 1.

In Figure 12 we show impulse response functions to a 5 percent decline in the supply of oil. The shock to oil supply increases the price of oil and causes core inflation to rise somewhat, which in normal times increases real interest rates. At the ZLB, the rise in inflation *reduces* the real interest rate because the monetary authority does not respond to the change in inflation. The decline in the real interest rate leads to an increase in consumption and an increase in equity prices.

In Figure 13 we show impulse response functions to a 5 percent increase in $\mu_{C^O,t}$. As was the case for a negative oil supply shock, in normal times, the shock to government consumption increases the price of oil and inflation. In normal times, the increase in inflation increases real interest rates, leading to a decrease in equity prices. At the ZLB, the real interest rate falls as inflation rises causing equity prices to rise. In Figure 14 we repeat the experiment for a 5 percent shock to government consumption. Qualitatively, the results are similar as in the case of the oil demand shock.

The key message from our DSGE model is that the response of equity and oil prices is different at the ZLB as compared to away from the ZLB. Furthermore, at the ZLB, equity and oil prices become more responsive to both types of shocks, and they move in the same direction at the ZLB.

Figures 12, 13, and 14 also show the effect of shocks under a "short duration" ZLB, under which the initial increase in η_t is smaller than in our baseline calibration. For each shocks, the response under the short duration ZLB is more similar to the response in normal times than to the response under the long duration ZLB. Our findings thus imply that the response of equity and oil prices under the ZLB should be most divergent from their responses in normal times when the ZLB is most binding.

6 Robustness

6.1 Oil-specific news announcements

We have shown that oil and equity returns respond in a time-varying way to macroeconomic news surprises, and that the time variation appears to be related to the monetary constraints for responding to news about global demand in the Zero Lower Bound (ZLB) era. To better differentiate between the ZLB and other theories such as financialization, we now test whether oil and equity prices have also had a time-varying response to oil-specific news. We measure oil-specific news by using surprises in U.S. crude inventories, which are announced weekly. Since news about crude inventories would likely not lead to monetary policy action at any time, we would not expect to see a change in the sensitivity of oil prices to inventory news before, during, or after the ZLB era. Rather, we would expect a positive inventory surprise to lower oil returns. Additionally, we would expect equity returns to show little sensitivity to oil inventories over the entire sample.

As mentioned in section 2, our measure of oil-specific news is the difference between announced realizations of the weekly change in U.S. crude oil inventories and the ex-ante survey expectations. Since the surveys are available only for the period June 2003 to December 2015, we run an additional analysis using just inventory changes, over the sample November 1988 to December 2015. Finally, because the dates of federal funds rate announcements generally coincide with inventories data releases, we control for funds rates surprises as well. The inventories and federal funds rate variables are each standardized by demeaning over the regression sample and dividing by the full sample standard deviation of the surprises. Lastly, in order to ensure that the results for equities are not driven by energy sector firms, we use the daily returns on the S&P 500 Ex-Energy Index for this analysis.

Figure 15 plots the rolling β for the regression of oil returns on oil inventory surprises and changes. The visual evidence that there is no structural break in this relationship around the start of the ZLB period is confirmed by the Chow test, which finds no significant break date for the regression of oil returns on the oil inventory surprise series.[6] We also find no significant break for the regression of equity returns on inventory surprises or changes.

As with the macroeconomic surprises, we also test whether the responsiveness to oil-specific news changed after our estimated break date of September 2008 using the model

$$\begin{aligned} Oil_t &= \beta_0 + \beta_1 S_t + \beta_2 D_t^{crisis} + \beta_3 S_t D_t^{crisis} + \beta_4 D_t^{post} + \beta_5 S_t D_t^{post} + \beta_6 S_{FF_t} + \varepsilon_t \\ Equity_t &= \beta_0 + \beta_1 S_t + \beta_2 D_t^{crisis} + \beta_3 S_t D_t^{crisis} + \beta_4 D_t^{post} + \beta_5 S_t D_t^{post} + \beta_6 S_{FF_t} + \varepsilon_t. \end{aligned}$$

[6]We do find a significant break date for the regression of oil returns on oil inventory changes, but this break is on December 16, 2000, and therefore is not likely related to the ZLB period.

The estimated results in table 10 support the finding that there is no change in the responsiveness of oil or equities to oil-specific news around our break date. The pre-break response of oil to inventory surprises of -0.45 is statistically significant, but the coefficient on the interaction of the post-break dummy and the surprise variable (β_5) is not significant. Additionally, the post-break response of oil prices to inventory surprises of -0.46 is similar in magnitude to the pre-break response, and is also statistically significant according to the associated F-test. In contrast, we find that equities have no statistically significant response to the inventory surprises before or after the break. The final two columns of the table show that the results are broadly similar when using inventory changes and a longer sample.

6.2 Reconciling our work with Wieland (2015)

In a related paper, Wieland (2015) also studies whether different rules apply at the zero lower bound. Similar to what we report in section 5, he shows that, in a standard New Keynesian model, negative supply shocks can theoretically be expansionary at the zero lower bound. However, in contrast to our evidence of responses changing with the ZLB, Wieland estimates that global oil supply shocks have similar effects in the ZLB era as before. Using this evidence as motivation, he develops a new model under which oil supply shocks are not expansionary. Finally, he notes that this new model does not predict large fiscal multipliers at the ZLB, and exhorts caution among policy-makers who might expect large positive outcomes from fiscal stimulus at the ZLB.

We would argue that our empirical evidence, which supports the model predictions for expansionary effects of both negative oil supply shocks and fiscal stimulus at the ZLB, is more salient for identifying the effects of the ZLB for two main reasons. First, we note that the small price effects found in Wieland's empirical exercise on oil shocks provide only weak motivation to develop the new model. Second, Wieland uses evidence on only supply shocks in his paper to motivate the new model, but emphasizes the new model's implications for fiscal multipliers. In contrast, we provide evidence for demand shocks using our macroeconomic surprises, and find large differences when comparing the effects before and during the ZLB era.

We first note that Wielands estimated price effects are too small. He uses the Kilian (2009) identifying assumption that an oil supply shock is measured as an unexpected change in global oil production. However, the resulting supply shocks have almost no effect on oil prices. In particular, "a one-standard-deviation negative oil supply shock raises real oil prices by just over 1 percent after 6 months" (Wieland (2015), page 22). A 1 percent change in oil prices over 6 months would seem to have little signal to economic agents, given that the standard deviation of daily oil price changes is over 2 percent (our table 2). Under Wieland's standard New Keynesian model, the contractionary effects of an oil supply shock would come through prices and inflation expectations. Given Wieland's empirically small price effects, it is perhaps not surprising that he does not find any contractionary effects.

Second, one of the main implications of Wieland's new model is its lack of large fiscal multipliers. While the development of that model is motivated by evidence on supply shocks, we find more direct evidence for the question of fiscal multipliers in our empirical work on macroeconomic surprises. In contrast to Wieland's new model prediction, as reported in our tables 6 and 7, we find that during the ZLB era, a one standard deviation surprise in nonfarm payrolls is associated with an almost 1 percent increase in oil prices, and a 0.2 percent increase in equity prices. These effects are much larger than those found by Wieland, as a one-standard deviation surprise in nonfarm payrolls has a similar oil price effect in one day as does the 6 month effect of an oil supply shock in Wieland's work. Our direct evidence on the effect of demand shocks would call into question both the empirical motivation for Wieland's development of a new model and the predictions of that model for the effects of fiscal stimulus at the ZLB.

7 Conclusion

Before 2008, idiosyncratic factors were more likely the primary drivers of oil prices and equity prices, and neither were responsive to macroeconomic news announcements. After 2008, the surprises have had a major impact on both equity prices and oil prices. These results suggest that oil price movements have been driven by macroeconomic fundamentals rather than financial factors. Consequently, the higher responsiveness of both oil and equity prices is translated into the high correlation between the two.

Further evidence against the financialization hypothesis comes from our finding that oil prices have had a relatively constant responsiveness to oil inventory surprises and that equity prices did not response to this series.

Our result that oil prices and equities became more responsive to news in the ZLB era may have important implications for macroeconomic modeling. The results are consistent with Christiano et al. (2011), which argues that fiscal policy became more effective during the ZLB period when monetary policy is constrained, using an New Keynesian model we find that oil and equity prices have different behavior at the ZLB than in normal times. Swanson and Williams (2014) argue that their finding that the 2-year rate only became unresponsive to news in 2011 means that the ZLB was not yet binding until 2011 and as such fiscal policy would not have been additionally potent in 2009-10. However, because oil prices, equity returns and metals prices are all responsive in 2009, perhaps monetary policy was constrained and fiscal policy was more effective.

References

Ron Alquist and Olivier Gervais. The role of financial speculation in driving the price of crue oil. Technical report, Bank of Canada Discussion Paper, 2011.

David Altig, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde. Firm-specific capital, nominal rigidities and the business cycle. *Review of Economic Dynamics*, 14(2): 225–247, April 2011.

Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Clara Vega. Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, pages 38–62, 2003.

Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Clara Vega. Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics*, 73(2):251–277, 2007.

Robert B Barsky, Susanto Basu, and Keyoung Lee. Whither news shocks? *NBER Working Paper 20666*, 2014.

Paul Beaudry and Franck Portier. Stock prices, news, and economic fluctuations. *American Economic Review*, 96(4):1293–1307, 2006.

Meredith J Beechey and Jonathan H Wright. The high-frequency impact of news on long-term yields and forward rates: Is it real? *Journal of Monetary Economics*, 56(4):535–544, 2009.

Ben Bernanke. The taylor rule: A benchmark for monetary policy? *Ben Bernanke's Blog*, April 28, 2015.

David S Bizer and Kenneth L Judd. Taxation and uncertainty. *The American Economic Review*, 79(2):331–336, 1989.

Martin Bodenstein, Luca Guerrieri, and Christopher J Gust. Oil shocks and the zero bound on nominal interest rates. *Journal of International Money and Finance*, 32:941–967, 2013.

John H. Boyd, Jian Hu, and Ravi Jagannathan. The stock market's reaction to unemployment news: why bad news is usually good for stocks. *Journal of Finance*, 60(2):649–672, 2005.

R Anton Braun, Lena Boneva, and Yuichiro Waki. Some unpleasant properties of loglinearized solutions when the nominal rate is zero. *Bank of England Staff Working Paper No. 553*, 2015.

Bahattin Buyuksahin, Michael S Haigh, Jeffrey H Harris, James A Overdahl, and Michel A Robe. Fundamentals, trader activity and derivative pricing. In *EFA 2009 Bergen Meetings Paper*, 2008.

Lawrence Christiano, Martin Eichenbaum, and Sergio Rebelo. When is the government spending multiplier large? *Journal of Political Economy*, 119:78–121, 2011.

Lawrence J Christiano and Martin Eichenbaum. Notes on linear approximations, equilibrium multiplicity and e-learnability in the analysis of the zero lower bound. *Manuscript, Northwestern University*, 2012.

Lawrence J Christiano, Martin Eichenbaum, and Charles L Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy*, 113(1): 1–45, 2005.

Wilbur John Coleman. Equilibrium in a production economy with an income tax. *Econometrica*, pages 1091–1104, 1991.

Gauti Eggertsson. *Monetary and Fiscal Coordination in a Liquidity Trap*. PhD thesis, Princeton University, 2004. Chapter 3.

Gauti B Eggertsson and Michael Woodford. Zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 2003(1):139–233, 2003.

Erkko Etula. Broker-dealer risk appetite and commodity returns. *Journal of Financial Econometrics*, 11(3):486–521, 2013.

Bassam Fattouh, Lutz Kilian, and Lavan Mahadeva. The role of speculation in oil markets: what have we learned so far? *CEPR Discussion Paper No. DP8916*, 2012.

Jonas DM Fisher. On the structural interpretation of the smets–wouters risk premium shock. *Journal of Money, Credit and Banking*, 47(2-3):511–516, 2015.

Reuven Glick and Sylvain Leduc. Central bank announcements of asset purchases and the impact on global financial and commodity markets. *Journal of International Money and Finance*, 31(8):2078–2101, 2012.

Francois Gourio and Phuong Ngo. Risk premia at the zlb: a macroeconomic interpretation. Technical report, Working paper, 2016.

Joseph W Gruber and Robert Vigfusson. Interest rates and the volatility and correlation of commodity prices. *FRB International Finance Discussion Paper No. 1065*, 2012.

Refet S Gürkaynak, Brian Sack, and Jonathan H Wright. The tips yield curve and inflation compensation. *American Economic Journal: Macroeconomics*, pages 70–92, 2010.

Christopher Gust, David López-Salido, Edward Herbst, and Matthew E Smith. The empirical implications of the interest-rate lower bound. *FEDS Working Paper*, 2016.

Marketa W Halova. Gas does affect oil: Evidence from intraday prices and inventory announcements. Technical report, Working paper, Washington State University, 2012.

James Douglas Hamilton. *Time series analysis*, volume 2. Princeton university press Princeton, 1994.

Benjamin K. Johannsen and Elmar Mertens. The expected real interest rate in the long run: Time series evidence with the effective lower bound. *FEDS Notes*, 2016.

Kenneth L Judd, Lilia Maliar, Serguei Maliar, and Rafael Valero. Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. *Journal of Economic Dynamics and Control*, 44:92–123, 2014.

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *The American Economic Review*, 99(3):1053–1069, 2009.

Lutz Kilian and Clara Vega. Do energy prices respond to us macroeconomic news? a test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics*, 93 (2):660–671, 2011.

Carlo Rosa. The high-frequency response of energy prices to us monetary policy: Understanding the empirical evidence. *Energy Economics*, 45:295–303, 2014.

Julio J Rotemberg. Sticky prices in the united states. *The Journal of Political Economy*, pages 1187–1211, 1982.

Frank Smets and Rafael Wouters. Shocks and frictions in us business cycles: A bayesian dsge approach. *The American Economic Review*, pages 586–606, 2007.

James H. Stock and Mark W. Watson. *Introduction to Econometrics*. Pearson/Addison Wesley Boston, 2003.

Eric T Swanson and John C Williams. Measuring the effect of the zero lower bound on medium-and longer-term interest rates. *The American Economic Review*, 104(10):3154–3185, 2014.

Ke Tang and Wei Xiong. Index investment and the financialization of commodities. *Financial Analysts Journal*, 68(5):54–74, 2012.

Johannes Wieland. Are negative supply shocks expansionary at the zero lower bound? inflation expectations and financial frictions in sticky-price models. *manuscript UC Berkeley*, 2015.

A Tables and Figures

Table 1: Summary Statistics

Variable	Obs.	Start Date	Mean	Std. Dev.	Min.	Max.
Oil returns (WTI nearby futures)	8084	1983-Apr-06	0.00	2.40	-40.05	22.80
Equity returns	8084	1983-Apr-06	0.05	1.11	-17.41	10.40
Metals returns	8084	1983-Apr-06	0.01	0.86	-10.29	9.40
Δ Inflation compensation (1 year)	4196	1999-Jan-05	0.02	13.35	-243.90	159.53
Δ Inflation compensation (5 year)	4196	1999-Jan-05	0.00	4.71	-53.45	30.88
Δ Interest rate (1 year)	8084	1983-Apr-06	-0.11	5.59	-83.00	52.00
Δ Interest rate (2 year)	8084	1983-Apr-06	-0.11	6.29	-84.00	38.00
Δ Interest rate (10 year)	8084	1983-Apr-06	-0.10	6.50	-75.00	39.00
Implied federal funds rate	6045	1992-Jan-29	2.06	3.00	-4.60	6.26
VIX (rolling 90-day avg.)	6242	1992-Jan-29	19.76	7.23	10.84	55.03
Open interest (rolling 90-day avg.)	6242	1992-Jan-29	0.87	0.49	0.30	1.87

Notes: Oil and metals returns are calculated as log-differenced prices, times 100. Equity returns are expressed in percentage points. Inflation compensation and interest rates are expressed in basis points.

Table 2: Additional Summary Statistics

Variable	Obs.	Start Date	Mean	Std. Dev.	Min.	Max.
Panel A: Oil Returns						
WTI nearby futures returns	8084	1983-Apr-06	0.00	2.40	-40.05	22.80
WTI far futures returns	8084	1983-Apr-06	0.01	1.38	-10.35	10.80
WTI physical spot returns	7410	1986-Jan-03	0.01	2.56	-40.64	21.70
Brent nearby futures returns	7888	1983-May-17	0.00	2.33	-40.71	27.82
Panel B: Equity Sector Returns						
Equity returns (full index)	8084	1983-Apr-06	0.05	1.11	-17.41	10.40
Equity returns ex. energy	4446	1998-Jan-02	0.02	1.27	-8.70	10.62
Consumer nondurables	8084	1983-Apr-06	0.06	0.96	-17.03	9.23
Consumer durables	8084	1983-Apr-06	0.04	1.48	-18.35	9.55
Manufacturing	8084	1983-Apr-06	0.05	1.23	-20.24	10.02
Energy	8084	1983-Apr-06	0.05	1.46	-19.43	18.82
Chemicals	8084	1983-Apr-06	0.05	1.11	-19.20	9.86
Business equipment	8084	1983-Apr-06	0.05	1.57	-20.09	15.36
Telecommunications	8084	1983-Apr-06	0.05	1.25	-16.69	14.09
Utilities	8084	1983-Apr-06	0.05	0.98	-12.86	13.50
Shops	8084	1983-Apr-06	0.05	1.17	-16.74	10.99
Healthcare	8084	1983-Apr-06	0.06	1.17	-17.89	10.83
Finance	8084	1983-Apr-06	0.05	1.44	-14.84	16.89
Other	8084	1983-Apr-06	0.04	1.20	-16.58	9.93
Panel C: Macroeconomic News Surprises						
Capacity utilization	333	1988-Apr-18	-0.01	0.35	-1.57	1.40
Consumer confidence	292	1991-Jul-30	0.13	5.17	-14.00	13.30
Core CPI	317	1989-Aug-18	0.00	0.11	-0.34	0.40
GDP (advance)	95	1992-Jan-29	0.10	0.76	-1.68	1.80
Initial claims	1216	1991-Jul-18	-0.06	18.42	-85.00	94.00
ISM manufacturing	309	1990-Feb-01	-0.01	2.01	-6.30	7.40
Leading indicators	431	1980-Feb-29	0.01	0.31	-1.80	2.00
New home sales	330	1988-Mar-29	4.60	57.42	-166.00	249.00
Nonfarm payrolls	371	1985-Feb-01	-8.62	102.39	-328.00	408.50
Core PPI	316	1989-Aug-11	-0.02	0.25	-1.20	1.07
Retail sales ex. autos	430	1980-Feb-13	-0.03	0.68	-2.40	5.13
Unemployment rate	429	1980-Feb-07	0.04	0.16	-0.60	0.60
Federal funds rate	253	1988-Nov-02	-0.02	0.08	-0.47	0.16
Oil inventories (surprises)	651	2003-Jun-18	20.77	3209.36	-10034.29	9078.11
Oil inventories (changes)	1690	1983-Jan-19	103.19	4289.57	-15222.00	12490.00

Notes: Oil returns are calculated as log-differenced prices, times 100. Equity returns are expressed in percentage points. News surprises are defined as the difference between the announced realization of the macroeconomic aggregates and the survey expectations, with the exception of Oil inventories (changes), which reflect the weekly change in the level of crude oil inventories.

Table 3: Chow Test Results

			Full Sample		Pre-Break		Post-Break	
	Obs.	Break Date	β	t-stat	β	t-stat	β	t-stat
Panel A: Regressions on the Full Index of Equity Returns								
WTI nearby futures returns	8084	2008-Sep-27	0.17	(6.93)	-0.16	(-5.26)	0.75	(19.62)
WTI far futures returns	8084	2008-Aug-30	0.09	(6.56)	-0.06	(-3.39)	0.35	(18.38)
WTI physical spot returns	7410	2008-Sep-25	0.18	(7.05)	-0.15	(-4.64)	0.76	(20.04)
Brent nearby futures returns	7888	2008-Sep-27	0.23	(9.75)	-0.11	(-3.69)	0.82	(25.54)
Metals returns	8084	2008-Oct-04	0.10	(11.43)	0.02	(2.22)	0.24	(13.92)
Δ Inflation compensation (1 year)	4196	2001-Aug-01	1.09	(6.75)	-0.27	(-0.64)	1.38	(7.94)
Δ Inflation compensation (5 year)	4196	2008-Nov-12	0.77	(13.84)	0.33	(4.52)	1.42	(16.64)
Panel B: Regressions of Oil on Various Equity Sectors								
Consumer nondurables	8084	2008-Sep-27	0.01	(0.22)	-0.30	(-9.32)	0.87	(16.33)
Consumer durables	8084	2008-Sep-27	0.09	(5.18)	-0.17	(-7.31)	0.49	(17.76)
Manufacturing	8084	2008-Sep-27	0.19	(8.94)	-0.14	(-4.88)	0.67	(21.06)
Energy	8084	2008-Nov-22	0.58	(34.09)	0.47	(22.88)	0.90	(31.67)
Chemicals	8084	2008-Sep-27	0.10	(4.03)	-0.21	(-7.41)	0.79	(19.28)
Business equipment	8084	2008-Sep-27	0.07	(3.85)	-0.07	(-3.84)	0.65	(17.29)
Telecommunications	8084	2008-Sep-27	0.08	(3.64)	-0.16	(-6.50)	0.66	(17.29)
Utilities	8084	2008-Jul-05	0.25	(9.20)	-0.05	(-1.44)	0.71	(16.49)
Shops	8084	2008-Sep-27	-0.05	(-2.15)	-0.27	(-10.54)	0.65	(14.13)
Health care	8084	2008-Sep-27	-0.01	(-0.53)	-0.21	(-8.07)	0.63	(13.41)
Finance	8084	2008-Sep-27	0.08	(4.27)	-0.21	(-8.53)	0.40	(15.32)
Other	8084	2008-Sep-27	0.13	(5.73)	-0.17	(-6.32)	0.65	(18.51)

Notes: Panel A reports equity betas from the regression $Y_t = \alpha + \beta Equity_t + \varepsilon_t$. Panel B reports equity sector betas from the regression $Oil_t = \alpha + \beta EquitySector_t + \varepsilon_t$. Full sample observations, β, and t-statistics are reported. The break date, pre-break, and post-break results are estimated after applying the standard Chow test to determine the break date which minimizes the sum of squared errors for regressions run on the pre- and post-break samples. All of these break dates were found to be statistically significant at the 10% level ($F_{crit} = 5.0$) when using the standard Andrews supremum-Wald critical value based upon 15% trimming of the sample as in Stock and Watson (2003).

Table 4: Oil and Equity Covariance Decomposition

	Covariance / Variance	Obs.	Positive Product Days Frequency	Positive Product Days Magnitude	Negative Product Days Frequency	Negative Product Days Magnitude
Panel A: Oil-Equity Covariance Decomposition						
Full Sample	0.08	8089	0.52	0.61	0.48	-0.51
Pre-break	-0.07	6285	0.48	0.47	0.52	-0.57
Post-break	0.42	1804	0.65	0.84	0.35	-0.35
Panel B: Oil Variance Decomposition						
Full Sample	5.77	8089	0.52	0.90	0.48	1.10
Pre-break	5.63	6285	0.48	0.81	0.52	1.17
Post-break	6.26	1804	0.65	1.10	0.35	0.82
Panel C: Equity Variance Decomposition						
Full Sample	1.23	8089	0.52	1.20	0.48	0.79
Pre-break	1.02	6285	0.48	1.09	0.52	0.92
Post-break	1.96	1804	0.65	1.23	0.35	0.57

Notes: We report the covariance and variance decompositions for the demeaned and standardized series for oil and equity returns. Positive product days are defined as the days on which these standardized oil and equity series move in the same direction; negative product days are defined as the days on which the two series move in opposite directions. The last date in the pre-break sample is September 27, 2008. See section 3.2 for further details.

Table 5: Oil, Equity, Metals, and Interest Rates Surprise Beta Estimates (full sample)

	Oil		Equity		Metal		Inflation comp.		Interest rate	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Capacity utilization	0.30	2.27	0.04	0.61	0.06	0.99	-0.16	-0.45	1.66	4.72
Consumer confidence	0.17	1.25	-0.07	-1.04	-0.09	-1.66	0.67	1.88	1.21	3.43
Core CPI	0.06	0.40	-0.19	-2.36	-0.08	-1.19	0.92	2.18	1.52	3.77
GDP (advance)	-0.26	-1.14	0.00	0.00	0.08	0.82	1.46	2.37	1.23	2.01
Initial claims	0.10	1.47	0.05	1.32	0.02	0.77	0.77	4.44	1.10	6.32
ISM manufacturing	0.12	0.93	0.17	2.46	0.06	1.11	0.93	2.59	2.96	8.42
Leading indicators	0.21	1.02	-0.16	-1.54	0.17	1.98	-0.74	-1.52	0.01	0.01
New home sales	0.05	0.35	-0.01	-0.19	0.03	0.47	0.71	2.06	0.66	1.92
Nonfarm payrolls	0.13	0.86	-0.04	-0.55	0.15	2.42	0.68	1.52	5.12	13.18
Core PPI	0.09	0.69	-0.07	-1.07	0.00	0.01	0.79	2.23	0.43	1.22
Retail sales ex. autos	-0.27	-1.66	0.16	1.86	0.04	0.59	1.50	3.82	2.65	6.27
Unemployment rate	-0.19	-1.22	0.02	0.21	0.08	1.26	-0.31	-0.74	1.36	3.36
Observations	2948		2948		2948		2088		2948	
R-squared	0.53		0.48		0.46		0.42		0.48	

Notes: We report the β_j from the regression $Y_t = \alpha + \boldsymbol{\beta S_t} + \varepsilon_t$, where S_{jt} refers to the standardized and demeaned news for announcement j on day t. The parameters β_j measure the responses of each dependent variable to a one standard deviation news surprise for announcement j. See section 4.1 for more detail.

Table 6: Oil Pre- and Post- Break Regression Results

	Cap. util.	Cons. conf.	Core CPI	GDP (adv.)	Init. claims	ISM manuf.	Leading ind.	New homes	Nonf. payrolls	Core PPI	Retail sales	Unemp. rate
Surprise	0.29	0.22	0.20	-0.19	-0.02	-0.05	0.57	0.06	0.00	0.00	-0.39	-0.27
	(1.61)	(1.29)	(1.11)	(-0.72)	(-0.28)	(-0.31)	(1.71)	(0.44)	(-0.01)	(-0.01)	(-2.17)	(-1.39)
Surprise	-0.37	-0.17	-0.48	-0.28	0.39	0.55	-0.42	-0.11	0.92	0.66	0.41	0.26
*Post-break dummy	(-1.27)	(-0.61)	(-1.42)	(-0.53)	(2.68)	(1.82)	(-0.98)	(-0.28)	(2.29)	(1.92)	(1.08)	(0.80)
Post-break response												
$\beta_{1j} + \beta_{2j}$	-0.08	0.05	-0.28	-0.47	0.37	0.50	0.15	-0.05	0.92	0.66	0.02	0.00
F-stat	0.12	0.04	0.97	1.08	9.08	3.80	0.32	0.02	6.21	4.50	0.00	0.00
p-value	0.72	0.84	0.32	0.30	0.00	0.05	0.57	0.90	0.01	0.03	0.96	0.99
Observations	280	283	283	95	1176	277	284	281	275	281	282	273

Notes: We report results from the regression $Oil_t = \alpha_0 + \boldsymbol{\beta_1 S_t} + \boldsymbol{\beta_2 S_t} D_t^{post} + \beta_3 D_t^{post} + \boldsymbol{\psi_1 S_t} D_t^{crisis} + \psi_2 D_t^{crisis} + \varepsilon_t$. The regression sample includes 2948 observations. The first two rows report the β_{1j} coefficients on the surprise variables and the associated t-statistics. These surprise variables, S_{jt}, are defined as the standardized news for announcement j on day t. The next two rows report the coefficients on the interaction of the post-break dummy and the surprise variable (β_{2j}). Additionally, we report the sum of the coefficients on the surprise and the surprise interacted with the post-break dummy ($\beta_{1j} + \beta_{2j}$), which represents the estimated effect of the surprise after the break. Finally, we also report the F-statistic and p-value associated with the test of whether this post-break response is significantly different from zero. See section 4.1 for more detail.

Table 7: Equity Pre- and Post- Break Regression Results

	Cap. util.	Cons. conf.	Core CPI	GDP (adv.)	Init. claims	ISM manuf.	Leading ind.	New homes	Nonf. payrolls	Core PPI	Retail sales	Unemp. rate
Surprise	0.12	0.06	-0.29	0.08	0.00	0.06	-0.03	-0.02	-0.10	-0.10	0.00	0.05
	(1.33)	(0.68)	(-3.15)	(0.62)	(0.08)	(0.77)	(-0.20)	(-0.22)	(-1.21)	(-1.31)	(0.03)	(0.54)
Surprise	-0.12	-0.12	0.39	-0.34	0.12	0.38	0.01	0.05	0.31	0.30	0.55	-0.06
*Post-break dummy	(-0.84)	(-0.80)	(2.25)	(-1.28)	(1.61)	(2.49)	(0.02)	(0.28)	(1.53)	(1.71)	(2.82)	(-0.37)
Post-break response												
$\beta_{1j} + \beta_{2j}$	0.00	-0.06	0.10	-0.26	0.12	0.44	-0.03	0.04	0.21	0.20	0.55	-0.01
F-stat	0.00	0.24	0.48	1.25	3.90	11.70	0.05	0.04	1.29	1.61	10.39	0.00
p-value	0.98	0.62	0.49	0.26	0.05	0.00	0.83	0.83	0.26	0.20	0.00	0.95
Observations	280	283	283	95	1176	277	284	281	275	281	282	273

Notes: We report results from the regression $Equity_t = \alpha_0 + \boldsymbol{\beta_1 S_t} + \boldsymbol{\beta_2 S_t} D_t^{post} + \beta_3 D_t^{post} + \boldsymbol{\psi_1 S_t} D_t^{crisis} + \psi_2 D_t^{crisis} + \varepsilon_t$. The regression sample includes 2948 observations. See also notes to table 6.

Table 8: Base Period Beta Estimates

	2009:2012 Base period								1992:2000	
	Oil		Equity		Metal		Inflation Comp.		Interest Rate	
	β	t-stat	β	t-stat	β	t-stat	β	t-stat	β	t-stat
Capacity utilization	-0.13	(-0.39)	-0.09	(-0.46)	-0.09	(-0.57)	-1.49	(-1.85)	3.00	(4.28)
Consumer confidence	-0.10	(-0.32)	0.02	(0.11)	-0.05	(-0.30)	1.07	(1.40)	2.32	(3.40)
Core CPI	-0.12	(-0.29)	-0.05	(-0.20)	-0.17	(-0.79)	-1.44	(-1.41)	2.74	(4.11)
GDP (advance)	-0.06	(-0.09)	-0.67	(-1.68)	0.75	(2.16)	2.30	(1.35)	-0.61	(-0.57)
Initial claims	0.28	(1.71)	0.16	(1.71)	0.00	(0.03)	1.02	(2.62)	0.52	(1.59)
ISM manufacturing	0.24	(0.70)	0.26	(1.34)	-0.18	(-1.04)	0.51	(0.62)	3.81	(6.04)
Leading indicators	0.29	(0.81)	0.14	(0.67)	0.19	(1.06)	-1.05	(-1.21)	2.02	(1.26)
New home sales	0.02	(0.04)	0.20	(0.63)	-0.04	(-0.16)	1.27	(0.93)	2.03	(3.19)
Nonfarm payrolls	0.86	(1.69)	0.69	(2.40)	0.12	(0.46)	1.15	(0.94)	4.17	(7.56)
Core PPI	0.83	(2.10)	0.31	(1.42)	0.24	(1.26)	2.02	(2.13)	0.84	(1.25)
Retail sales ex. autos	0.15	(0.36)	0.59	(2.51)	-0.17	(-0.81)	4.46	(4.48)	4.93	(4.54)
Unemployment rate	0.10	(0.29)	0.07	(0.34)	0.01	(0.07)	-0.20	(-0.24)	2.69	(3.71)
Observations	487		487		487		487		1109	
R-squared	0.52		0.53		0.49		0.46		0.46	

Notes: We report the β_j from the regression $Y_t = \alpha + \boldsymbol{\beta S_t} + \varepsilon_t$, where S_{jt} refers to the standardized an demeaned news for announcement j on day t. The parameters β_j measure the responses of each dependent variable to a one standard deviation news surprise for announcement j. For oil, equities, metals, and inflation compensation, the base period is 2009 to 2012. For interest rates, the base period is 1992 to 2000. See section 4.1 for more detail.

Table 9: Kernel Regression

Panel A: Varying equity beta for oil

$$Oil_t = \alpha(.) + \beta(.)Equity_t + \varepsilon_t, \quad \boldsymbol{\Gamma}(.) = \{\alpha(.), \beta(.)\}$$

Does the implied rate (Z_k) help explain the variation in the equity beta for oil?

Unrestricted	Restricted	p-value
$\boldsymbol{\Gamma}(Z_k)$	$\boldsymbol{\Gamma}$	0.00 (yes)
$\boldsymbol{\Gamma}(Z_k, VIX_k)$	$\boldsymbol{\Gamma}(VIX_k)$	0.00 (yes)
$\boldsymbol{\Gamma}(Z_k, OI_k)$	$\boldsymbol{\Gamma}(OI_k)$	0.00 (yes)

Do the alternatives (VIX_t, OI_t) help explain the variation in the equity beta for oil?

Unrestricted	Restricted	p-value
$\boldsymbol{\Gamma}(VIX_k)$	$\boldsymbol{\Gamma}$	0.02 (yes)
$\boldsymbol{\Gamma}(Z_k, VIX_k)$	$\boldsymbol{\Gamma}(Z_k)$	0.03 (yes)
$\boldsymbol{\Gamma}(OI_k)$	$\boldsymbol{\Gamma}$	0.00 (yes)
$\boldsymbol{\Gamma}(Z_k, OI_k)$	$\boldsymbol{\Gamma}(Z_k)$	0.00 (yes)

Panel B: Varying responsiveness to surprises

$$Y_t = \alpha(.) + \boldsymbol{\beta}(.)\boldsymbol{S_t} + \varepsilon_t, \quad \boldsymbol{\Gamma}(.) = \{\alpha(.), \boldsymbol{\beta}(.)\}$$

Does the implied rate (Z_k) help explain the variation in the responsiveness to surprises?

		Dependent Variable (Y_t)		
Unrestricted	Restricted	Oil	Equity	Interest Rate
$\boldsymbol{\Gamma}(Z_k)$	$\boldsymbol{\Gamma}$	0.00 (yes)	0.00 (yes)	0.00 (yes)
$\boldsymbol{\Gamma}(Z_k, VIX_k)$	$\boldsymbol{\Gamma}(VIX_k)$	0.00 (yes)	0.02 (yes)	0.00 (yes)
$\boldsymbol{\Gamma}(Z_k, OI_k)$	$\boldsymbol{\Gamma}(OI_k)$	0.01 (yes)	0.03 (yes)	0.00 (yes)

Do the alternatives (VIX_t, OI_t) help explain the variation in the responsiveness to surprises?

		Dependent Variable (Y_t)		
Unrestricted	Restricted	Oil	Equity	Interest Rate
$\boldsymbol{\Gamma}(VIX_k)$	$\boldsymbol{\Gamma}$	0.38 (no)	0.61 (no)	0.13 (no)
$\boldsymbol{\Gamma}(Z_k, VIX_k)$	$\boldsymbol{\Gamma}(Z_k)$	0.30 (no)	0.34 (no)	0.09 (yes)
$\boldsymbol{\Gamma}(OI_k)$	$\boldsymbol{\Gamma}$	0.51 (no)	0.30 (no)	0.11 (no)
$\boldsymbol{\Gamma}(Z_k, OI_k)$	$\boldsymbol{\Gamma}(Z_k)$	0.38 (no)	0.30 (no)	0.09 (yes)

Table 10: Inventory Surprises Pre- and Post- Break Regression Results

	Inventory Surprises		Inventory Changes	
	Oil	Equity	Oil	Equity
Oil-specific news	-0.45	0.11	-0.35	-0.03
	(-2.80)	(1.33)	(-3.90)	(-0.68)
Oil-specific news *Post-break Dummy	-0.02	-0.14	-0.02	-0.01
	(-0.11)	(-1.32)	(-0.09)	(-0.16)
Federal Funds Rate Surprise	-0.02	-0.15	-0.04	-0.15
	(-0.10)	(-1.49)	(-0.20)	(-1.65)
Post-break Response				
$\beta_1 + \beta_5$	-0.47	-0.03	-0.37	-0.04
F-stat	13.48	0.20	7.01	0.41
p-value	0.00	0.66	0.01	0.52
Observations	778	778	1318	1318

Notes: We report β_1, β_5, and β_6 from the regression $Y_t = \beta_0 + \beta_1 S_t + \beta_2 D_t^{crisis} + \beta_3 S_t D_t^{crisis} + \beta_4 D_t^{post} + \beta_5 S_t D_t^{post} + \beta_6 S_{FF_t} + \varepsilon_t$, where S_t refers to the standardized news for oil inventories on day t, or is proxied for by the change in the level of oil inventories. The parameter β_1 measures the response of each dependent variable to a one standard deviation news surprise in oil inventories before the break, while $\beta_1 + \beta_5$ measures the post-break response. See section 6 for more detail.

Figure 1: Oil and Equity Prices

Panel (a): Price of the front-month futures contract for WTI crude oil in dollars per barrel. Panel (b): Level of the equity price index, indexed to April 5, 1983 = 100. Panel (c): One-year rolling correlation between daily oil and equity returns.

(a) Price of WTI Oil ($ per barrel)

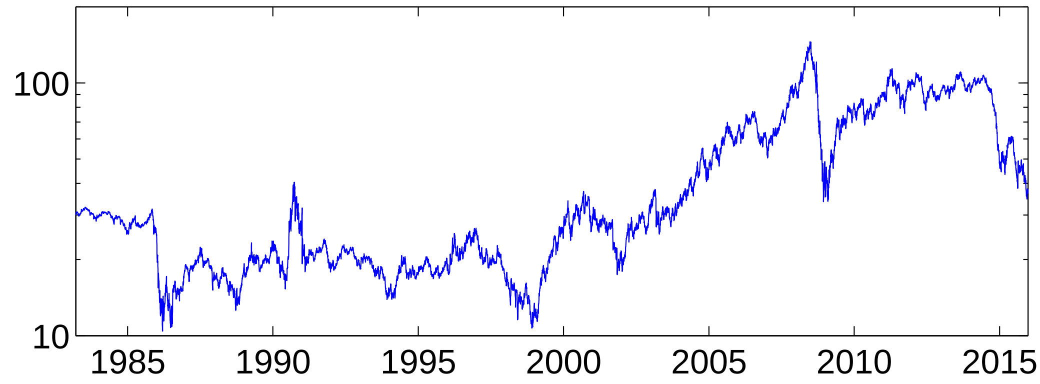

(b) Equity Index

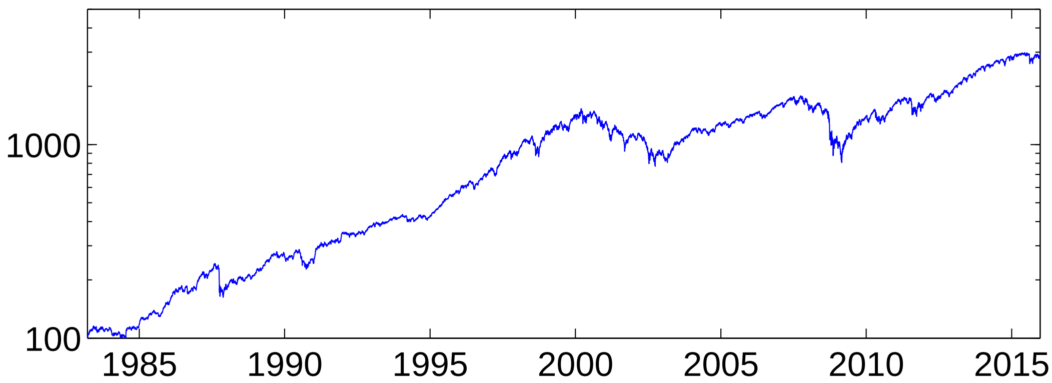

(c) Oil-Equity Rolling Correlation

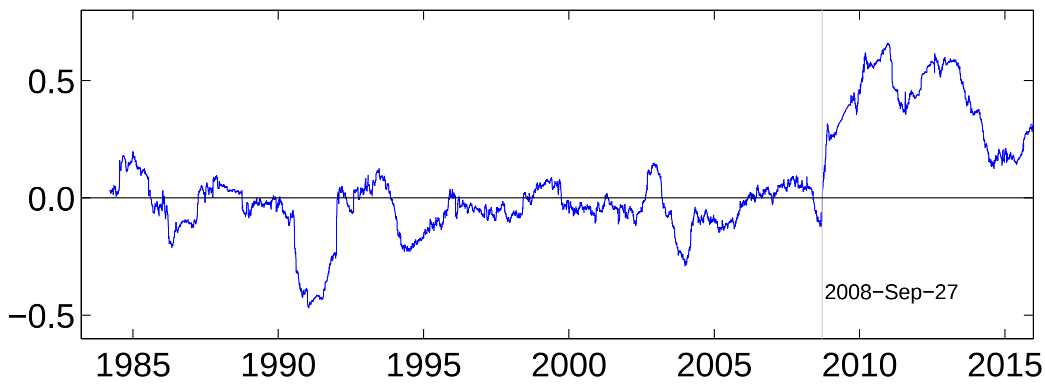

Figure 2: Rolling Equity Betas

Equity betas for $Y_t \in \{Oil_t, Metals_t, InflationCompensation_t\}$ are estimated over rolling samples of one year for the model $Y_t = \alpha + \beta Equity_t + \varepsilon_t$.

(a) Oil

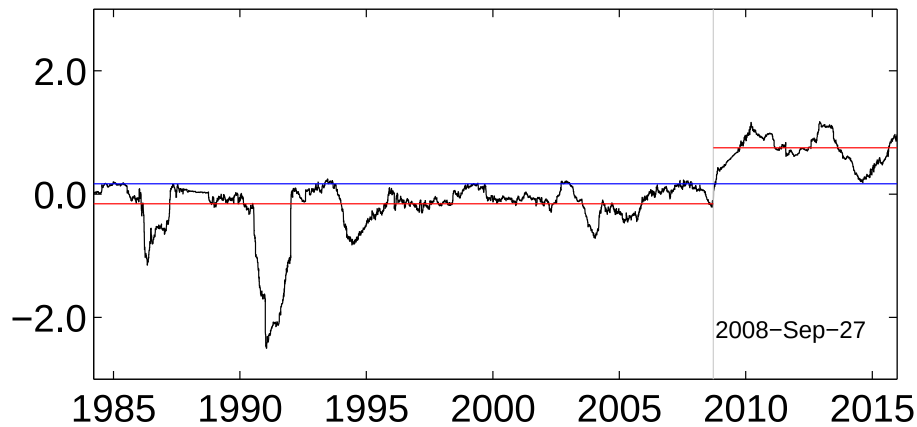

(b) Metals

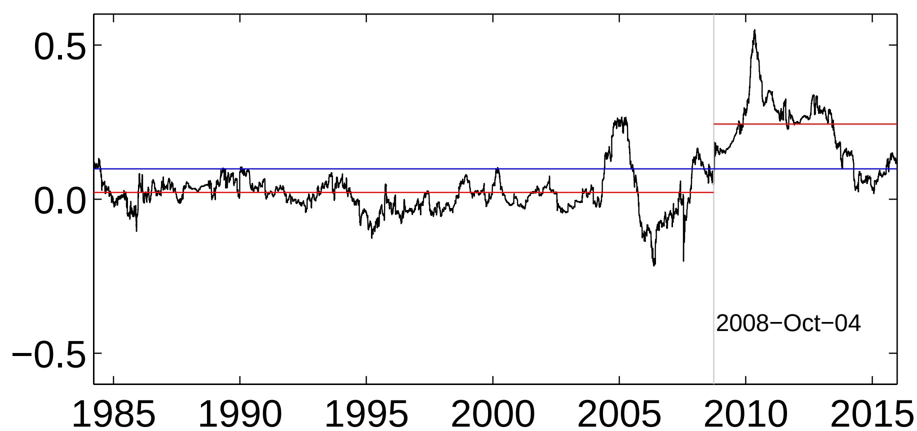

(c) Inflation Compensation (5 year)

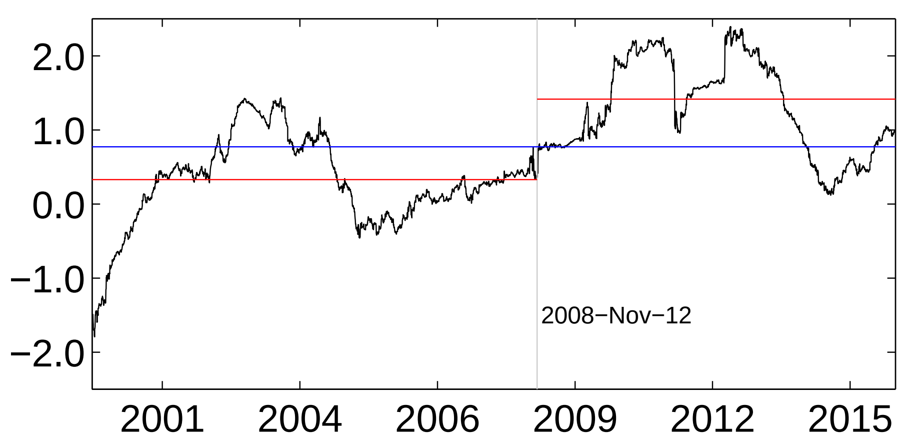

Figure 3: Oil and Equity Returns

The scatter plots of the standardized series for oil and equity returns are generated by subtracting the sample mean and dividing by the sample standard deviation for each series.

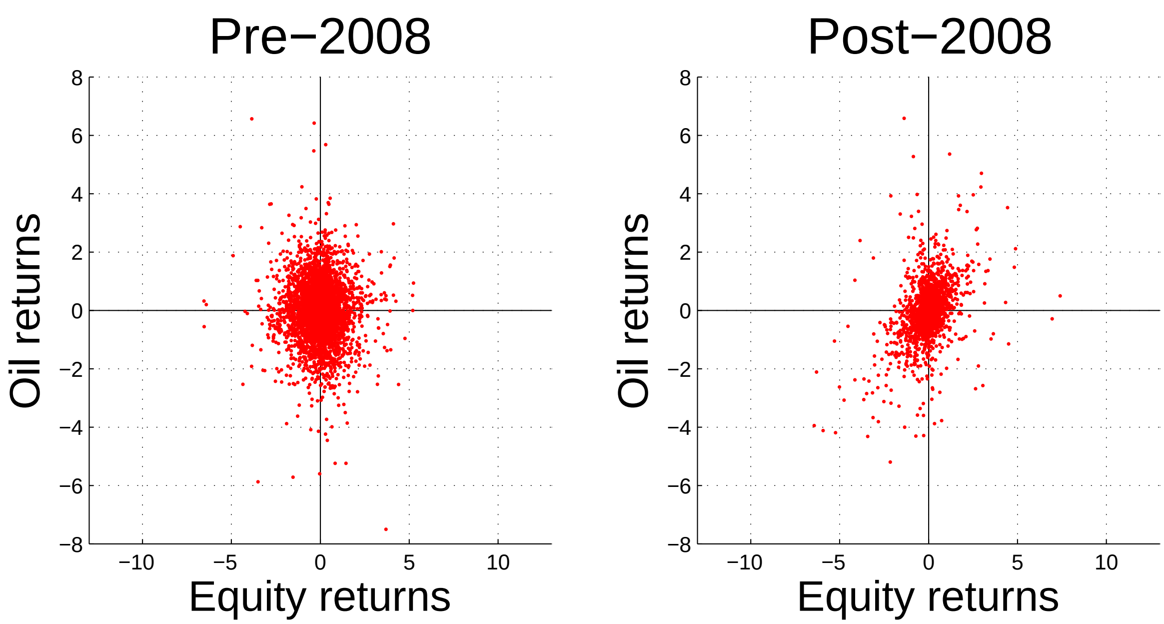

Figure 4: Oil and Equity Responsiveness to Surprises

For each dependent variable $Y_t \in \{Oil_t, Equity_t\}$, the monthly series for δ^τ is estimated using the regression $Y_t = \alpha^\tau + \delta^\tau \hat{\boldsymbol{\beta}} \boldsymbol{S_t} + \varepsilon_t^\tau$. $\hat{\boldsymbol{\beta}}$ is the vector of fixed parameters estimated from the regression $Y_t = \alpha + \boldsymbol{\beta} \boldsymbol{S_t} + \varepsilon_t$ using the subsample 2009 to 2012.

(a) Oil

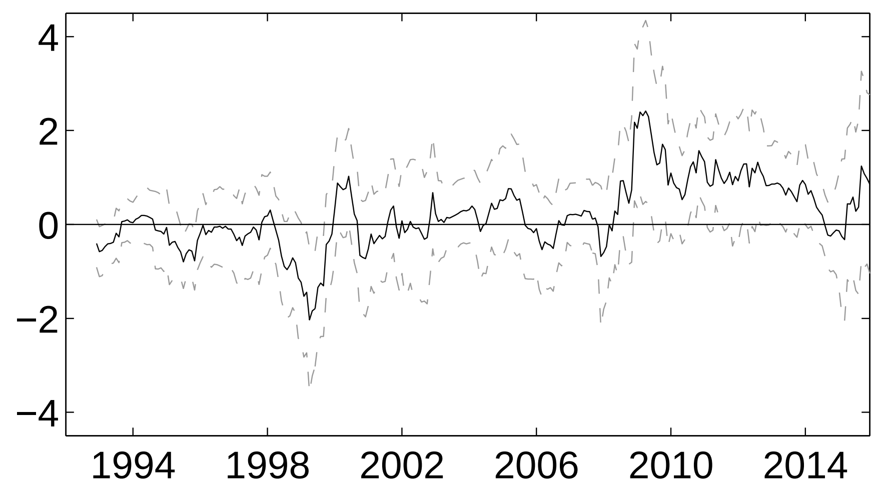

(b) Equity

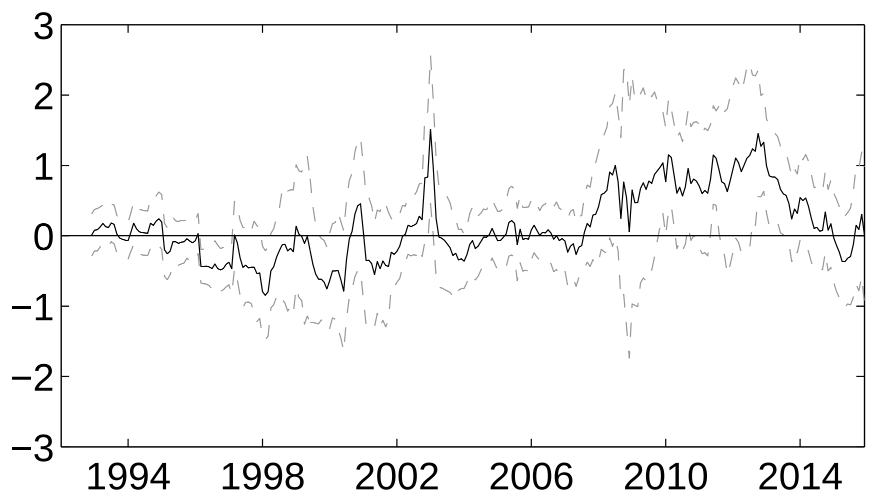

Figure 5: Responsiveness to Surprises

For each dependent variable Y_t, the monthly series for δ^τ is estimated using the regression $Y_t = \alpha^\tau + \delta^\tau \hat{\boldsymbol{\beta}} \boldsymbol{S_t} + \varepsilon_t^\tau$. $\hat{\boldsymbol{\beta}}$ is the vector of fixed parameters estimated from the regression $Y_t = \alpha + \boldsymbol{\beta} \boldsymbol{S_t} + \varepsilon_t$ using the subsample 2009 to 2012.

(a) Oils

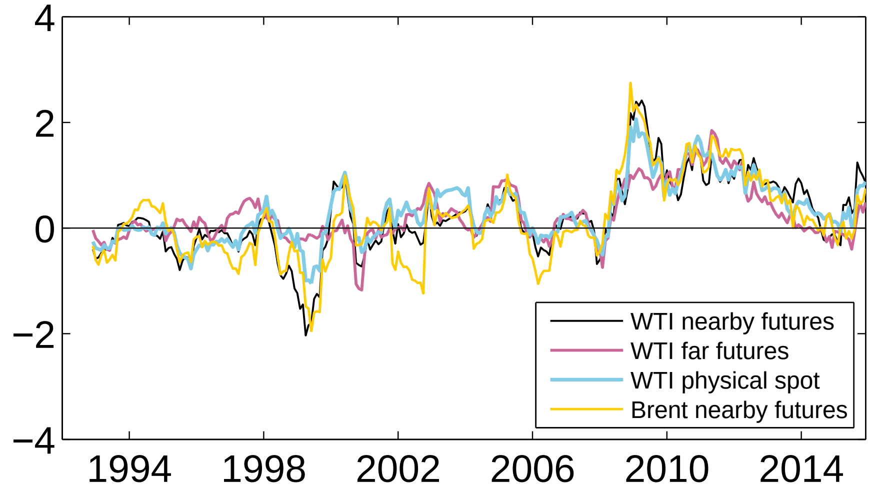

(b) Metals and inflation compensation

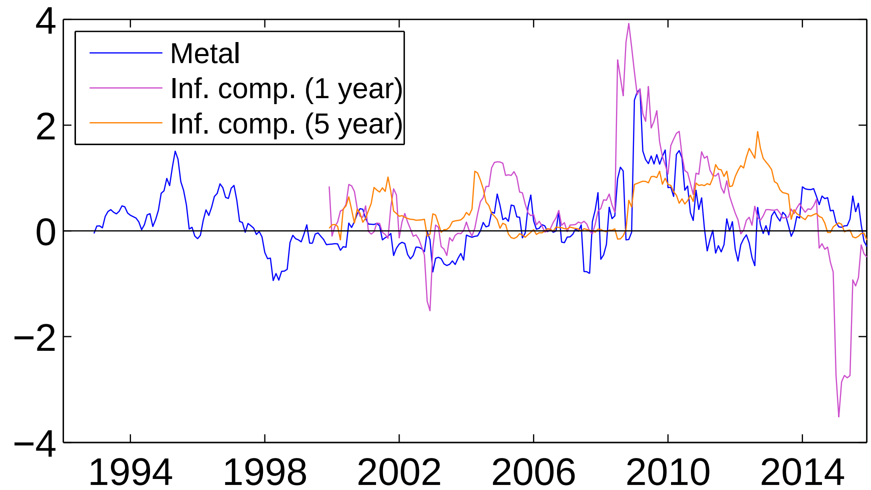

Figure 6: Interest Rate Responsiveness to Surprises

For each dependent variable Y_t, the monthly series for δ^τ is estimated using the regression $Y_t = \alpha^\tau + \delta^\tau \hat{\boldsymbol{\beta}} \boldsymbol{S_t} + \varepsilon_t^\tau$. $\hat{\boldsymbol{\beta}}$ is the vector of fixed parameters estimated from the regression $Y_t = \alpha + \boldsymbol{\beta} \boldsymbol{S_t} + \varepsilon_t$ using the subsample 1992 to 2000 for interest rates and 2009 to 2012 for oil and equity returns.

(a) Interest rates

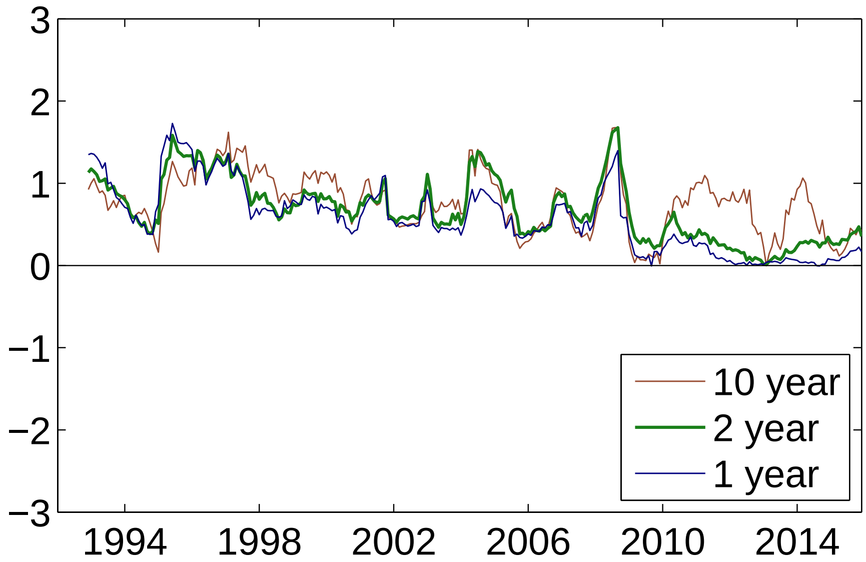

(b) Oil, equity, and interest rate

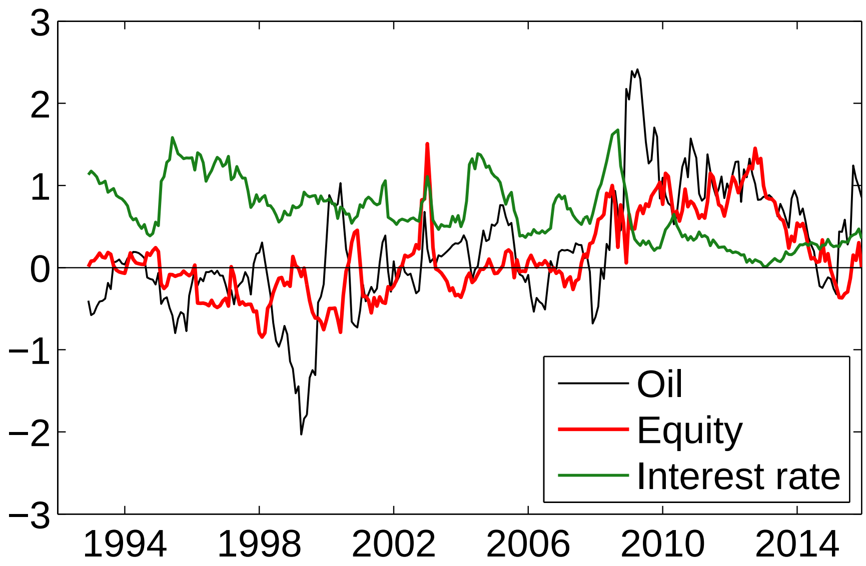

Figure 7: Kernel Estimation of the Equity Beta for Oil: Implied Interest Rate

Panel (a): The actual federal funds rate and the implied federal funds rate, defined as the rate implied by the modified Taylor rule in Bernanke (April 28, 2015). The implied rate is intended to capture the target federal funds rate as implied by the current state of the economy, without censoring due to the ZLB. Panel (b): Estimate of the equity $\beta(Z_k)$ from the regression $Oil_t = \alpha(Z_t) + \beta(Z_t)Equity_t + \varepsilon_t$.

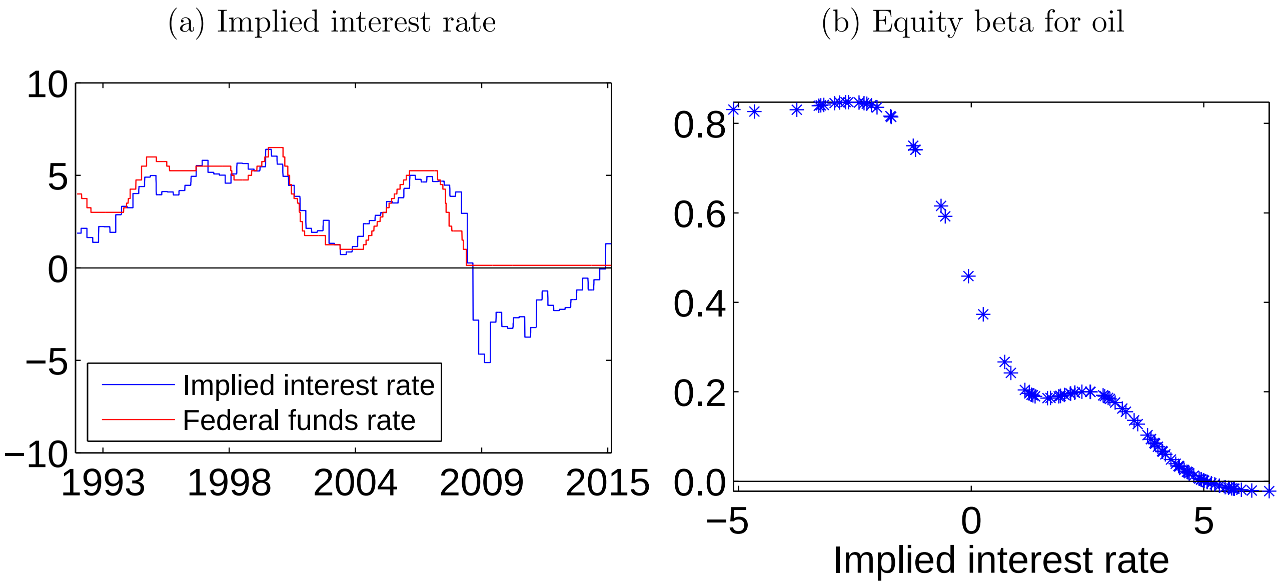

Figure 8: Varying Responsiveness to Surprises: Implied Interest Rate

For each of our dependent variables, $Y_t \in \{Oil_t, Equity_t, InterestRate_t\}$, we estimate the regression $Y_t = \alpha(Z_t) + \boldsymbol{\beta}(Z_t)\boldsymbol{S_t} + \varepsilon_t$. The average responsiveness to surprises, $\overline{\boldsymbol{\beta}}(Z_k) = \frac{1}{12}\sum_{j=1}^{j=12}\omega_j\hat{\beta}_j(Z_k)$, weights the estimated $\hat{\beta}_j(Z_k)$ by the frequency ω_j of announcement j in the estimation.

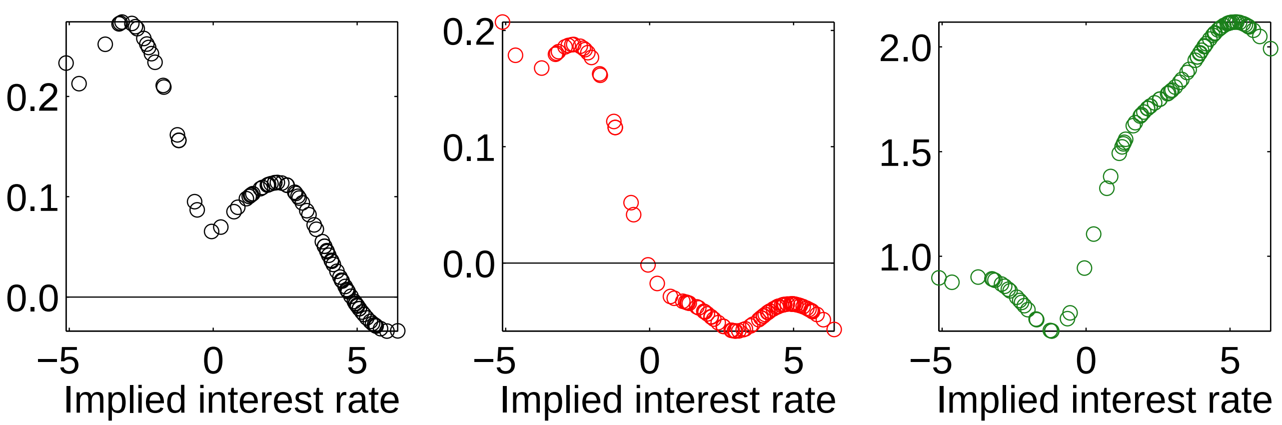

Figure 9: VIX and Open Interest

Panel (a): The 90-day rolling average of the VIX. Panel (b): The 90-day rolling average of open interest across all maturities of WTI futures contracts.

(a) VIX

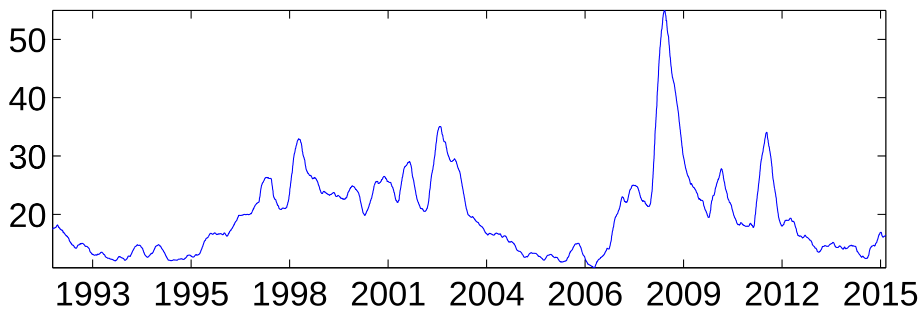

(b) Open interest

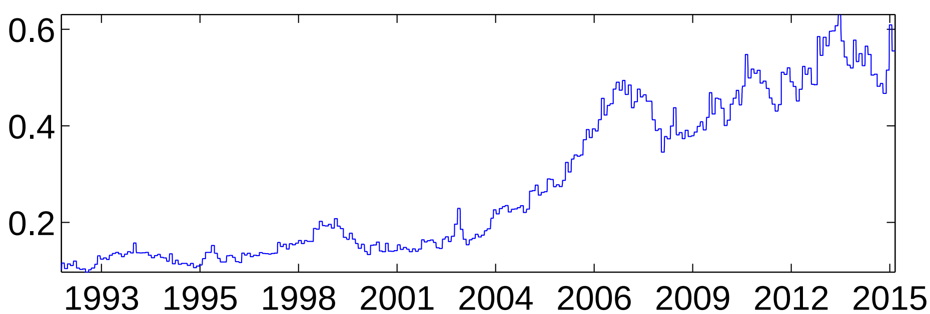

Figure 10: Kernel Estimation of the Equity Beta for Oil: VIX and Open Interest

Estimates of the equity $\beta(X_k)$ from the regression $Oil_t = \alpha(X_t) + \beta(X_t)Equity_t + \varepsilon_t$. X_k is the 90-day rolling average of the VIX for panel (a) and the 90-day rolling average of open interest for panel (b).

(a) Oil (b) Equity

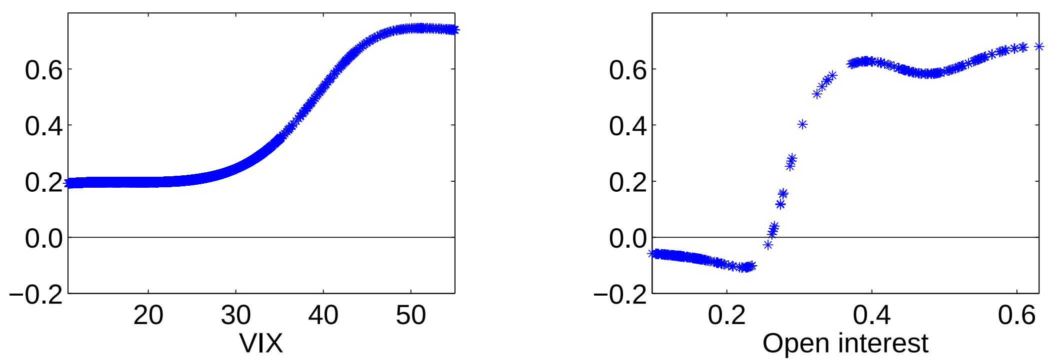

Figure 11: Varying Responsiveness to Surprises: VIX and Open Interest

The average responsiveness to surprises, is estimated using the 90-day rolling average of the VIX for panels (a) through (c) and using the 90-day rolling average of open interest for panels (d) through (f). See also notes to Figure 8.

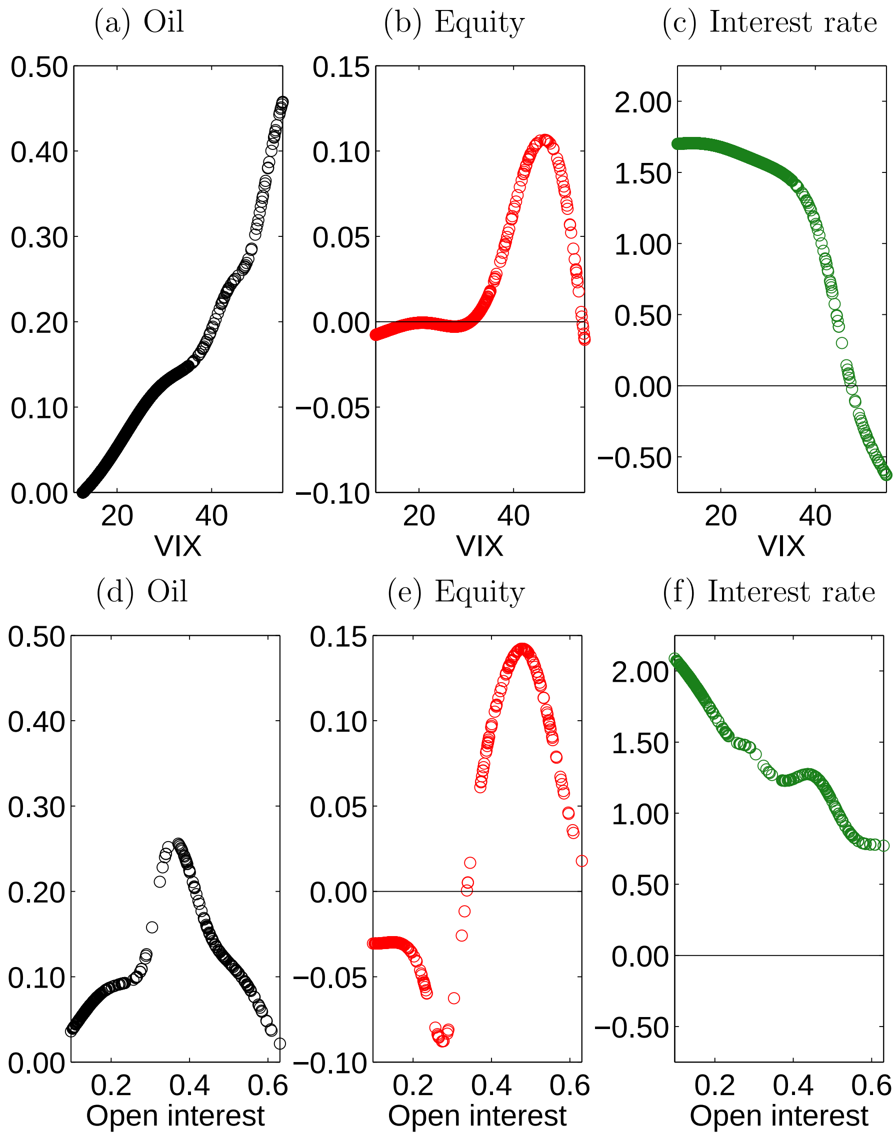

Figure 12: Response to Oil Supply Shock

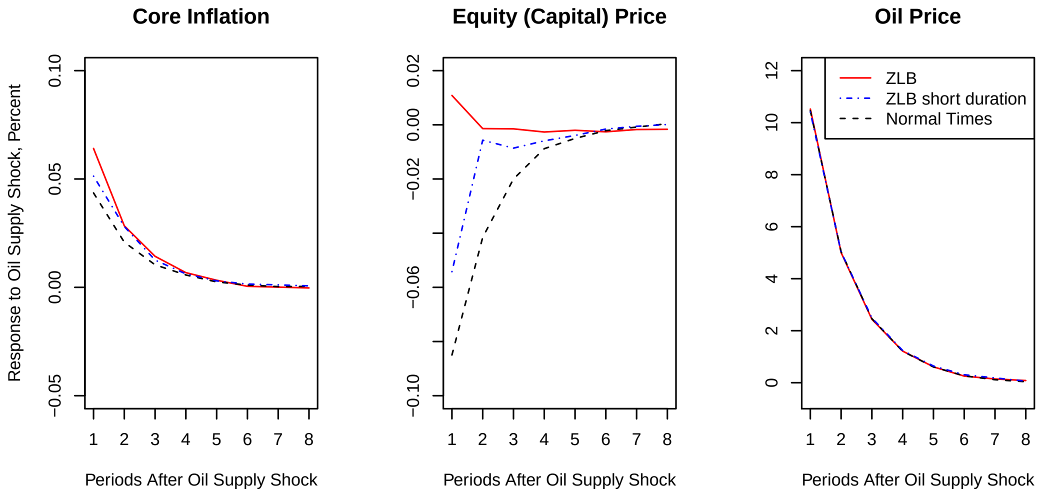

Figure 13: Response to Oil Demand Shock

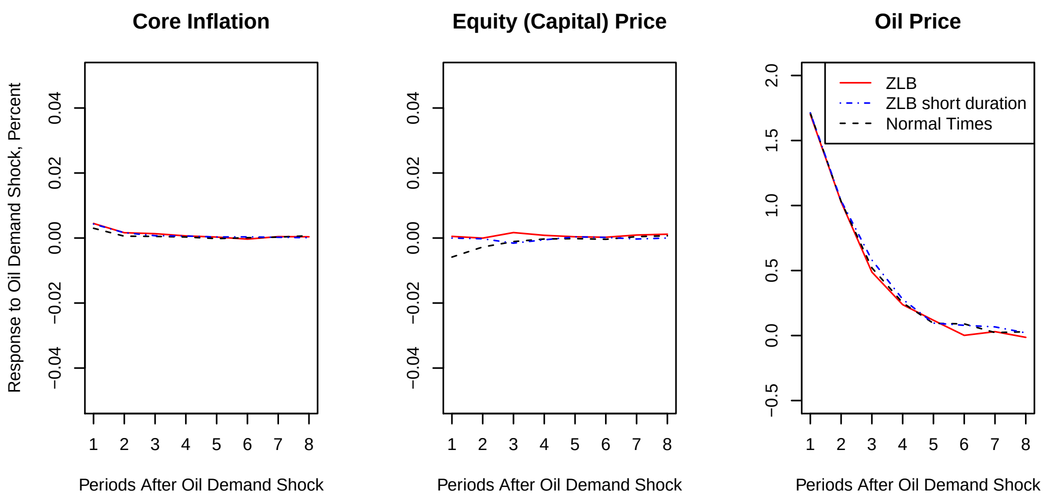

Figure 14: Response to Government Spending Shock

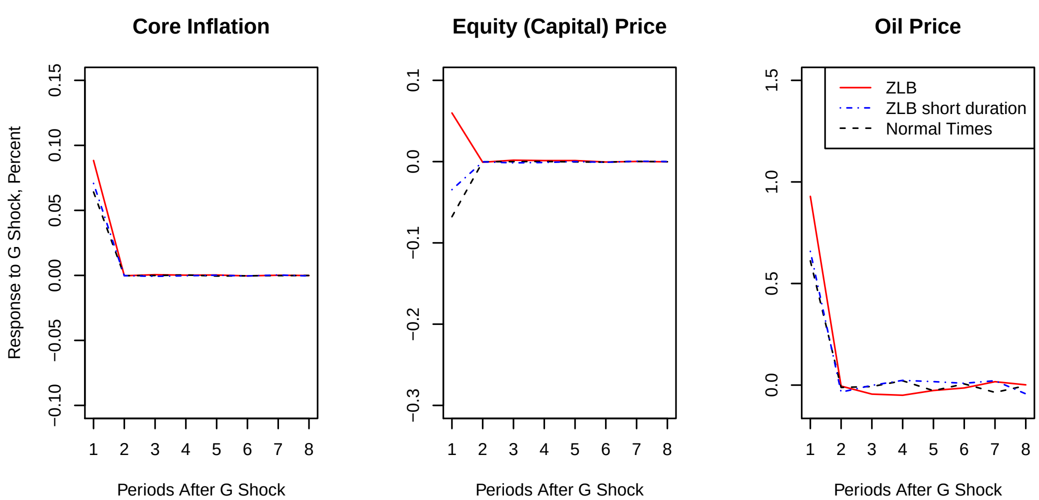

Figure 15: Rolling Oil Inventory Betas for Oil and Equity Returns

For each of our dependent variables, $Y_t \in \{Oil_t, EquityExEnergy_t\}$, we plot the one-year rolling β_1 from the regression $Y_t = \beta_0 + \beta_1 S_t + \beta_2 S_{FF_t} + \varepsilon_t$, where S_t refers to the standardized news for oil inventories on day t in the shorter sample and the change in the level of oil inventories in the longer sample.

(a) Oil

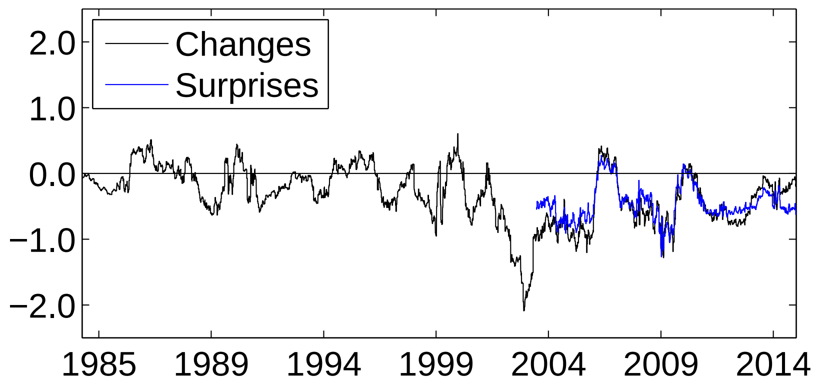

(b) Equity ex. energy