The Macroeconomic Impact of Oil Industry Uncertainty: New Evidence from Millions of Financial Analyst Forecasts

Xiaohan Ma* Roberto Samaniego†

September 2, 2019

Abstract

We develop measures of energy industry uncertainty using analyst forecasts drawn from a large firm-level dataset. We find that our uncertainty measure is related to future economic downturns, and that the dispersion of analyst forecasts may serve to forecast future downturns. An increase in our uncertainty also has adverse effects on the domestic oil sector as well as the world oil market. Our paper underlines the importance of the energy sector in understanding and forecasting broader macroeconomic dynamics, as well as the ways in which uncertainty affects oil markets.

Keywords:

Abstract

Oil industry uncertainty, business cycles, oil markets, analyst forecasts.

JEL Codes: D80 E32 E37 E47 E71 G17.

*Texas Tech University, Department of Economics, P.O. Box 41014, Lubbock, Texas 79409-1014 U.S.A. Email: xiaohan.ma@ttu.edu.

†Department of Economics, The George Washington University, 2115 G Street, NW Monroe Hall 340, Washington, DC 20052, U.S.A. Email: roberto@gwu.edu.

1 Introduction

The macroeconomic impact of events originating in or propagating through oil markets is a major research topic. An independent literature studies the impact of uncertainty shocks on business cycle dynamics. However, it remains an open question whether *uncertainty in oil markets* might have significant macroeconomic impact.

In this paper we develop several measures of oil industry uncertainty, and study their macroeconomic impact, at monthly frequency. Our measures are not based on traditional indicators such as oil prices, oil production, or oil price volatility. Oil prices and oil production are endogenous – responding to the forces of energy demand and supply, which themselves include macroeconomic variables. Instead, we use the *forecasts and forecast errors* from a large survey set of financial analysts regarding the financial variables of firms in the US oil and gas sector. The baseline measure is the median 12 month-ahead earnings-per-share (EPS) absolute forecast error, drawn from the Institutional Brokers' Estimate System or I/B/E/S.

Our reasoning is as follows. When analysts make a forecast regarding, for example, the EPS of a firm, they are using the best available information at the time that the forecast is made. Indeed, since their compensation and reputation are based on the usefulness of the forecasts they make, they have an incentive to incorporate and process as much information as is available in their forecasts. There will typically be a certain amount of background uncertainty in the environment which results in a non-zero forecast error one way or the other in normal times. However, when the absolute forecast errors are *larger* than normal, this indicates the impact of factors of uncertainty that were not adequately foreseen or processed at the time the forecasts were made.

We find that our oil industry uncertainty measure appears to be distinct from overall macroeconomic uncertainty – identified using the Jurado et al (2015) measure of aggregate uncertainty. While there are periods of time when the two appear to co-move, there are also periods when they do not co-move at all. This suggests that oil industry uncertainty and aggregate uncertainty are different concepts – although aggregate uncertainty might sometimes have an impact on the oil market, and viceversa.

To study the impact of our oil industry uncertainty measure on economic aggregates and on oil markets, we estimate a series of structural vector auto-regressions (SVAR). We also provide further evidence on whether oil industry uncertainty is different from aggregate uncertainty, and in what way. The VARs contain our baseline oil industry uncertainty measure, demand and supply factors for the oil market, as well as the oil price, in addition to some representative macroeconomic and policy variables. Since many of the channels through which aggregates and oil markets are affected are explicitly estimated in the procedure, any impact of oil industry uncertainty is more likely to be exogenous and thus capture causality, as the impact from or through these other variables is already considered. This is a distinct advantage of our approach to measurement and estimation over, for example, using indicators based on oil prices and/or estimating a VAR with only oil market related variables, where endogeneity might otherwise present a challenge to estimation and/or interpretation of econometric relationships.

We find that our baseline oil industry uncertainty lowers US output and the US price level, as well as the federal funds rate, which suggests that oil industry uncertainty behaves like a negative aggregate demand shock. In the oil market, an increase in oil industry uncertainty lowers US oil production and the oil price. We reach the same conclusions when we include both oil industry uncertainty and aggregate uncertainty measures in the same VAR. This implies that our oil industry uncertainty measure contains information about the energy market that is absent from aggregate uncertainty measures. Finally, the stock market responds little to changes in oil industry uncertainty when the estimation is conducted over the whole sample period ($1982 - 2018$). However, when the post-Great Recession period is excluded from the sample, oil industry uncertainty raises the stock index. This is likely because an increase in oil industry uncertainty lowers the oil price, which is a positive signal for the overall economic activity of sectors that use energy.

We develop several other measures of oil industry uncertainty. One is the measures of forecast dispersion among financial analysts regarding firms in the oil and gas sector. When analyst forecasts about the same firm are more dispersed than normal, this indicates the insufficiency of information for arriving at a conclusion about the future, possibly reflecting uncertainty of a different form than a measure based on forecast errors. These dispersion measures are particularly useful for forecasting because they do not require future data for their computation at all. When estimated in the same VAR model as the baseline, we find generally similar behavior, indicating that our approach to measuring oil industry uncertainty can be computed in real time as new forecasts are made and entered into the I/B/E/S database.

Other versions of our baseline oil industry uncertainty measure are developed such as the median absolute forecast errors constructed using 3-month, 6-month and 9-month forecasts. Again, VAR exercises with these alternative quarter-based measures find that they generally behave similarly to our benchmark measure. The advantage of these measures is that they do not require waiting as long as the baseline measure in order to observe the realization of the EPS forecast.

Finally, we calculate the median absolute forecast error regarding oil producers *relative* to the median absolute forecast error regarding firms *outside* the oil sector, or oil industry "relative" uncertainty. This is a way of ensuring that any increases in our measure are due to events originating from the oil industry or disproportionately affecting the oil industry relative to other industries in the economy. We find that the behavior of this measure is similar to our baseline measure. However, there is one difference: oil industry relative uncertainty has a clear positive impact on the stock market, similar with that obtained from the estimation with the baseline oil industry uncertainty over the pre-2007 sample. This result reassures us that oil uncertainty can have an independent impact on the macroeconomy that is different from aggregate uncertainty and mainly reflects information originating in and/or specifically influencing the oil industry, considering the fact that aggregate uncertainty rose dramatically before the Great Recession and had significantly negative impact on the stock market.[1]

Our study is related to three strands of literature. First, it is widely documented that changes in economic uncertainty are an important driving force of business cy-

[1] This is found in Bloom (2009), Jurado et al. (2015), Ma and Samaniego (2019), among others.

cle fluctuations and oil market dynamics.[2] However, even though the oil market is a major component of global markets, the literature has devoted relatively little attention to the question of how uncertainty originating in the oil industry accounts for macroeconomic and oil market fluctuations. Our paper contributes to the literature by developing measures of oil industry uncertainty based on millions of financial analyst stock forecast errors, and empirically investigating the dynamic impact of oil industry uncertainty – by itself, and conditional on the impact of aggregate uncertainty.

Second, while an extensive literature focuses on oil market uncertainty, almost all papers study the impact of oil *price* uncertainty. Examples include Elder and Serletis (2010), which assumes oil price shocks can lead to variations in oil price volatility and uses a GARCH model to estimate the impact of oil price volatility on the US real output; Maghyereh et al. (2016), which uses the crude oil implied volatility index as a measure of oil price uncertainty; and Yin and Feng (2019), which measures oil market uncertainty using the volatility risk premium as in Carr and Wu (2009), computed using oil futures prices. A drawback of using oil prices to measure uncertainty is that the oil price itself responds to changes in other macroeconomic aggregates, so that the direction of causality between oil prices and macroeconomic variables is unclear – as underlined by Kilian (2009). This suggests that understanding the role of uncertainty in oil markets requires some exogenous measure of uncertainty originating in energy markets. As well as measuring oil industry uncertainty in a manner that uses neither oil prices nor oil production as an input, we also estimate a VAR containing several macroeconomic and oil industry indicators to study how our oil uncertainty measures affect macroeconomic aggregates. Thus, our uncertainty measure is exogenous, and the causal relationship between oil industry uncertainty and macroeconomic variables is well-defined in our paper. In addition, volatility is at best a noisy proxy for uncertainty, as discussed in Jurado et al (2015). Our measures

[2]See Bloom (2009), Bachmann et al. (2013), and Jurado et al (2015), inter alia, on the impact of aggregate uncertainty on the macroeconomy, and see Van Robays (2016), Joëts et al. (2017), and Bakas and Triantafyllou (2018) on oil market dynamics. The potential impact of industry-specific uncertainty is studied in Carriero et al. (2018) and Shin and Zhong (2018) for financial markets, and in a comprehensive breakdown of industries in Ma and Samaniego (2019).

capture uncertainty in the form of an increase in the difficulty of analysts arriving at accurate or agreed forecasts, rather than volatility.

Given that we develop some of our oil uncertainty measures based on forecast errors, our work is perhaps most closely related to Jo (2014) in this strand of literature. Jo (2014) models oil price uncertainty as the time-varying standard deviation of the one-quarter-ahead oil price forecasting error, interpreted as an exogenous process that is independent from the level of oil price shocks, and examines how it affects the global economy. Aside from the potential issues related to measuring oil market uncertainty using statistics based on oil prices, a drawback of this measure is that, in the author's own words, "predictable variation in a monthly or quarterly volatility measure is negligible at the horizons relevant to the cash flow of investment decisions, removing the theoretical rationale for a non-negligible contractionary effect on real activity."[3] Our baseline forecast horizon is 12-months, which is likely to be a more relevant horizon for investment decisions, and therefore captures more accurately the effect on real activity.

Finally, there is a large body of literature investigating the impact of oil shocks and uncertainty shocks on the oil market and macroeconomic aggregates. A common model used in these papers is the SVAR, with either macroeconomic variables or oil market variables. For example, Bloom (2009) uses an eight-variable VAR to estimate the macroeconomic impact of uncertainty. Kilian (2009) differentiates oil demand and oil supply shocks in a VAR with four variables related to the oil market to explain oil price variations. In contrast, our baseline model is a combination of both macroeconomic variables and oil market variables. In this way, the impact of oil industry uncertainty on the oil market is conditional on how it is affected by the shocks to important macroeconomic aggregates such as industrial production and the stock market, and policy variables such as the federal funds rate. At the same time, the model can also be used to study how oil industry uncertainty affects the macroeconomy, again conditional on the impact of various other shocks.

[3] Jo (2014) adds that "one of the interesting questions for future research will hence be the development of alternative specifications based on extraneous survey-based measures of oil price uncertainty."

Our paper is structured as follows. Section 2 describes our measure of uncertainty. Section 3 describes in detail the data that we use, and some basic properties of the measure. Section 4 shows the impact of our measures of oil industry uncertainty on macroeconomic and oil market dynamics. Section 5 concludes with a discussion of potential future work.

2 The Measure

The premise behind our measurement strategy is that changes in uncertainty, and thus the predictability of the economic environment – at the aggregate or at the industry level – will be reflected in that analyst forecasts are of lower accuracy than usual, or that analysts display excessive disagreement. We use this idea to develop a measure of uncertainty for energy markets. The approach to measurement is similar to that in Ma and Samaniego (2019), and our presentation of the methodology follows theirs.

Time is discrete and divided into days which are collected into months. Let $M \subset \mathbb{N}$ be the set of months, numbered consecutively, and and let $t \in M$ be a month. Then, define $D_t \subset [t, t+1)$ as the set of days in the month t, so that $d \in D_t$ represents a day in month t. Let $S_{i,t}$ be a statistic about a firm i observed on day d, and let $F\left[S_{i,d^*}|I_{j,d}\right]$ be the forecast about the realization of statistic S at firm i on a future day d^*, using the information $I_{j,d}$ available to them on day d to forecaster j. This means that d^* minus d is the forecast horizon. Note that d^* will not be in the same month if the forecast horizon is longer than a month: this will be the case in general in our data. We define the firm-level forecast error as the difference between the forecast made on day d about statistic S at date d^*, and the actual realization of the statistic on day d^*:

$$FE_{i,d} = F\left[S_{i,d^*}|I_{j,d}\right] - S_{i,d^*}. \tag{1}$$

In our benchmark measure, the forecast period is a year, but we also look at quarterly forecasts. If more than one analyst makes a forecast about firm i on day d, we define

$F\left[S_{i,d^*}|I_{j,d}\right]$ as the average forecast error made about firm i on day d.[4]

There are thousands of forecasts made every day about different firms. To measure uncertainty U_t in month t, we will focus on the uncertainty experienced by a typical firm. In particular we look at the *median* absolute forecast error across all firms within the month. We focus on the median in order to avoid being swayed by individual outliers, which a large data set of forecasters will inevitably have.[5] In addition, we define uncertainty based on the median of the *absolute value* of the forecast error. This way uncertainty is measured as lack of accuracy – regardless of the direction. Not doing so would lead to a measure of relative optimism or pessimism compared to the realization, not uncertainty. As discussed below, we also try several other approaches for robustness. In practice, all our measures will be monthly, the highest frequency for which we have data on industrial production. Thus, on each date within month t, we compute the median value of $\|FE_{i,d}\|$ within the month, pooling all firm-day forecasts within the month, which gives our baseline uncertainty measure for month t:

$$U_t = \text{median}\left\{\|FE_{i,d}\| : \forall d \in [t, t+1), i \in \Upsilon\right\}. \tag{2}$$

Notice that definition 2 restricts firms to be from some set Υ. For most of this paper, we will define Υ to be the set of firms in the oil and gas producing sector, i.e. firms reporting SIC codes between 1300 and 1389. We refer to this as *oil industry uncertainty* or just *oil uncertainty*.[6] Sometimes, however, we will define Υ to be the set of firms *outside* the oil and gas producing sector. We refer to this as *non-oil uncertainty*.

The specific statistics that we look at are forecasts of the earnings-per-share ratios

[4] In practice, only about 14 percent of all firm-day combinations have more than one analyst making a forecast about it, ranging from 2 to 5 analysts. For robustness, we also repeat this procedure looking at the average forecast about a particular firm *within a month*, in which case about 87 percent of firm-month combinations have more than one forecast, ranging up to 42. Results are similar.

[5] Indeed, we found that uncertainty measured using the mean rather than the median was extremely volatile and had no meaningful properties.

[6] According to the US Energy Information Agency, two thirds of energy in the US comes from oil and natural gas.

(EPS) of individual companies. We use EPS forecasts because they are the most widely available in our database, and also because EPS ratios are a basic indicator of the profitability of a share, and are thus widely understood and followed both by financial analysts and their clients.

A concern with the measure is that some variation in EPS ratios could be due to the fact that firms have different scales – or rather that the granularity of their share size may vary. As a result, we divide all of our forecast errors $FE_{i,d}$ by the price of the share of company i on the day d when the forecast was made. Conceptually, this measure has the interesting property that it can be interpreted as a forecast of inverse price-earnings ratios, a common statistic used for share valuations. To produce such a measure, we combine our forecasting data with data on share prices, which allows us to divide the EPS forecast error by the corresponding security prices.[7]

Later in the paper we will study the behavior of *relative* oil uncertainty. If we define Υ as the set of firms reporting SIC codes between 1300 and 1389, relative oil unertainty is defined simply as:

$$U_t = \frac{\text{median}\left\{\|FE_{i,d}\| : \forall d \in [t, t+1), i \in \Upsilon\right\}}{\text{median}\left\{\|FE_{i,d}\| : \forall d \in [t, t+1), i \notin \Upsilon\right\}}. \tag{3}$$

These measures are useful for measuring uncertainty as we can see whether or not forecasts made at a particular date were less accurate than usual, in an absolute or relative sense. However, a drawback is that they require future information for their computation – i.e. the econometrician must observe the realization of the forecasted variable S_{i,d^*}. As a result, we will also use measures of uncertainty based on forecast *dispersion*, i.e. based on the extent of disagreement among forecasters. These measures can be computed month-to-month in real time, as soon as the forecasts are reproted. For each firm i and within each month t, we compute

$$D_{it} = \text{Disp}\left\{F\left[S_{i,d^*}|I_{j,d}\right] : \forall d \in [t, t+1)\right\}$$

[7] Later on we also use different moments of these forecasts. In particular, in Section 4.2 we study several measures based on the extent of forecast *dispersion*.

where "Disp" indicates a measure of dispersion. Then, we measure uncertainty in the subset of firms $i \in \Upsilon$ (e.g. firms in the oil industry) using the formula:

$$U_t = \text{median}\left\{D_{it} : i \in \Upsilon\right\}. \tag{4}$$

3 Data

Our forecasts are drawn from the Institutional Brokers' Estimate System or I/B/E/S, available through a WRDS subscription and managed by Thomson Reuters. It contains analyst forecasts of several measures of interest to investors and researchers, the most widely-available being earnings per share (EPS) forecasts. I/B/E/S also reports realizations of the forecast data, collected from a variety of public data sources. Companies are included in the database as long as at least one analyst provides a forecast for that company. Forecasts are not included unless they are confirmed within 6 months.[8] Forecasts are collected each day as they are released by analysts.[9]

We focus on US firms. This yields about 4.7 million forecasts issued by about $1,500$ different brokers, who make forecasts about many firms over time. For each firm on each day we compute the average forecast error.[10] We take the absolute value of this average forecast error, and divide it by the share price of the forecasted firm on the day that the forecast was taken. Share price data are available from CRSP. The absolute forecast error normalized using share prices in this manner will be our empirical counterpart of the term $\|FE_{i,d}\|$ in equation 2. Our measure

[8]For further details, see https://wrdsweb.wharton.upenn.edu/wrds/support /Data/_001Manuals%20and%20Overviews/_003I-B-E-S/Release%20Notes/, last checked 3/20/2018.

[9]Later we also look at EPS forecasts made over different horizons. No forecasts for different horizons are made about the same firm on the same day by the same analysts. However, all analysts that make an annual forecast make a quarterly forecast about a given firm sometime that month. About 46 percent of forecasters who make an annual forecast about a given firm make a 2-quarter ahead forecast the same month, and about 39 percent for 3-quarter ahead forecasts.

[10]86 percent of them are single forecasts about a firm on a given day. The rest have 2 forecasters making forecasts about a firm on the same day, except for 0.29 percent of the sample which has $3-5$ forecasts. Averaging when there are multiple forecasters yields about 3 million day-firm observations.

of uncertainty is the median absolute value of these forecast errors across all firms within each month, starting in March 1982.[11] Thus it is the median forecast error by firm-day pair.[12]

As well as share price information, CRSP reports NAICS and SIC codes of these firms. This allows us to compute our uncertainty measure for subsets of the firms inside or outside the oil/gas industries, based on industry classification. We use SIC codes for these purposes because NAICS codes did not exist early in our sample. Based on this information we narrow our sample down to about 4 million observations, of which about a third of a million are analyst forecasts about oil and gas firms.

[11] This is the first month after which continuous series may be computed for oil industry uncertainty. The date is based on the month and year of the variable *anndats*.

[12] We find that some of our measures appear to have seasonal effects. In particular, our measure tends to decline from October to January, possibly due to the forecasters being better informed about firms' financial conditions as annual statements are compiled and delivered towards the end of the year. As a result we remove the mean value for each month from the data to remove any such seasonal effects.

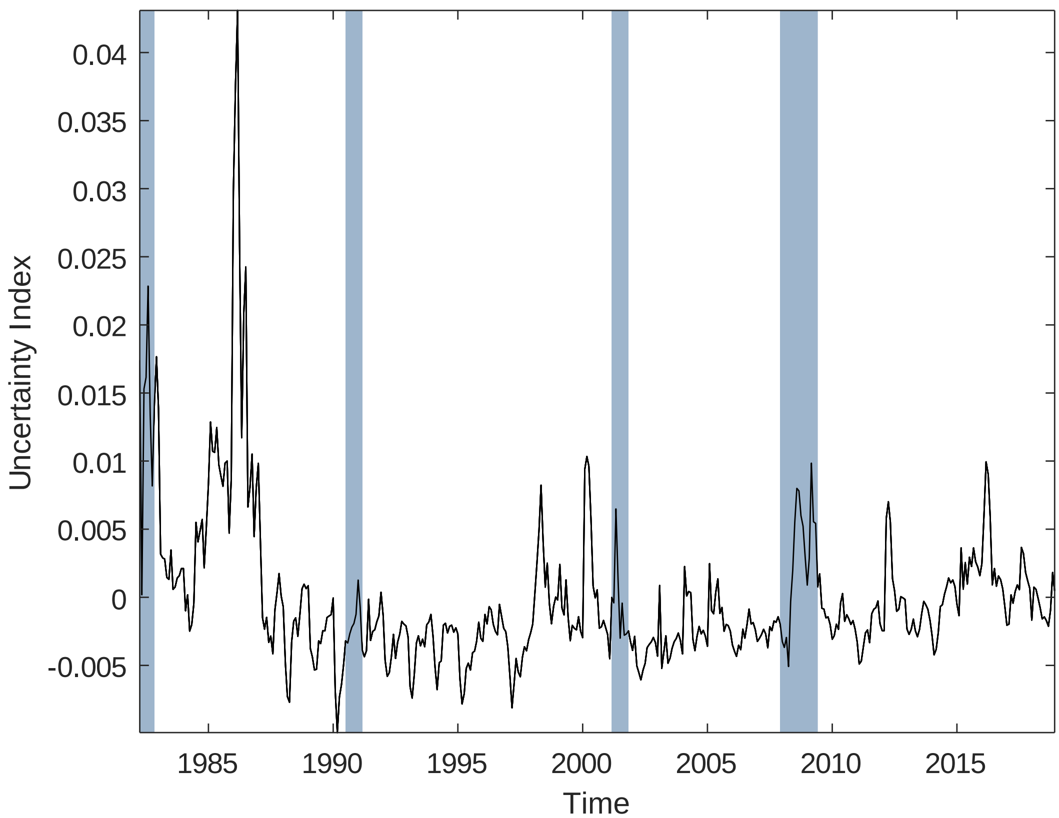

Figure 1 – oil industry uncertainty and the business cycle, 1981-2016. Bands represent NBER recession dates. The measure is the median absolute value of the forecast error from I/B/E/S by month. The forecast error is the difference between the 12-month EPS forecast and the realized EPS, deflated by share price, for firms reporting SIC codes between 1300 and 1389.

Figure 1 displays the series for oil industry uncertainty. Several observations stand out. First, the series appears to have a more or less stable level of uncertainty, punctuated by sharp spikes. This is consistent with the notion that there is a background level of oil industry uncertainty which is subject to occasional shocks.

Second, while some of these spikes coincide with recessions, many do not, including the largest spikes. Given that the literature suggests that aggregate uncertainty is related to recessions, this suggests that oil industry uncertainty is different from overall uncertainty, and thus a factor of uncertainty that could potentially have distinct effects on the oil industry and on aggregates.

To verify this conjecture, Figure 2 compares our oil industry uncertainty measure (OIU) to the aggregate uncertainty measure of Jurado et al (2015).[13] There are times when oil uncertainty co-moves with aggregate uncertainty, and in fact the correlation between the two series is 0.41 and significant. On the other hand, it is also clear that spikes in one series do not always coincide with spikes on the other. This suggests that, while there is a relationship between oil industry uncertainty and aggregate uncertainty, they are distinct forms of uncertainty, the economic impact of which an appropriate econometric specification should be able to tease apart.

[13] As discussed in Ma and Samaniego (2019), other popular measures of aggregate uncertainty behave similarly.

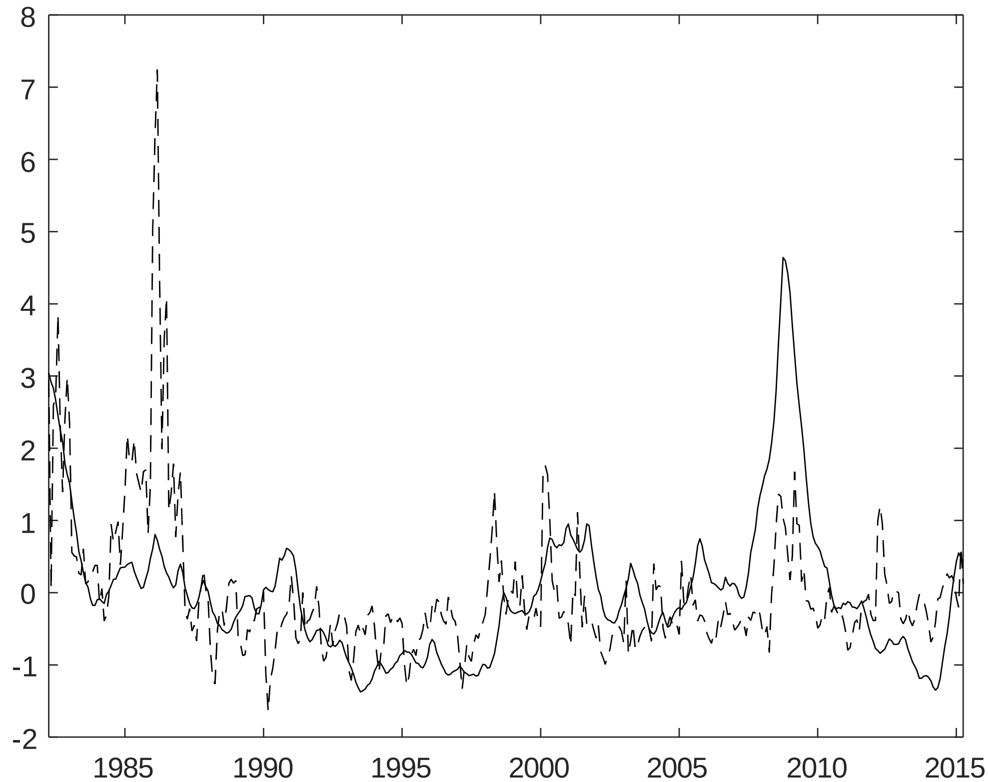

Figure 2 – Comparison of different uncertainty measures. The dashed line is our oil industry uncertainty measure and the thick line is the measure of Jurado et al (2015). measure. The sample period is 1982m3-2015m4.

4 Impact of Oil Industry Uncertainty

To investigate the role of oil industry uncertainty in characterizing the dynamics of macroeconomic aggregates and oil market, we use vector-autoregression (VAR) method to estimate the responses of key macro and oil market variables to innovations in our oil industry uncertainty, which we refer as oil uncertainty shocks.

The specification of the baseline VAR includes representative macroeconomic and oil market variables, where the macroeconomic elements are similar to that studied in Bloom (2009) and Samaniego and Ma (2019), and the oil market elements are similar to that studied in Kilian (2009), as to what variables to include and how to order them in the VAR. Following Bloom (2009), we include log S&P 500 index, federal funds rate, log CPI, and log US real economic activity approximated by US industrial production. Following Kilian (2009), we use log US crude oil production, log world crude oil production, world real economic activity approximated by Kilian Index introduced by Kilian (2009), and log real oil price, which is the nominal WTI crude oil prices deflator by CPI. We use 12 lags of monthly data of these variables between 1981m10 and 2018m12:

$$\begin{bmatrix} \log(\text{S\&P 500 Index}) \\ \textit{oil uncertainty} \\ \text{federal funds rate} \\ \log(\text{CPI}) \\ \log(\text{US oil production}) \\ \log(\text{US real activity}) \\ \log(\text{world oil production}) \\ \text{world real activity} \\ \log(\text{real oil price}) \end{bmatrix}$$

Unlike Bloom (2009) and Kilian (2009), we use all the variables in levels in the estimation as Jurado et al. (2015). One exception is the world real activity, as it is measured as a percentage deviation from trend. As suggested by Sims (1980), Sims et al. (1990) and others, stationarity of the variables is not necessary if the results of interest are dynamic impulse responses, and keeping the levels of variables can shed light on long-run relations between variables. In addition, as discussed in Toda and Yamamoto (1995), including more lags (and in our case, 12 lags) in the VAR can generate consistent estimates even when we use variables in levels. For robustness, we re-estimate the baseline VAR with stationary variables obtained from HP-filtering and log-differencing. The results shown Appendix *A* suggest that qualitatively, oil uncertainty shocks have similar impact.

4.1 Results

We estimate the VAR model with the baseline oil industry uncertainty measure using recursive ordering identification. The results are plotted in Figure 3, where the solid lines display the impulse responses of macroeconomic aggregates and oil market to one-standard deviation shock to oil industry uncertainty, and the shaded area represents +/- one standard error confidence bands. As shown, macroeconomic aggregates respond negatively to an increase in oil uncertainty, with the maximum decline in real activity of 0.4 percent and in the price level of 0.13 percent. The monetary authority responds to the declines by lowering the federal funds rate. The stock market, however, hardly responds to oil uncertainty innovations in the baseline VAR.

In addition, the increase in oil uncertainty negatively affects the US and world oil markets. On impact, US oil production decreases immediately, with a peak effect of 4 percent that occurs around 7 months after initial impact. Interestingly, world real activity and oil production also react to the increase in US oil uncertainty. Specifically, world real activity significantly decreases for almost 10 months. This could be due to the possibility that US oil uncertainty reflects to some extent an increase in aggregate uncertainty, which is shown to have an adverse impact on the worldwide economic activity as in Mumtaz and Theodoridies (2015), among others. World oil production slightly increases, and then returns to the long run trend. As a result, it is not surprising that the real oil price declines, as world oil demand decreases as reflected by lower world real activity, and world oil supply increases (or remains unchanged).

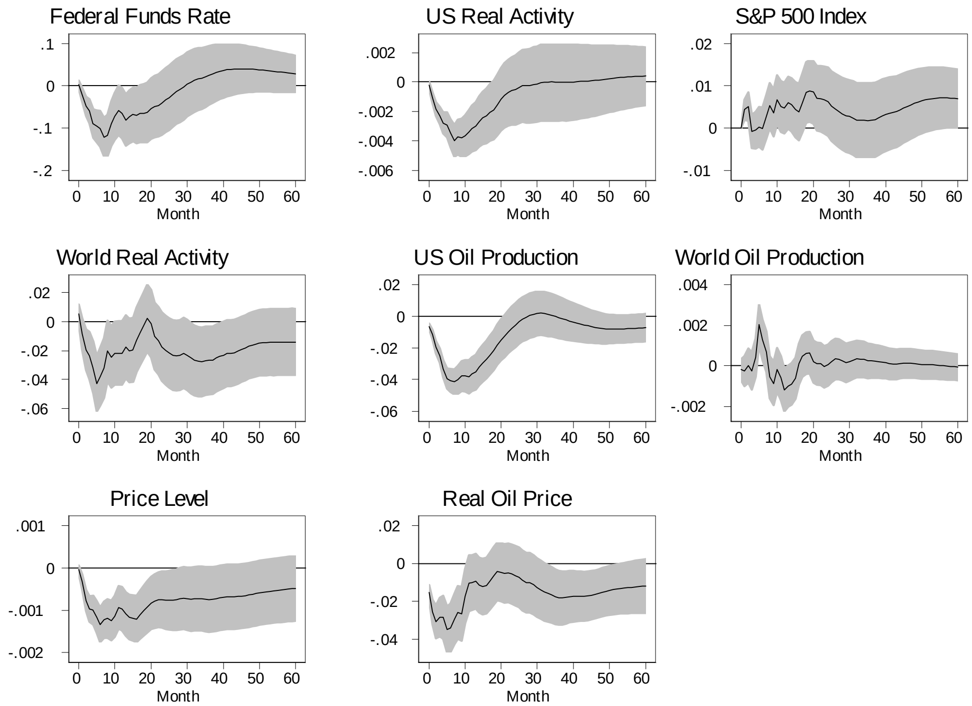

Figure 3 – Impulse response of the federal funds rate, the price level, the stock index, US and world real economic activity, US and world oil production, and real oil price to oil uncertainty shock. The results are estimated from VAR with the baseline oil industry uncertainty. The shock is one standard deviation, and the shared areas represent +/- one standard error confidence bands. The unit in the vertical axis is proportional deviation relative to their respective long run trends. The sample period is 1981m10-2018m12.

In Section 4.2 and Appendix *A*, we report robustness check results estimated from alternative VAR specifications, ordering and variables, including VARs with dispersions as oil uncertainty measure, VARs with uncertainty measures constructed from shorter horizon forecasts, a VAR based on the sample excluding post-2007 periods, a VAR with different identification scheme, and VARs with stationary variables.

4.2 Alternative measures

Measures constructed from dispersions

Our baseline oil uncertainty measure is built on ex-post absolute forecast errors, which is the absolute difference between expected EPS and realized EPS 12 months later. Therefore, we are not able to use it to forecast the near future (within 12 months), as the future EPS is not realized yet. Alternatively, if an oil industry measure is constructed from information that does not rely on future realization of EPS, it can potentially be used as a leading indicator of the dynamic impact of oil uncertainty on macroeconomic aggregates and oil markets. To this end, we use I/B/E/S forecast data to construct two dispersion measures that may also capture uncertainty prevailing in oil industry but which do not depend on forecast errors using equation (4): the median *interquartile range* and the median *standard deviation* of analyst forecasts about a particular oil/gas firm in a month. These uncertainty indices measure the extent of disagreement among forecasters rather than the extent to which forecasts are correct. Consequently, they may serve as forecasting tools in future work. Figures 4 and 5 plot the impulse responses estimated from VARs with these alternative uncertainty measures replacing the baseline measure. As displayed, the impact of uncertainty on macroeconomic and oil sector dynamics is qualitatively similar compared to the baseline.

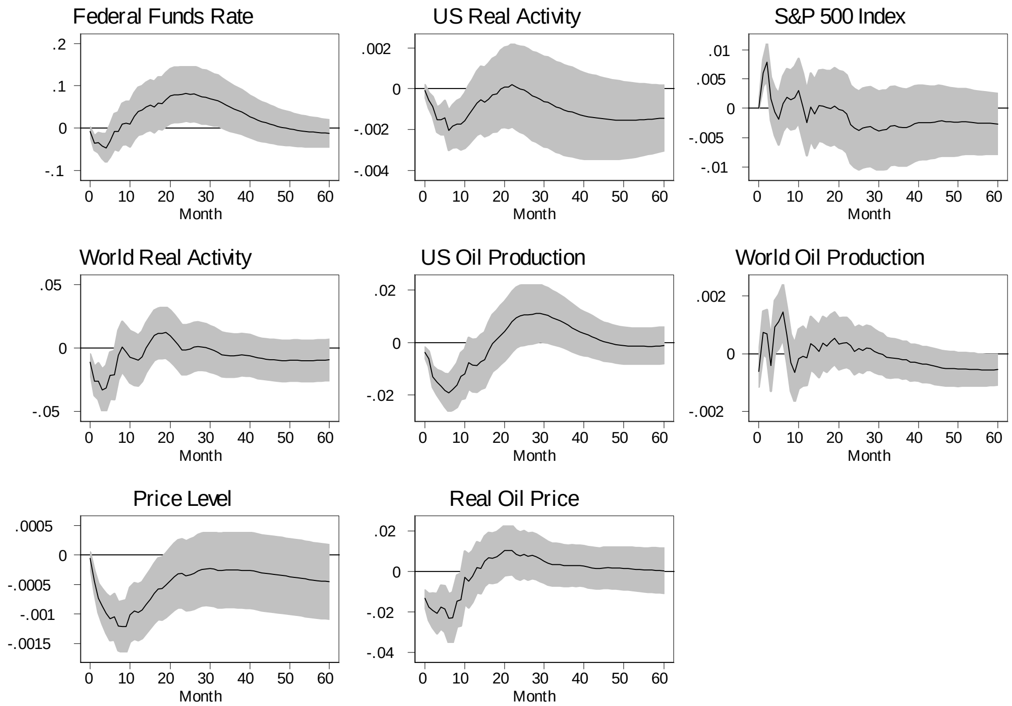

Figure 4 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with interquartile uncertainty measure.

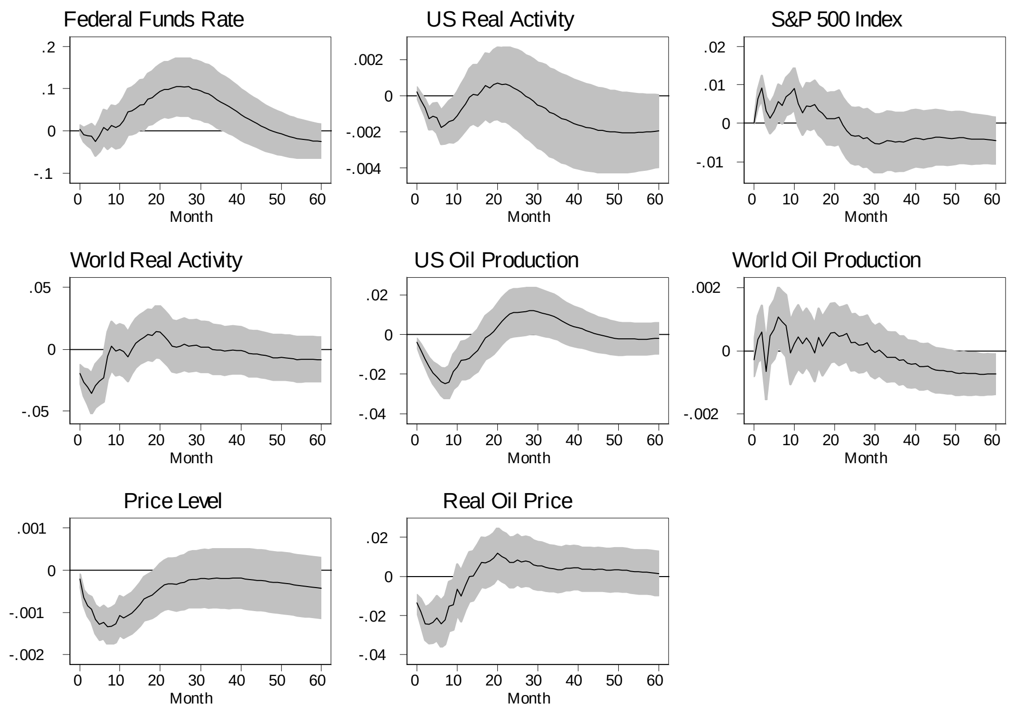

Figure 5 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with standard deviation uncertainty measures.

Measures with various forecast horizons

As the I/B/E/S dataset also provides forecasts with shorter horizons, another way of constructing oil industry uncertainty measures that does not require realized EPS values 12 months later is to use shorter-horizon forecasts. Specifically, we construct three alternative oil uncertainty measures, where the uncertainty about the current month's oil industry is respectively measured as the absolute difference between the forecasted EPS and the realized EPS of oil and gas firms 1 quarter, 2 quarters, and 3 quarters later, leading to 1-, 2-, and 3-quarter ahead oil industry uncertainty measures. The advantage of these measures is that they can be measured earlier, providing earlier insight into future macroeconomic and oil market events.

Figures 6, 7, and 8 show the impulse responses to innovations in these uncertainty measures. Increases in oil uncertainty based on one quarter-ahead absolute forecast errors have a different impact on the federal funds rate, US real activity, and the stock market, which suggests that short-run and long-run oil industry uncertainty may contain different information that in turn has different impact on the macroeconomy. In particular, as the forecast horizon based on which our oil industry uncertainty is constructed increases from 1 quarter to 4 quarters (baseline), the impact of oil uncertainty on US real activity and the price level becomes increasingly significantly negative, as does the impact on the federal funds rate. This suggests that oil uncertainty regarding longer horizons has a stronger and more persistent negative impact on the macroeconomy, as longer-horizon uncertainty may matter more for economic agents' decisions.[14] On the other hand, all the measures generally lead to significant decreases in US oil production and oil prices. This implies that these measures can also serve as forecasting tools, just as the uncertainty measures based on dispersion.[15]

[14] Jurado et al. (2015) have similar findings for aggregate uncertainty. They show that uncertainty based on shorter forecasting horizon (1 and 3 months) generally explains less variation in real economic activity compared to that based on longer forecasting horizon (12 months).

[15] Using quarter-ahead measures as forecasting tools allow us to incorporate more recent information. For example, suppose we are in 2019m8, if we want to forecast the oil production in 2019m9, the last available data point of the year-ahead oil uncertainty measure would be *2018m8* given the definition of our uncertainty measure (the uncertainty in 2018m8 equals to the absolute difference between the forecasted EPS *made in 2018m8* for 2019m8 and the realized EPS in 2019m8), whereas the last available data point of the two-quart-ahead oil uncertainty measure is *2019m2* (the uncertainty in 2019m2 equals to the absolute difference between the forecasted EPS *made in 2019m2* for 2019m8 and the realized EPS in 2019m8).

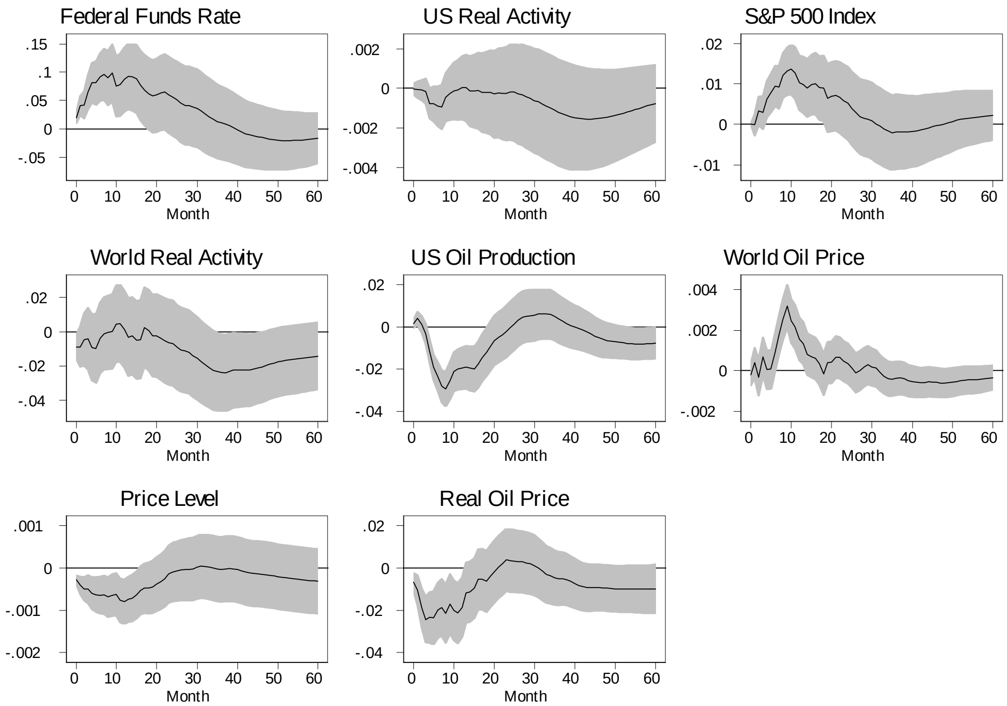

Figure 6 – Impulse responses of macroeconomic variables and oil sector variables from estimation of VAR with 1-quarter ahead oil uncertainty measure.

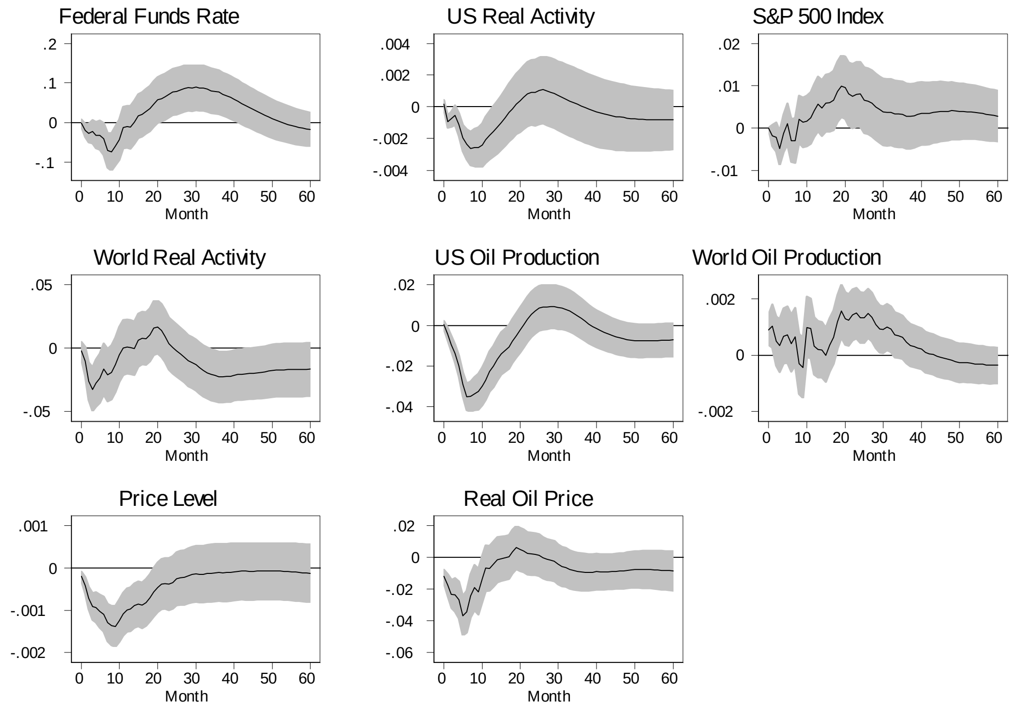

Figure 7 – Impulse responses of macroeconomic variables and oil sector variables from estimation of VAR with 2-quarter ahead oil uncertainty measures.

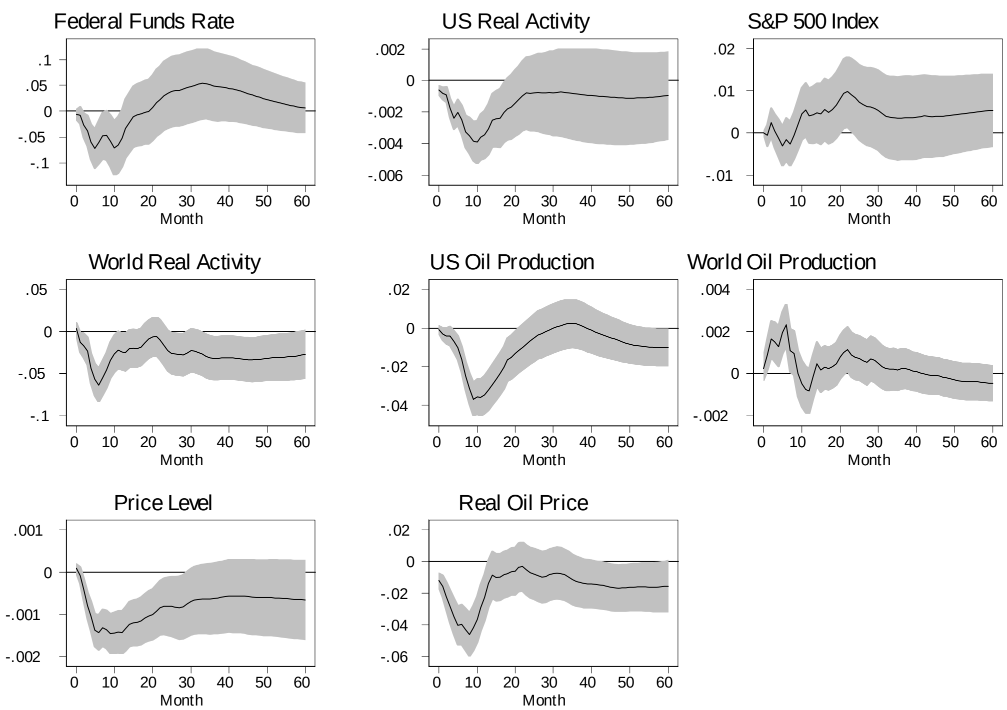

Figure 8 – Impulse responses of macroeconomic variables and oil sector variables from estimation of VAR with 3-quarter ahead oil uncertainty measure.

Before the Great Recession

The 2007 financial crisis may contribute to some of the impact of oil uncertainty shocks, as it is a major demand-driven event in our sample period. To see whether our results are dominated by that event, we re-estimate the VAR system with data spanning between 1981m10 and 2006m12. It is clear from Figure 9 that oil uncertainty still has an adverse impact on US real activity, US oil production and the real oil price. An interesting difference, though, is that the stock market index rises more significantly relative to the baseline estimation. It is likely that, during the historical periods before the Great Recession, an increase in oil uncertainty drives up the stock market on average, as it can lower the oil price and therefore lower production costs in the economy at large.

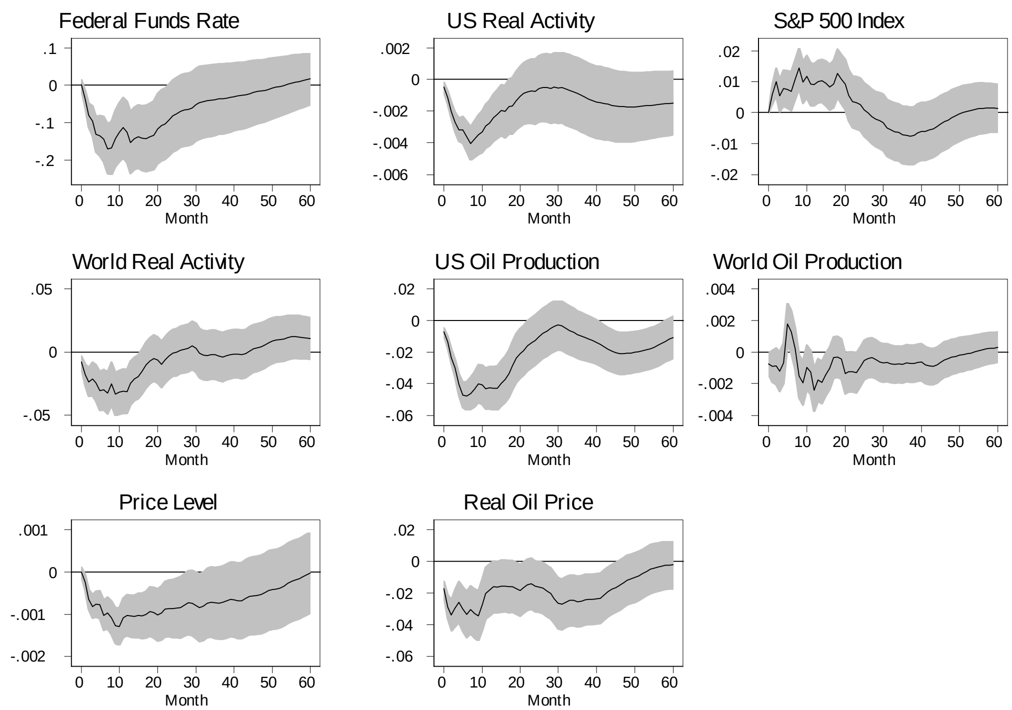

Figure 9 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with sample 1981m10 to 2006m12.

4.3 Oil and Non-Oil Uncertainty

To investigate the difference in the specific information conveyed by uncertainty in the oil sector versus uncertainty in the rest of the economy, we construct a measure that is also based on financial analysts forecasts, but using firms in all the sectors *excluding* the oil and gas industries. We call this *non-oil* uncertainty.

First, we estimate the same VAR system as in the benchmark, except that oil uncertainty is replaced by non-oil uncertainty, to see how this measure behaves on its own. Figure 10 shows that non-oil uncertainty has more significant impact on US and world real activity and the stock market than oil uncertainty. This is not surprising, as non-oil uncertainty captures uncertainty that originates in or is transmitted through non-oil sectors, and therefore it is likely to have more impact on macroeconomic variables, especially if other types of uncertainty shocks occur more often than oil uncertainty shocks during the sample period. On the other hand, non-oil uncertainty does *not* affect world oil production as significantly. Finally, it has a delayed and smaller impact on US oil production and on the real oil price, suggesting that *oil uncertainty* contributes more to the dynamics of US oil production and oil prices.

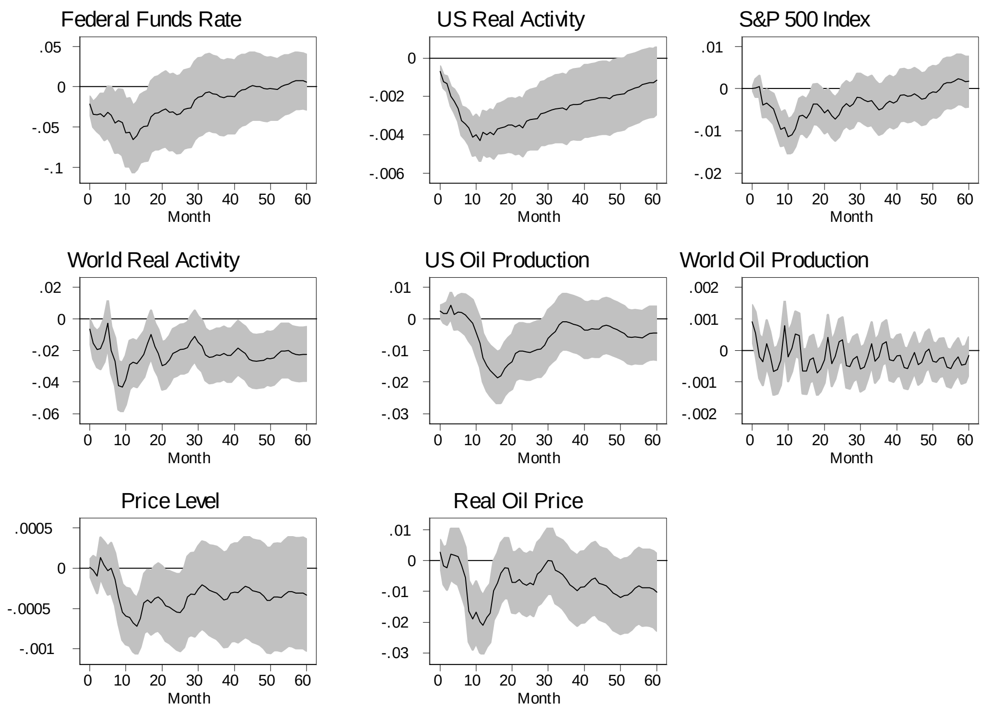

Figure 10 – Impulse response of macroeconomic variables and oil sector variables from estimation of benchmark VAR with non-oil uncertainty.

We also estimate a VAR that simultaneously includes both oil and non-oil uncertainty measures. We order non-oil uncertainty *before* oil uncertainty, which allows us to examine whether uncertainty from the oil sector still matters for macroeconomic and oil sector dynamics *conditional* on the impact of uncertainty in the rest of the economy.

Figure 11 shows the impulse responses to oil and non-oil uncertainty shocks. As before, we find that the macroeconomic impact of non-oil uncertainty is significantly greater than that of oil uncertainty, as shown by the more persistent decrease in US real economic activity. Interestingly, world real activity reacts similarly to oil uncertainty and to non-oil uncertainty in the *short run* (within 6 months), but *more* significantly to non-oil uncertainty in the medium and long run. In addition, world oil production reacts more to oil uncertainty relative to non-oil uncertainty, though the effects quickly become insignificant within one year. Finally, oil uncertainty has more significant impact on US oil production and on the real oil price at almost all horizons. Note that, as oil uncertainty is placed after non-oil uncertainty, the effects of oil uncertainty are measured *after* we have accounted for all the variation in oil uncertainty attributable to shocks to non-oil uncertainty. The finding that US oil production and the oil price still significantly fall due to oil uncertainty shocks reinforces the conclusion that our oil uncertainty measure reflects unique information that has important implications for the oil sector, whereas non-oil uncertainty contains more information about macroeconomic activity.[16]

[16]We obtain a similar conclusion if we change the ordering of oil and non-oil uncertainty. The results are shown in Appendix B.

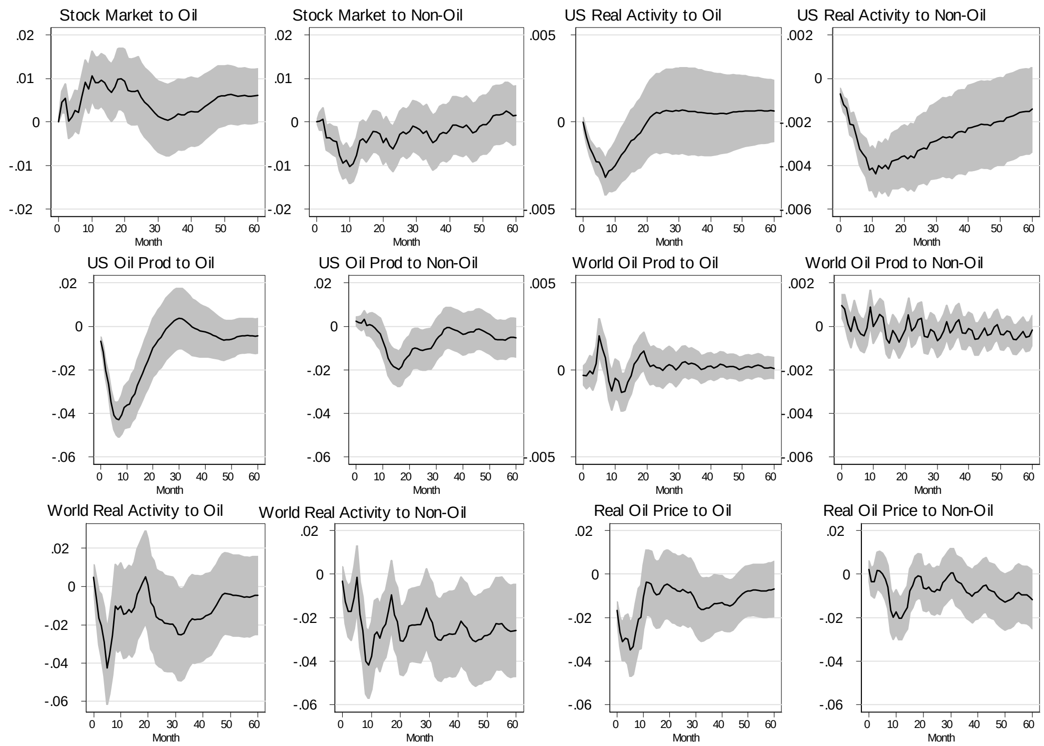

Figure 11 – Impulse responses from estimation of VAR with Oil uncertainty and Non-Oil uncertainty. The responses are estimated from the VAR including both forms of uncertainty, with non-oil uncertainty ordered first. The gray areas are +/- one standard error confidence bands. The unit in the vertical axis is proportional deviation relative to long run trend.

Similar findings are suggested by the results of variance error decomposition over horizons of 3, 12, 36, and 60 months for macroeconomic and oil market variables due to oil and non-oil uncertainty. Table 1 shows that, the baseline oil uncertainty explains more variation in US and world oil production in almost all horizons compared to non-oil uncertainty, whereas the non-oil uncertainty has stronger impact on US real economic activity relative to non-oil uncertainty. Oil uncertainty has a similar effect on world activity in the short run, but less significant effect in the long run. Finally, the impact of oil uncertainty on oil prices is more significant than that due to non-oil uncertainty in all horizons.

	US oil prod		World oil prod		US activity		World Activity		Oil Price	
Horizon	Oil	Non-Oil	Oil	Non-Oil	Oil	Non-Oil	Oil	Non-Oil	Oil	Non-Oil
$h = 3$	11.11	0.21	0.09	0.64	4.10	5.21	0.52	0.72	7.59	0.11
$h = 12$	33.86	0.47	1.85	0.63	7.94	14.27	3.08	3.88	9.45	1.79
$h = 36$	20.00	4.07	2.46	1.09	2.75	13.81	3.37	5.96	7.40	2.18
$h = 60$	17.41	3.81	2.26	1.38	2.26	12.98	3.00	7.76	6.93	2.75

Table 1 – Macroeconomic Variables and Oil Market Variables Forecast Variance Due to Oil Uncertainty and Non-Oil Uncertainty (in percent)

4.4 Oil Uncertainty and Aggregate Uncertainty

The previous exercises show that our oil uncertainty measure can contain information exclusive to the oil sector that has an impact on oil markets. To further shed light on whether oil uncertainty originates from the oil sector or instead reflects the transmission of aggregate uncertainty to the oil market, we conduct additional exercises in this section with aggregate uncertainty and oil uncertainty.

First, we re-estimate the same VAR system as in the benchmark, but replace oil uncertainty with aggregate uncertainty. The aggregate uncertainty measure is the macroeconomic uncertainty measure created in Jurado et al. (2015). As in the literature, we find that aggregate uncertainty lowers economic activity, as shown in Figure 12. A new result is that it also lowers US oil production, world oil production, world economic activity, and oil prices.

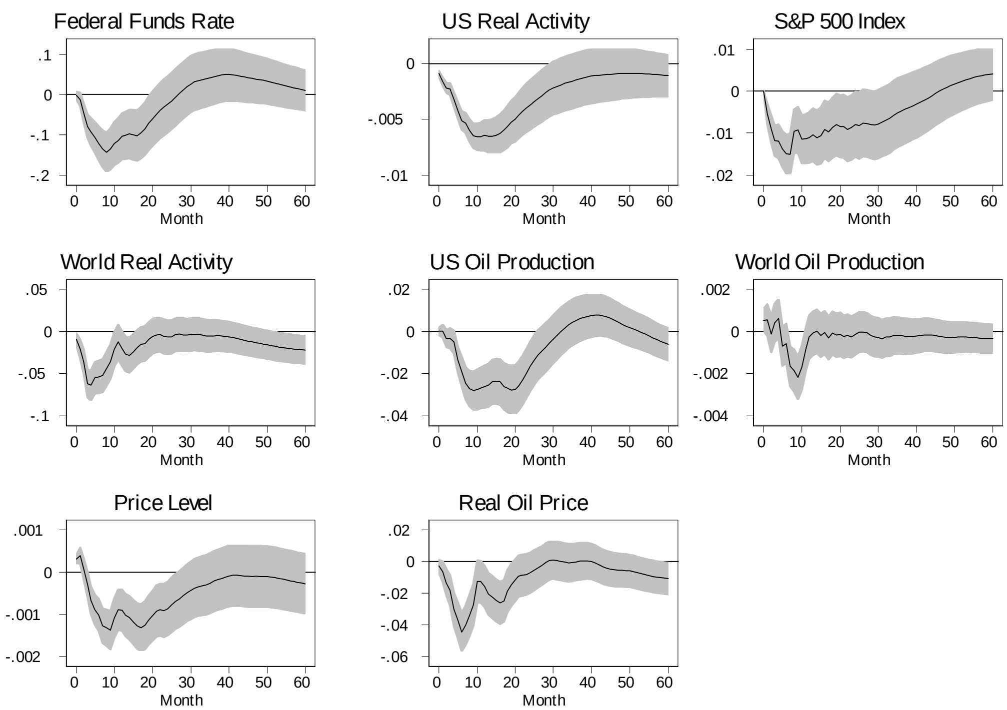

Figure 12 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with aggregate uncertainty.

To investigate whether the impact of oil uncertainty is still significant *conditional* on aggregate uncertainty, or equivalently, whether our oil uncertainty measure contains independent information compared to aggregate uncertainty, we estimate an alternative VAR system by including both oil uncertainty and aggregate uncertainty measures while keeping other variables the same. Figure 13 displays the impulse responses of macroeconomic variables and oil sector variables to a one standard deviation innovation in oil uncertainty and aggregate uncertainty. Like the baseline results analyzed before, macroeconomic variables, such as US and world real activity, still decline following higher oil uncertainty, even though the impact of aggregate uncertainty is already accounted for. The magnitude of impact is smaller, though. The stock market response to oil uncertainty is similar to that estimated from sample up to 2006, suggesting that, generally, uncertainty specific to oil industry has a *positive* influence on the stock market.

Turning to the oil market, oil uncertainty still has more significant impact on US oil production, generating a 1 percent greater variation at the peak level. Interestingly, for world oil production, oil uncertainty tends to raise it, whereas aggregate uncertainty lowers it. Finally, oil uncertainty still lowers oil prices immediately. Relative to aggregate uncertainty, the magnitude of the effect of oil uncertainty is bigger in the short run, but smaller in the medium and long run. These results imply that, at least to some extent, uncertainty originateing from the oil sector alone can serve as a driving force of macroeconomic and oil sector dynamics, as the effects of uncertainty spreading from the aggregate economy to oil sector have already been captured by the inclusion of aggregate uncertainty in the VAR. This conclusion can also be inferred from a Granger causality test, which is performed on the previously described VAR system. The p-value for the statistical test that all coefficients of lags of the aggregate uncertainty in the oil uncertainty equation are zero is 0.38. This result suggests that we fail to reject the null hypothesis that lags of aggregate uncertainty do not affect oil uncertainty, or equivalently, we can conclude that aggregate uncertainty does not Granger-cause oil uncertainty.[17]

It is worth underlining that in Figure 13 the magnitude of the impact of oil uncertainty on US oil production is similar to that in Figure 3, when aggregate uncertainty is not in the VAR. This indicates that, while aggregate uncertainty does have an impact on US oil production, the impact of oil uncertainty on the domestic oil market is independent and is preserved when we condition on aggregate uncertainty. On the other hand, regarding the oil price, the impact of oil uncertainty becomes less significant when conditioning on the impact of aggregate uncertainty. This suggests that oil uncertainty contains some of the same the information as aggregate uncertainty in accounting for the world oil market dynamics.

[17] In fact, for all the variables in the VAR, the null hypothesis that they Granger-cause oil uncertainty can be rejected.

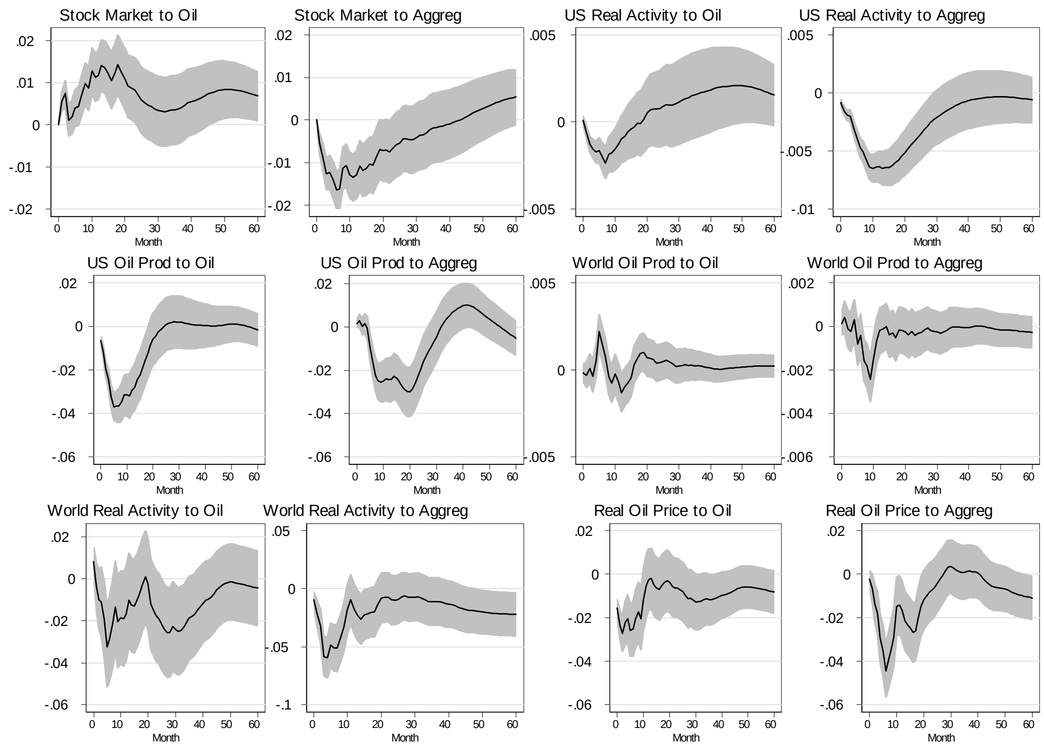

Figure 13 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with oil uncertainty and aggregate uncertainty.

We also calculate the VAR forecast error variance decomposition for oil uncertainty and aggregate uncertainty, and show the results in Table 2. Similarly to our previous analysis, oil uncertainty explains more variation in US oil production and in the oil price (in the short run), whereas aggregate uncertainty explains more variation in US and World real economic activity, and in the oil price (in the medium and long run). The two aggregate measures explain similar variation of world oil production. These results show that oil uncertainty has important implications for the oil market, and is therefore especially informative for studying the behavior of variables closely linked with the US oil sector.

	US oil prod		World oil prod		US activity		World Activity		Oil Price	
Horizon	Oil	Aggr	Oil	Aggr	Oil	Aggr	Oil	Aggr	Oil	Aggr
$h = 3$	10.69	0.21	0.05	0.09	3.05	10.24	0.26	2.32	6.21	0.93
$h = 12$	28.55	9.01	2.01	3.38	4.95	39.00	2.38	10.16	6.74	11.55
$h = 36$	19.27	17.40	2.69	2.54	1.83	25.53	4.08	8.02	6.22	11.80
$h = 60$	16.26	15.83	2.44	2.28	3.73	19.33	3.43	7.94	6.02	10.09

Table 2 – Macroeconomic Variables and Oil Sector Variables

Forecast Variance Due to Oil Uncertainty and Aggregate Uncertainty

(in percent)

We also obtain the conclusion that oil uncertainty contains exclusive information from an alternative exercise, i.e., an estimation of a VAR with both non-oil uncertainty and aggregate uncertainty, as well as other same macroeconomic and oil market variables. The impulse responses are displayed in Figure 14. The results reassure us that information contained in our oil uncertainty measures is not contained in non-oil sector uncertainty measures, as US oil production and the oil price barely respond to non-oil uncertainty shocks.

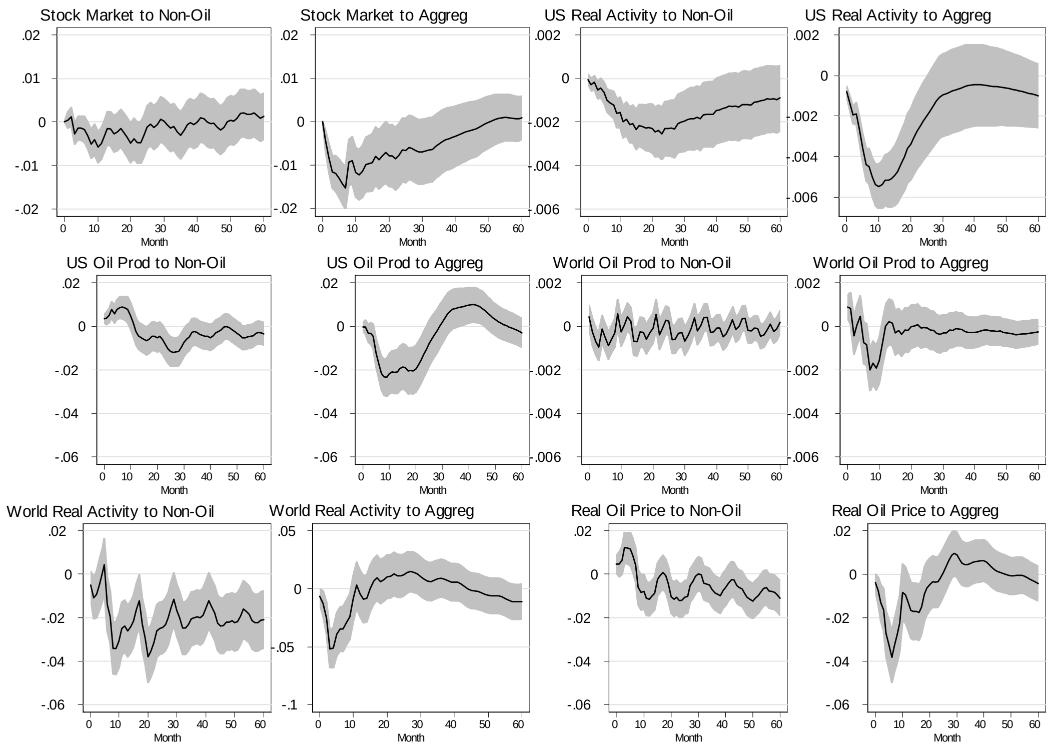

Figure 14 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with non-oil uncertainty and aggregate uncertainty.

4.5 Relative Oil Uncertainty

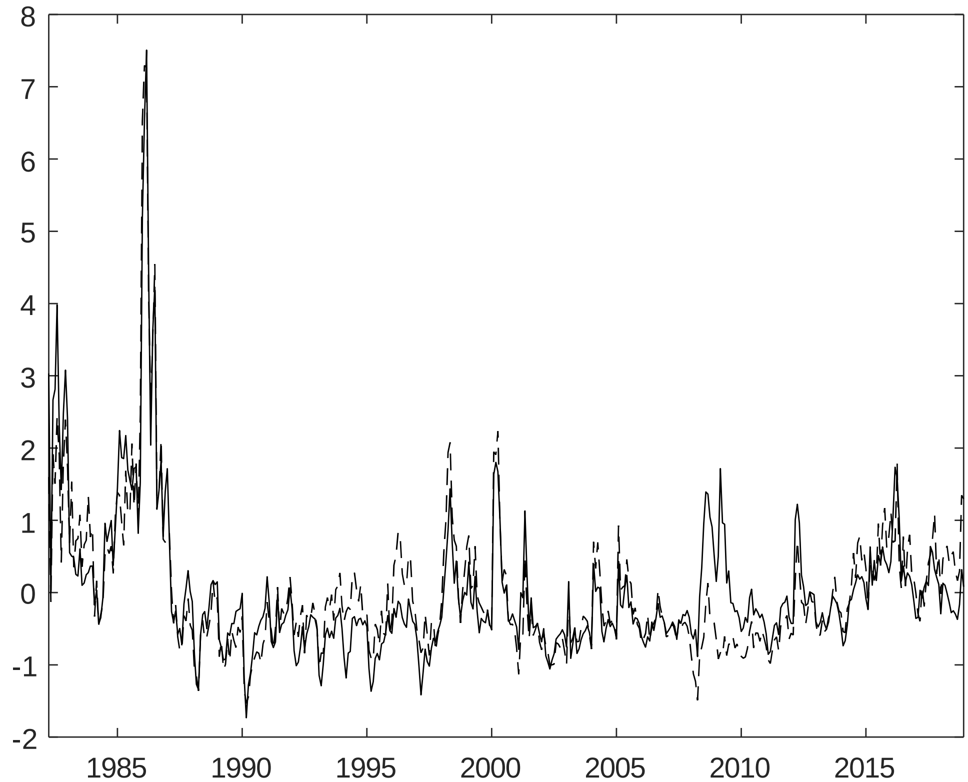

Figure 15 – Comparison of different uncertainty measures. The thick line is our baseline oil industry uncertainty measure and the dashed line is the relative measure i.e. oil industry uncertainty divided by uncertainty outside the oil market. The sample period is 1982m3-2018m12.

It would be insightful to explore the impact of uncertainty during the periods when oil sector experiences *relatively* higher uncertainty than the non-oil sector, which may reflect to some extent the exclusive impact on the macroeconomy and the oil sector of exceptional increases in uncertainty *particularly* in the oil sector. To see this, we construct a *relative* oil uncertainty measure, which is the ratio of oil uncertainty to non-oil uncertainty, capturing periods of uncertainty that are particularly large for the oil sector. The time series of this relative uncertainty measure is shown in Figure 15. Notably, oil relative uncertainty coincides with baseline uncertainty most of the time, but during some crises that are known not to originate from the oil sector, such as the Great Recession, oil relative uncertainty is lower than the baseline oil uncertainty measure.

Intuitively, this measure should have similar impact compared to the oil uncertainty measure, and especially capture the information contained in the oil sector but not in the non-oil sector. This intuition is verified by the results shown in Figure 16, which show that an increase in uncertainty in oil sector relative to non-oil sector lowers US oil production, raises world oil production, lowers US and world economic activity, and lowers oil prices. Again, notably, the measure of relative uncertainty in the oil sector that is not in the non oil sector has *positive* impact on the stock market. These findings are consistent with our previous analysis on oil uncertainty.

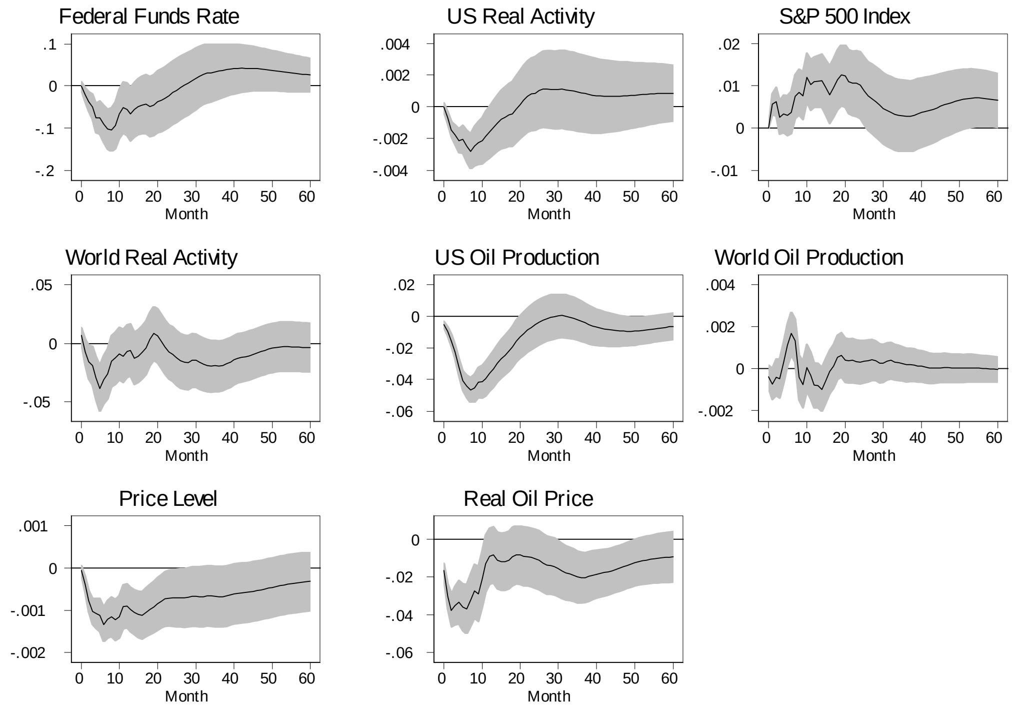

Figure 15 – Impulse response of macroeconomic variables and oil sector variables from estimation of VAR with oil relative uncertainty.

5 Concluding Remarks

We construct several measures of oil industry uncertainty based on a large number of financial analysts forecasts on oil and gas industry. The measure we develop turns out to have unique implications for oil market dynamics. Variants of the baseline measure are also suitable for forecasting purposes.

Given the empirical findings in this paper that underline the importance of oil uncertainty shocks, it is interesting in the future to develop a structural quantitative model where uncertainty originates in the oil industry and spreads to the rest of the economy in the ways we uncover. This could be especially useful for policy analysis if incorporated in a monetary DSGE framework.

6 References

Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims. 2013. Uncertainty and Economic Activity: Evidence from Business Survey Data. American Economic Journal: Macroeconomics, 5(2): 217-49.

Bakas, Dimitiros, Athanasios Triantafyllou. 2018 The Impact of Uncertainty Shocks on the Volatility of Commodity Prices. Journal of International Money and Finance, 87: 96-111.

Basu, Susanto, and Brent Bundick. 2017. Uncertainty Shocks in a Model of Effective Demand. Econometrica, 85(3): 937-958.

Bloom, Nicholas. 2009. The Impact of Uncertainty Shocks. Econometrica 77: 623-685.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino. 2018. Measuring Uncertainty and Its Impact on the Economy. Review of Economics and Statistics, Forthcoming.

Elder, John, and Apostolos Serletis. 2010. Oil Price Uncertainty. Journal of Money, Credit, and Banking 42(6): 1137-1159.

Jo, Soojin. 2014. The Effects of Oil Price Uncertainty on Global Real Economic Activity. Journal of Money, Credit, and Banking, 46(6): 1113-1135.

Joëts, Marc, Valérie Mignon, and Tovonony Razafindrabe. 2017. Does the Volatility of Commodity Prices Reflect Macroeconomic Uncertainty? Energy Economics 68: 313-326.

Jurado, Kyle, Sydney C. Ludvigson and Serena Ng. 2015. Measuring Uncertainty. American Economic Review 105(3): 1177-1216.

Ma, Xiaohan and Roberto Samaniego. 2019. Deconstructing Uncertainty. European Economic Review 119, 22-41.

Maghyereh, Aktham I., Basel Awartani, and Elie Bouri. 2016. The Directional Volatility Connectedness between Crude Oil and Equity Markets: New Evidence from Implied Volatility Indexes. Energy Economics 57: 78-93.

Mumtaz, Haroon, and Konstantinos Theodoridies. 2015. The International Transmission of Volatility Shocks: An Empirical Analysis. Journal of the European Economic Association 13(3): 512-533.

Punzi Maria Teresa. 2019. The Impact of Energy Price Uncertainty on Macroeconomic Variables. Energy Policy 1306-1319.

Shin, Minchul and Molin Zhong. 2018. A New Approach to Identifying the Real Effects of Uncertainty Shocks. Forthcoming, Journal of Business and Economic Statistics.

Sims, Christopher A. 1980. Macroeconomics and Reality. Econometrica 48(1): 1-48.

Sims, Christopher A., James Stock, and Mark Watson. 1990. Inference in Linear Time Series Models with Some Unit Roots. Econometrica 58(1): 113-144.

Sims, Christopher A., and Tao Zha. 1999. Error Bands for Impulse Response Functions. Econometrica 67(5): 1113-1155

So, Eric C. 2013. A New Approach to Predicting Analyst Forecast Errors: Do Investors Overweight Analyst Forecasts? Journal of Financial Economics 108(3): 615-640.

Toda, Hiro Y. and Taku Yamamoto. 1995. Statistical inference in vector autoregressions with possibly integrated processes. Journal of Econometrics 66(1-2): 225-250.

Van Robays, Ine. 2016. Macroeconomic Uncertainty and Oil Price Volatility. Oxford Bulletin of Economics and Statistics 78(5): 671-693.

Yin, Libo and Jiabao Feng. 2019. Oil Market Uncertainty and International Business Cycle Dynamics. Energy Economics, 81: 728-740.

Appendix for online publication only

Appendix A: Robustness for baseline VAR results

We explore the robustness of the baseline results through three additional VAR estimations.

First, in the benchmark, oil uncertainty is placed as the second variable, assuming that innovations in oil uncertainty have contemporaneous effects on all the macroeconomic variables and oil market variables in the VAR, except for the stock market index. To investigate whether the baseline results are driven by this ordering, we re-estimate a VAR with the same variables, but place oil uncertainty as the last variable. In this way, it is assumed that oil uncertainty shock does not affect other variables in the VAR in the same period. The estimation results are shown in Figure $A1$. Compared with Figure 3, a difference is that oil uncertainty shock does not affect US oil production and oil price on impact. This is not surprising by the assumption of the timing of the shock. However, the main results that higher oil uncertainty depresses macroeconomic and oil market activity remain with this new ordering.

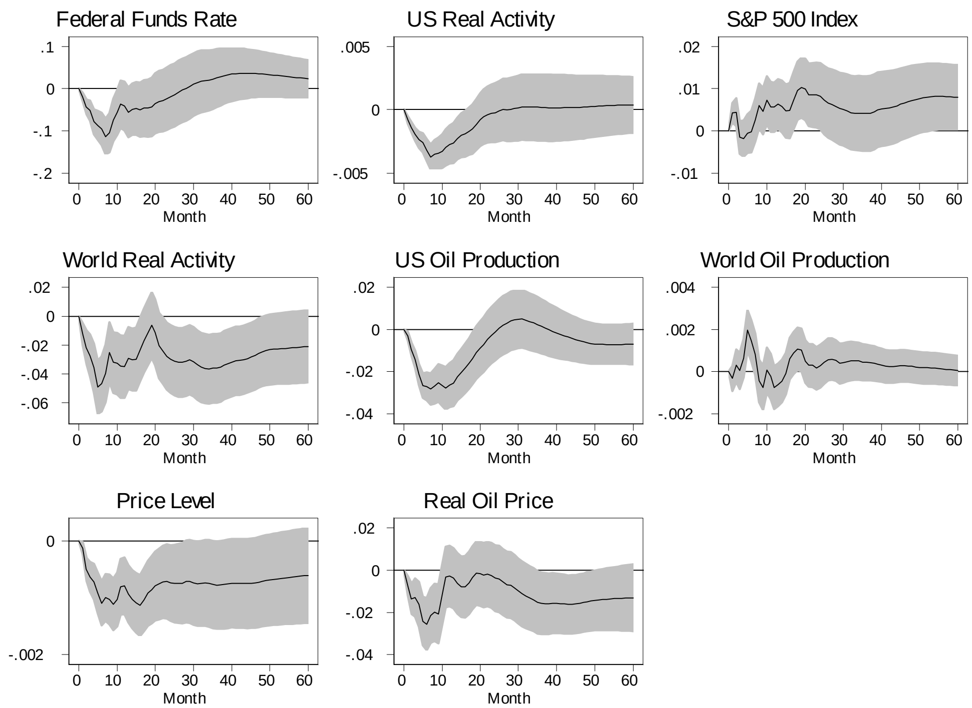

Figure $A1$ – Impulse response of the federal funds rate, the price level, the stock index, US and world real economic activity, US and world oil production, and real oil price to oil uncertainty shock. The results are obtained from estimation of benchmark VAR, with oil uncertainty measured using the 12-month forecast error and placed last in the VAR The shock is one standard deviation, and the gray areas are +/- one standard error confidence bands. The unit in the vertical axis is proportional deviation relative to their respective long run trends. The sample period is 1981m10 to 2018m12.

Second, we use the log of some variables in the baseline VAR for estimation to not exclude possible medium and long run dynamics between variables. To see whether the results are robust to stationarity of variables, we convert whichever variable that is not stationary to stationary before estimating the VAR with the same variables. In particular, we use two method for stationarity purpose: i) HP-filter as in Bloom (2009) and ii) log-differencing as in Kilian (2009). The estimation results with the new variables are shown in Figure $2A$ and $3A$ respectively. Qualitatively, the results are very similar with the benchmark results.

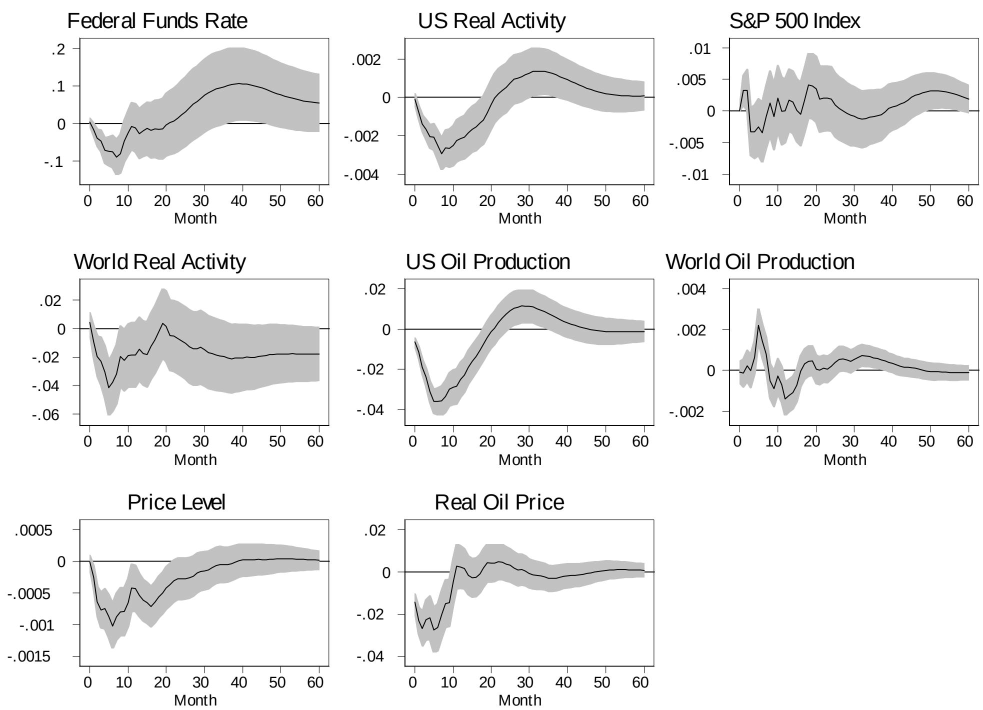

Figure $A2$ – Impulse response of macroeconomic variables and oil sector variables from estimation of benchmark VAR, with oil uncertainty measured using the 12-month forecast error and HP-filtered variables in the VAR.

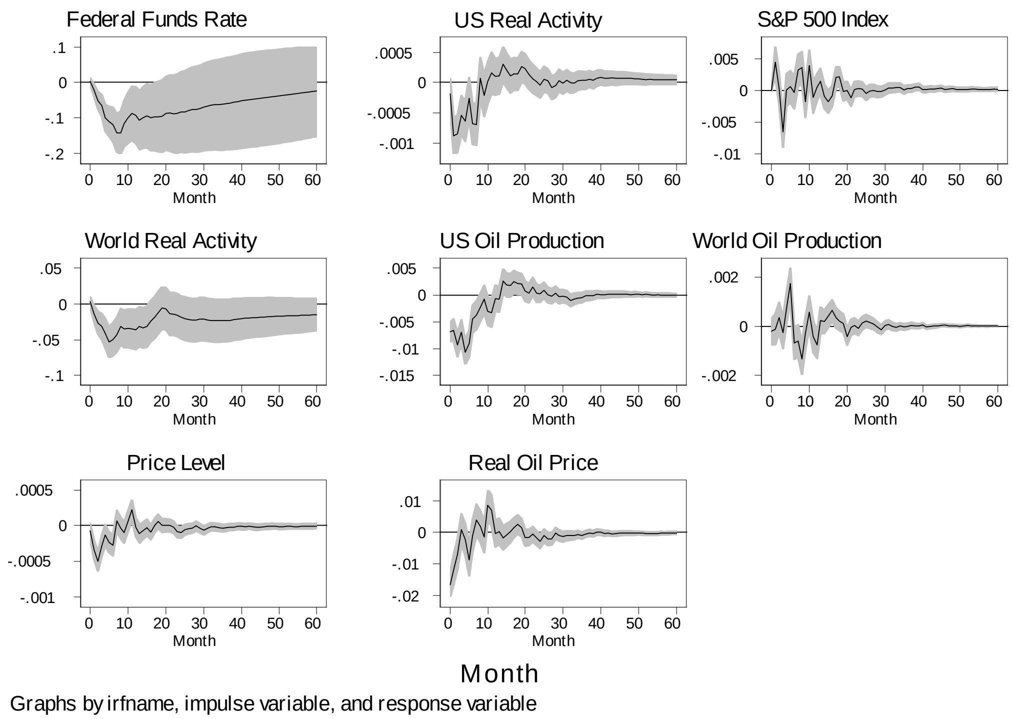

Figure $A3$ – Impulse response of macroeconomic variables and oil sector variables from estimation of benchmark VAR, with oil uncertainty measured using the 12-month forecast error and log-differenced variables in the VAR.

Appendix B: Robustness for VAR results with Oil and Non-Oil Uncertainty

In the main text, we estimate the VAR with both oil and non-oil uncertainty by ordering non-oil uncertainty before oil uncertainty, and find that oil uncertainty contributes more to the fluctuations in US oil production and oil price. In this robustness check, we explore whether the conclusion is related to the ordering of the oil and non-oil uncertainty by estimating an alternative VAR that orders non-oil uncertainty after oil uncertainty. This implies that the effects of non-oil uncertainty on macroeconomic variables and oil market variables are measured after we have removed all the variation in non-oil uncertainty that is attributable to shocks to oil uncertainty.

The impulse responses are shown in Figure $B1$. As before, increases in oil uncertainty and non-oil uncertainty still generate significant impact on US real activity. The ordering is also irrelevant when analyzing the contribution of these two uncertainty measures to the dynamics of oil market.

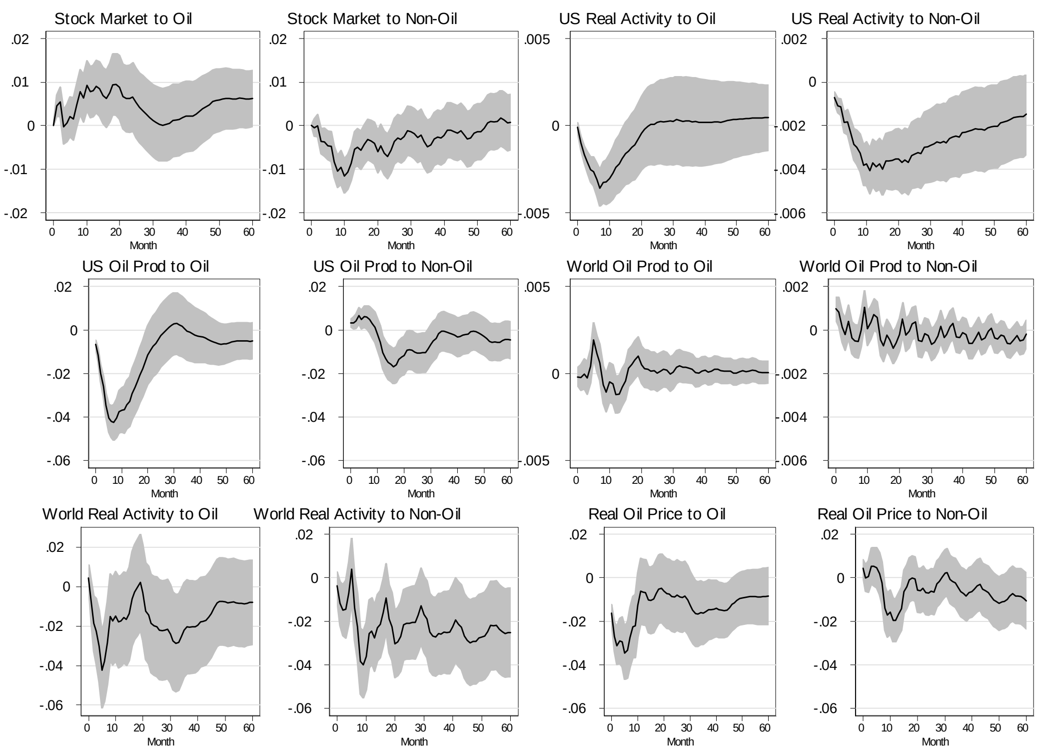

Figure $B1$ – Impulse responses from estimation of VAR with Oil uncertainty and Non-Oil uncertainty. The solid lines show the responses estimated from the VAR using the baseline data sample and including both forms of uncertainty, with oil uncertainty ordered first.