Oil News Shocks and the U.S. Stock Market

Zeina Alsalman[*] Ana María Herrera[†] Sandeep Kumar Rangaraju[‡]

Abstract

We study the effect of oil news shocks on U.S. aggregate and industry-level stock returns. Using a Proxy-VAR and a sample from January 1973 to December 2019, we find no significant effect on aggregate stock price index on impact, but a persistent and significant drop at longer horizons. An important degree of heterogeneity is found in the industry-level responses: stock returns for precious metals, coal, petroleum and natural gas, and utilities increase significantly and persistently after the shock, whereas consumer goods, rubber and plastic, automobiles, and trucks fall shortly. Moreover, our estimates indicate that oil news shocks pose a risk for IT sectors whose returns exhibit losses. When we extend the sample to December 2022, we uncover a crucial change in the dynamics of the oil surprise measure: it is contaminated by lags of the oil price. We illustrate how using the oil surprise as an instrumental variable in the extended sample produces puzzling responses of industry-level stock returns, whereas a purged measure of the oil surprise leads to more stable estimates.

JEL Classification: C32, E32
Keywords: Oil prices, U.S. stock returns, OPEC, proxy-VAR.

[*]Department of Economics, 413 Elliott Hall, School of Business Administration, Oakland University, Rochester, MI 48309, USA; e-mail: alsalman@aokland.edu

[†]Department of Economics, Gatton College of Business and Economics, University of Kentucky, Lexington 40206-0034; phone: (859) 257-1119; e-mail: amherrera@uky.edu

[‡]Department of Economics, Goddard School of Business & Economics, Weber State University, Ogden 84408-3801; e-mail: srangaraju@weber.edu

1 Introduction

In October 2022, as OPEC+ announced the first production cut since the onset of the pandemic in March 2020, President Biden urged U.S. oil companies to use their "record-breaking profits to increase production and refining" Hooper (2022). Some analysts argued that policies put in place by the Biden's Administration –such as drawing down the Strategic Petroleum Reserve– amounted to manipulating prices as OPEC had been accused of doing when cutting production (Editorial (2022)). This discussion came against the backdrop of high inflation rates, the ongoing war in Ukraine, global supply-chain disruptions and a looming global recession, which bolstered concerns regarding the possible negative effects of high oil prices on economic activity. How do changes in oil market expectations around these announcements affect U.S. stock returns?

Recent work by Känzig (2021) pioneered the use of high-frequency variation in oil price futures around OPEC announcements as an external instrument to identify the effect of oil price shocks. He interprets these shocks as measuring news regarding future oil supply and estimates these shocks to have a negative impact on U.S. aggregate economic activity and on some stock indices. An advantage of using high-frequency identification is that it can mitigate issues that arise when oil price shocks are identified via short-run or sign restrictions (see (Herrera and Rangaraju, 2020)). Yet, as the instrument-based identification (proxy-VAR or SVAR-IV) approach has gained momentum, concerns regarding the role played by information frictions have emerged. For instance, in the context of monetary policy shocks, it has been shown that failing to account for information frictions can lead the researcher to estimate dynamic responses that "confound the effects of a monetary policy shock with the endogenous response of the central bank to changes in the economy" (Miranda-Agrippino and Ricco (2021)). Consequently, alternative choices of instruments, empirical specification and sample period can result in important differences in the estimated responses (Miranda-Agrippino and Ricco (2021) and Plagborg-Møller and Wolf (2021)). In the context of oil price shocks, Degasperi (2021) argues that surprises in oil futures prices around OPEC announcements capture not only changes in the market's expectations regarding oil supply, but they also reflect changes in expectations regarding oil demand. He finds evidence that information frictions also play a role in the identification of oil price shocks in a proxy-VAR setting.

In this paper, we take a fresh look at the relationship between oil market expectations and industry-level stock returns. Our contribution to the literature is threefold. First, we build on the work by Känzig (2021) and Degasperi (2021) by exploring how heterogeneity in the response of industry-level portfolios to oil market expectations contributes to our understanding of the transmission mechanism of oil news shocks. We find that, in some industries, the response of stock returns resembles that to uncertainty shocks in that news of higher future oil prices lead to losses, possibly due to declines in investment, employment, and production.[1] In particular, oil news shocks lead to losses for shares in computer software and hardware, as well as in electronic equipment. On the contrary, precious metals experience gains, a result that is consistent with such goods playing a "safe-haven" role when uncertainty hits.

Second, for industries that have been traditionally thought to be sensitive to oil price shocks, we find that oil news shocks operate in a similar fashion as oil-specific demand shocks (Kilian and Park, 2009): they result in losses for industries that are intensive in the use of energy in production (e.g., rubber and plastics) or consumption (e.g., automobiles and shipping containers), and gains for the energy sector (e.g., petroleum and gas).

Last but not least, we illustrate how ignoring informational frictions may lead to estimates of the impulse responses that are sensitive to changes in the sample period. This problem is found to be particularly severe when we extend the sample to include post-2019 data, a period of time when the oil surprise instrument is contaminated by lagged oil prices. That is, spot oil prices are informative for future oil surprises. We argue, that information asymmetries between market participants and the OPEC Conference confounds the effect of oil price shocks and the endogenous response of OPEC to conditions of the oil market. We then demonstrate how, using the residual of a regression of the oil surprises on the oil market variables (i.e., "purging" the instrument) as an instrument, allows the researcher to recover the dynamic causal response of stock returns to oil price shocks.

This paper is organized as follows. The following section describes the data and section 3 discusses the empirical strategy. Estimation results for the pre-COVID sample and the extended sample are reported in sections 4 and 5, respectively. The last section presents our conclusions.

[1] Work by Bloom (2009) and many others find that uncertainty shocks lead to decline in investment, hiring, and productivity through a precautionary channel. Recently, Bretscher *et al.* (2023) uncovers evidence of a risk premium channel in the response of stock returns.

2 Data

To study the effect of oil news shocks on stock returns we first use monthly data spanning the period between January 1973 and December 2019. We then extend the data to December 2022 and explore the stability of our results to including the years of the COVID-19 pandemic and the onset of the war in Ukraine. Our data include the standard four oil market variables – the real price of oil, world oil production, world oil inventories, global real economic activity– as well as aggregate and industry-level U.S. real stock returns. In what follows we describe the sources and construction of the time series used in this paper.

2.1 Oil market and stock return variables

For ease of comparison with Känzig (2021), our baseline regressions use a similar specification where the variables, unless otherwise specified, are measured in logarithms. We obtain data on world oil production (including lease condensate production) from the Energy Information Administration (EIA). The data is measured in thousands of barrels per day and enters the baseline specification in logs.

The real price of oil is computed using the refiners' acquisition cost (RAC) of imported crude oil deflated by the U.S. consumer price index (CPI). This allows us to obtain a consistent time series of oil prices starting in the early 1970s. However, we note that our estimation results are robust to using the WTI as in Känzig. The data on the RAC is collected from the EIA starting from January 1974 and extrapolated to January 1973 as in Barsky and Kilian (2002). The CPI series is obtained from the FRED database.

Regarding inventories, we collect data on total US crude oil inventories, US petroleum stocks, and OECD petroleum stocks from EIA. Given that a measure of crude oil inventories for non-U.S. countries is unavailable, we follow Hamilton (2009) and Kilian and Murphy (2014) in using data for total U.S. crude oil inventories as a proxy for global oil inventories and scaling the total U.S. crude oil inventories by the ratio of OECD petroleum stocks to U.S. petroleum stocks.[2] The scaled ratio in our sample spans between 2.23 and 2.61, which is identical to the scale factor in Herrera and Rangaraju (2020) for a sample spanning January 1973 to December 2016 and very close to the range of 2.23 to 2.59 in Kilian and Murphy (2014) for a sample ending in

[2] Given the lack of consistency in the OECD petroleum stocks series prior to December 1987, we apply backward extrapolation to derive the rate of change in OECD inventories from the rate of growth in U.S. petroleum inventories Kilian and Murphy (2014)).

August 2009.

As a measure of global economic activity we use the world industrial production (WIP) computed for the OECD together with the prime economies of Brazil, China, India, Indonesia, Russian Federation, and South Africa. The WIP is based on the real output in manufacturing, mining, along with the electric and gas industries.[3] However, we show in the robustness checks that specifications using the rate of growth of the WIP or an an alternative measure of global economic activity, Kilian's real economic activity index (Kilian, 2009) lead to the same conclusions.[4]

The data on monthly nominal stock returns, at the aggregate and industry level, are available from Kenneth French's data library.[5] Aggregate stock returns are measured as the excess return on the market, which is constructed by subtracting the one-month Treasury bill rate from the Center for Research in Security Prices (CRSP) value-weighted portfolio return. This latter is computed as the value-weighted return on all NYSE, AMEX, and NASDAQ portfolios from the CRSP. The 49 industry-level portfolios are constructed by assigning each NYSE, AMEX, and NASDAQ stock to an industry portfolio based on its four-digit Standard Industry Classification (SIC) code provided by Compustat, or by CRSP if Compustat is unavailable. Real stock returns are calculated by subtracting the CPI inflation from the log of nominal stock returns.[6]

2.2 The instrument

To identify the oil news shock, we follow Känzig (2021) and compute an oil surprise series using daily data on oil futures prices. Specifically, we gather OPEC press releases for the period 1983-2022 (a total of 150 announcements). Since 2002, press releases are available from the official OPEC webpage.[7] Prior to 2002, we collect the announcement dates from OPEC official resolutions and press releases OPEC (1990), as well as from Bloomberg news (see Table A.2 of the online appendix).

Oil futures prices starting in April 1984, when oil futures started trading, were collected from Bloomberg. Then, we computed oil surprises as the log difference between the oil futures price on the day of the announcement, $F_{t,d}^{h}$ and the price of the final trading day prior to the

[3]The index is available at: https://sites.google.com/site/cjsbaumeister/datasets

[4]See Kilian and Zhou (2018), Kilian (2019), Funashima (2020), and Nonejad (2020) for a thorough evaluation of the Kilian index. A revised version of the index (Kilian, 2019), is updated constantly and is available from the St. Louis Federal Reserve Bank Economic Data (FRED) at https://fred.stlouisfed.org/series/IGREA

[5]The data is available at: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data$_{l}$*ibrary.html*

[6]A list of the 49 portfolios with their 4-digit SIC code is available in the online appendix Table A.1.

[7]Collected from the press release archives at: http://www.opec.org/opec$_{w}$*eb/en/press*$_{r}$*oom/28.htm*.

announcement, $F^h_{t,d}$.

$$OSS^h_{t,d} = log(F^h_{t,d}) - log(F^h_{t,d-1}) \tag{1}$$

where $F^h_{t,d}$ is the oil futures price for the h months contract on day d of month t, $OSS_{t,d}$ denotes the oil surprise series at a particular maturity and t and d represent the month and the day of the announcement, respectively.[8] Regarding the maturity of the futures contract, h, we follow Känzig (2021) and use futures contracts with maturities spanning from one to 12 months, which is suitable to the time lag between OPEC announcements and their implementation.

After computing the daily series of surprises at each maturity, we extract their first principal component. Then, to match the frequency of the oil market data, we compute a monthly oil surprise series in the following manner. For months where there are no announcements, the oil surprise series is set to zero. When several announcements take place in a single month, the monthly oil surprise series is the sum of all the daily surprises in the month. In the case of a single announcement in a given month, the monthly oil surprise series is equivalent to the solo daily one.

We note that while Känzig refers to the instrument as *oil supply surprises*, throughout the text we refer to it as the *oil surprise series.* After all, changes in expectations regrading future oil prices around the OPEC announcements may reflect agents beliefs not only about future supply but also demand. Indeed, work by Kilian and Zhou (2023) argues that OPEC announcements may reflect decisions to adjust oil production stemming from a decline in demand. For instance, the press release following the meeting of OPEC on October 24, 2008 stated: *"The Conference observed that the financial crisis is already having a noticeable impact on the world economy, dampening the demand for energy, in general, and oil in particular. This slowdown in oil demand is serving to exacerbate the situation in a market that has been over-supplied with crude for some time [...] Accordingly, the Conference has decided to decrease the current OPEC-11 production ceiling of 28.808 million barrels a day by 1.5 mb/d, effective 1 November 2008.* It seems clear that in this and other situations, the decision to cut production was driven by anticipation of lower demand and oil prices.

[8] When the conference ends on a weekend, holiday, or after the NYMEX crude close, we use the date of the next trading day.

3 Empirical Strategy

To investigate the effect of shocks to oil price expectations on U.S. stock returns, and to gauge the effect of informational asymmetries, we employ two estimation strategies. First, we follow the specification put forward by Känzig (2021) and estimate a proxy-VAR model. We then compare these estimates with responses estimated from a VARX model where we first purge in the oil surprise variable. This section describes both empirical strategies.

3.1 Proxy-VAR model

We first estimate a proxy-VAR model (Stock and Watson (2012) and Mertens and Ravn (2013)) à la Känzig (2021) where oil surprises are employed as an instrument. That is, we use the oil surprise series derived from the daily oil futures data, aggregated to a monthly frequency (see subsection 2.2), as an instrument to identify the oil news shock.[9]

To estimate the effect of a shock to oil market expectations (an oil news shock) we consider the following reduced-form VAR(p)

$$y_t = \boldsymbol{c} + \boldsymbol{B_0} + \boldsymbol{B_1} y_{t-1} + \ldots + \boldsymbol{B_k} y_{t-p} + \epsilon_t \tag{2}$$

where y_t is a 5×1 vector that contains the oil market series followed by the stock return of interest (e.g., the aggregate index or one of the 49 industry-level portfolios); c is a 5×1 vector of intercepts; $\boldsymbol{B_i}$ denote 5×5 matrices of autoregressive coefficients; and $\boldsymbol{\epsilon_t}$ denotes a 5×1 vector of Gaussian innovations with mean $\mathbf{0}$ and variance $\mathbf{\Sigma}$. The oil market variables include, in order, the real price of oil, world oil production, world oil inventories, and a measure of real economic activity. We rotate each of the stock returns, one at a time, after the oil variables. We follow Känzig (2021) and set the lag order p equal to 12 and estimate the model in log levels.[10]

The Wold representation of y_t is given by

$$y_t = C(L)\epsilon_t \tag{3}$$

where $C(L)$ is a matrix lag polynomial and we assume the innovations ϵ_t are linear combi-

[9] Proxy-VARs (SVAR-IV) have also been increasingly used to identify the effects of other types of shocks (e.g., monetary policy, uncertainty, etc.) See, among others, Carriero *et al.* (2015), Gertler and Karadi (2015), Caldara and Kamps (2017), and Stock and Watson (2018)).

[10] As we will show in section 6 the results are robust to various alternative specifications.

nations of the structural shocks u_t at time t via some framing $\boldsymbol{S}$ such that:

$$\epsilon_t = \boldsymbol{S}u_t = \begin{bmatrix} s_{11} & s_{12} & s_{13} & s_{14} & s_{15} \\ s_{21} & s_{22} & s_{23} & s_{24} & s_{25} \\ s_{31} & s_{32} & s_{33} & s_{34} & s_{35} \\ s_{41} & s_{42} & s_{43} & s_{44} & s_{45} \\ s_{51} & s_{52} & s_{53} & s_{54} & s_{55} \end{bmatrix} \begin{bmatrix} u_{1t} \\ u_{2t} \\ u_{3t} \\ u_{4t} \\ u_{5t} \end{bmatrix} \tag{4}$$

where u_t ~WN (0, I) and $\boldsymbol{S}$ is a nonsingular matrix. Therefore, the variance-covariance matrix of the reduced-form shocks ϵ_t is given by

$$\boldsymbol{\Sigma} = E[\epsilon_t \epsilon_t'] = E[\boldsymbol{SS'}]. \tag{5}$$

Assume, without loss of generality, that the oil news shock (the structural shock of interest) is given by u_{1t} and let the $(n-1) \times 1$ vector $u_{qt} = [u_{2t}, u_{3t}, u_{4t}, u_{5t}]'$ collect the remaining structural shocks. Similarly, let ϵ_{qt} collect the second to the last reduced-form innovations. Then $\boldsymbol{s_1}$, the first column in matrix $\boldsymbol{S}$, corresponds to the impact of the structural shock u_{1t} on each of the reduced-form residuals ϵ_t. To compute the impulse responses of interest, it suffices to estimate:

$$y_t = \boldsymbol{c} + \boldsymbol{B_0} + \boldsymbol{B_1} y_{t-1} + ... + \boldsymbol{B_k} y_{t-p} + \boldsymbol{s_1} u_{1t} \tag{6}$$

Note that, because our sole interest is the effect of an oil news shock, we only need to identify the elements of the column $\boldsymbol{s_1}$.

Therefore, to estimate the impulse response functions, we proceed as follows. Let z_t denote the instrument (the oil surprise series) and assume the relevance and exogeneity conditions hold so that:

$$Relevance : E[z_t u_{1t}] = \phi \neq 0 \tag{7}$$

$$Exogeneity : E[z_t u_{qt}] = 0. \tag{8}$$

In other words, assume the oil surprise series carries information about the oil news shock u_{1t} (relevance), but is not driven by any other shock u_{qt} (exogeneity).

Then, to obtain estimates of $\boldsymbol{s_1}$ in equation (6), we first estimate by OLS the reduced-form

VAR model in (2) and obtain the fitted residuals $\hat{\epsilon}_t$. Then, given that s_{i1} for $i = 2, ..., 5$ corresponds to the response of ϵ_{it} to a unit increase in u_{1t} (i.e., a unit oil news shock), we obtain an estimate of the ratio s_{i1}/s_{11} from a two stage least square regression of the reduced-form residuals $\hat{\epsilon}_{qt}$ on $\hat{\epsilon}_{1t}$ using the instrument z_t where $\hat{\epsilon}_{1t}$ is the first element of $\hat{\epsilon}_t$. To ease interpretation, we normalize the responses so that an oil news shock corresponds to an 10% increase in real oil prices.

An advantage of this estimation method is that it allows us to first obtain more precise estimates of the reduced-form residuals by employing the sample starting in 1973:1. And then use the shorter sample starting in 1983:4 (when future oil prices become available) to compute the two stage least square regressions to get $\boldsymbol{s_1}$ and the impulse response functions.

3.2 VARX Model

We employ the Proxy-VAR as our baseline because its two-step procedure allows us to estimate the responses using a longer time series (1973:1-2019:12) even though the instrument is only available starting in 1983:4. However, to investigate the robustness of the results and the impact of purging the instrument, we estimate a VARX model using a sample that starts in 1983:4, when oil futures prices became available. The VARX model is given by:

$$y_t = \tilde{\boldsymbol{B}}_{\boldsymbol{0}} + \tilde{\boldsymbol{B}}_{\boldsymbol{1}} y_{t-1} + ... + \tilde{\boldsymbol{B}}_{\boldsymbol{k}} y_{t-k} + \tilde{\boldsymbol{A}} z_t + \tilde{u}_t \tag{9}$$

We collect the regressors in $\tilde{x}_t = [1, y'_{t-1}, ...y'_{t-k}, \tilde{z}_t]'$ and the coefficients are given by the matrices $\tilde{\boldsymbol{B}} = [\tilde{\boldsymbol{B}}_{\boldsymbol{0}}, \tilde{\boldsymbol{B}}_{\boldsymbol{1}}, ..., \tilde{\boldsymbol{B}}_{\boldsymbol{K}}, \tilde{\boldsymbol{A}}]'$. The series, $\tilde{z}_t$ corresponds to the purged oil surprise series, which is the fitted residual obtained from projecting the original oil surprise series on the oil market variables. We assume the relevance and exogeneity IV conditions (i.e., equations (7 and 8) for $\tilde{z}_t$ are met.[11] Let

$$\underset{T\times nk+2}{\tilde{X}} = \begin{bmatrix} \tilde{x}'_1 \\ \tilde{x}'_2 \\ . \\ . \\ \tilde{x}'_T \end{bmatrix} \quad \underset{T\times 1}{Z} = \begin{bmatrix} \tilde{z}'_1 \\ \tilde{z}'_2 \\ . \\ . \\ \tilde{z}'_T \end{bmatrix} \quad \underset{T\times n}{\hat{\tilde{U}}} = \begin{bmatrix} \hat{\tilde{u}}'_1 \\ \hat{\tilde{u}}'_2 \\ . \\ . \\ \hat{\tilde{u}}'_T \end{bmatrix} \quad \underset{T\times n}{Y} = \begin{bmatrix} y'_1 \\ y'_2 \\ . \\ . \\ y'_T \end{bmatrix},$$

[11] See Paul (2020) for discussion on the conditions for consistency of the relative impulse response functions.

and $\hat{\tilde{\boldsymbol{B}}} = (\tilde{X}'\tilde{X})^{-1}\tilde{X}'Y$ and $\hat{\tilde{U}} = Y - \tilde{X}\hat{\tilde{\boldsymbol{B}}}$ denote the ordinary least squares coefficients and error terms, respectively. Note that $\tilde{X} = [X \ Z]$.[12]

We can rewrite the regression coefficient matrix $\hat{\tilde{B}}$ as follows:

$$\hat{\tilde{B}} = \begin{pmatrix} X'X & X'Z \\ Z'X & Z'Z \end{pmatrix}^{-1} \times \begin{pmatrix} X'Y \\ Z'Y \end{pmatrix}$$

The last row of the $\hat{\tilde{B}}$ matrix associated with the exogenous variable Z is:

$$\hat{\tilde{A}}' = -(Z'Z - Z'X(X'X)^{-1}X'Z)^{-1}Z'X(X'X)^{-1}Z'Y + (Z'Z - Z'X(X'X)^{-1}X'Z)^{-1}Z'Y$$
$$\hat{\tilde{A}}' = \tilde{k} \times [Z'Y(I - X(X'X)^{-1}X')]$$

Where $\tilde{k} = (Z'Z - Z'X(X'X)^{-1}X'Z)^{-1}$. Note that $\hat{B} = (X'X)^{-1}X'Y$ and $\hat{\epsilon}_t = Y - X\hat{B}$ from equation 2. Hence the above equation for $\hat{\tilde{A}}'$ reduces to:

$$\hat{\tilde{A}}' = \tilde{k} * [Z'\hat{\epsilon}_t] \tag{10}$$

Then, the contemporaneous relative impulse response can be computed as the ratio of the elements in $\hat{\tilde{A}}$:

$$\frac{\hat{\tilde{A}}_i}{\hat{\tilde{A}}_j} = \frac{Z'\hat{\epsilon}_i}{Z'\hat{\epsilon}_j}$$

$\hat{\epsilon}_i$ and $\hat{\epsilon}_j$ are the columns i and j in the reduced form innovations $\hat{\epsilon}_t$ from equation 2. The subsequent impulse responses are obtained by using the impact response to the oil news shocks and iterating on equation (9).

4 Oil News Shocks and U.S. Stock Returns Before the COVID-19 Pandemic

This section first describes the estimation results obtained using the Proxy-VAR model.[13]. Given that we employ a longer sample than Känzig (2021), we start by reporting the results for the aggregate/oil market block and then describe the results for the aggregate and industry-level

[12]The proofs for the exogenous variable approach were originally produced in Paul (2020) online appendix A.2. We now recall those proofs for our purpose in this paper.

[13]Recent literature on the use of external instruments for causal inference in SVAR has underlined the perils associated with the use of news shocks that are noninvertible. To assuage such concerns, we pre-test for invertibility following Plagborg-Møller and Wolf (2022). We reject the null that oil news does not Granger cause y_t at a 0.00001 significance level, which indicates non-rejection of invertibility.

stock returns. A discussion comparing our findings to the earlier empirical literature follows and the section concludes with a comparison to results obtained from a VARX model.

4.1 Macroeconomic Aggregates

Figure 1 plots the response of the macroeconomic and oil market variables to the oil news shock estimated from the proxy-VAR. The 90% and 68% confidence bands computed using 10,000 bootstrap replications are denoted by the light and dark-blue shaded areas, respectively.[14] The responses have been normalized so that a negative oil news shock results in a 10% increase in the real price of oil. As the figure illustrates, an oil news shock leads to an immediate and persistent increase in the real price of oil, an immediate but only marginally significant increase in oil production followed by a persistent decline, and a long-lasting contraction in global economic activity.

These results are largely consistent with the estimates obtained by Känzig (2021) using a shorter sample period. However, a noticeable difference is that our estimates suggest a somewhat smaller and slower accumulation of crude oil inventories in response to the oil news shock. Twelve (36) months after the oil news shock we estimate inventories have increased by less than 0.5% (1%), compared to the 0.7% (1.5%) increase estimated by Känzig (2021) in the 1974:1-2017:12 sample.

All in all, the responses of oil prices and inventories to oil news shocks suggest uncertainty regarding future oil supply plays a key role in the transmission of oil news shocks to the global economy. In the wake of an OPEC announcement regarding future cuts in oil production, market participants immediately increase the demand of crude oil for storage and continue to accumulate inventories for prolonged period of time. This behavior is consistent with an increased demand driven by precautionary motives in expectation of higher oil prices.

As for the U.S. aggregate stock returns, we find a negative but insignificant response to an oil news shock. A similar decline in U.S. aggregate stock returns to oil-specific demand shocks is estimated by Kilian and Park (2009) using a shorter sample that spans from 1973:1 until 2006:12. While their specification and estimation method differs from ours,[15] the similarity between the response of U.S. stock returns depicted in Figure 1 and their estimated response

[14] The confidence bands are computed using the moving block bootstrap proposed by Jentsch and Lunsford (2019).

[15] Kilian and Park (2009) estimate a structural VAR that does not include crude oil inventories, uses an alternative measure of global economic activity and focuses on the impact of demand and supply driven oil shocks

to an oil-specific demand shock is indicative of the role played by increased storage demand in response to oil news shocks.

4.2 Industry-Level Portfolios

Does the insignificant response of aggregate U.S. stock returns to oil news shocks mask important differences across diverse industry portfolios? Work by Kilian and Park (2009) and Alsalman and Herrera (2015) suggests this might be the case. For instance, while Alsalman and Herrera (2015) do not estimate the response of stock returns to oil news shocks, their estimated response to oil price shocks (representing an average response to unexpected increases in oil prices stemming from different sources) indicate that while stock returns for industries such as rubber and plastic products and automobiles and trucks exhibit a decline following an unexpected oil price increase, such a shock constitutes a win for precious metals, coal and petroleum and natural gas. Similarly, Kilian and Park (2009) find that automobiles and trucks exhibit a negative response to oil-specific demand shocks, whereas the response of petroleum and gas is positive.

The aim of this section is twofold. First, we investigate whether the response to oil news shocks is heterogeneous across different industries. Then we explore whether the patterns we uncover resemble the response to oil supply shocks or storage demand shocks. Understanding how oil supply shocks affect stock returns across different industries sheds some light regarding the role of oil price expectations in the transmission of oil news shocks.

Panel A of Table 1 reports the cumulative response of each of the 49 industry-level portfolios to an oil news shock. As mentioned earlier, estimates are obtained by rotating the industry-level portfolios as the last variable in the proxy-VAR. For the sake of brevity, we report the cumulative response at three representative horizons (h = 0, 6, and 12 months). As the table illustrates, on impact, oil news shocks lead to a significant increase in stock returns for 28 portfolios. Yet, for most of the stocks where a negative oil news shock constitutes a win, the effect is short-lived and lasts less than a year. The exceptions where the gains derived from a 10% increase in the price of oil continue to be significant after twelve months are: coal (8.45%), precious metals (7.27%), healthcare (6.85%), petroleum and natural gas (6.81%), defense (5.13%), fabricated products (4.29%), utilities (3.60%), wholesale (2.53%), insurance, (2.50%), pharmaceutical products (2.35%) and, shipbuilding and railroad (2.14%). In contrast, the oil news shock results in a significant, immediate, and persistent decline in stock returns for computer hardware, electronic equipment, computer software, shipping containers, and retail. In addition, portfolios such as consumer goods, rubber and plastic and textile do not exhibit a significant decline on impact, but do experience a drop 12-months after the shock. One year after the shock, declines in stock returns resulting for a negative oil news shocks that leads to a 10% increase in oil prices range from -8.73% for computer hardware to -2.14% for textiles. Interestingly, whereas automobiles and trucking -an industry that has historically been hit the hardest by unexpected oil price increases- experiences a smaller decline in stock returns (-4.03%) than computer hardware (-8.73%) and electronic equipment (-6.78%).

Figures 2 and 3 depict the responses of stock returns for portfolios that have historically been shown to be very responsive to oil price shocks (e.g., automobiles and trucking, petroleum and gas) and portfolios that we find to be sensitive to oil news shock, but were not found to respond significantly to oil price shocks in the earlier literature (e.g., computer hardware, electronic equipment).[16] The 90% and 68% confidence bands computed using 10,000 bootstrap replications are denoted by the light and dark-blue shaded areas, respectively.

Three interesting patterns emerge from these figures. First, for most industries that exhibit a decline in stock returns, the trough occurs around ten months after the shock and the losses become insignificant about two years after the negative oil news shock hits the economy. Second, industries that experience an increase in stock returns show a greater degree of heterogeneity regarding the persistence of the response. In particular, the positive response of petroleum and gas, and utilities persist after four years, whereas the wins for coal, precious metals and steel works become insignificant less than three years after the shock.

To summarize, we find a significant degree of heterogeneity in the dynamic response of industry-level portfolios to oil news shocks. The fact that some industries suffer losses while other experience gains explains why the response of aggregate U.S. stock returns is statistically insignificant.

4.3 Discussion and comparison with VARX estimates

As Paul (2020) proves, impulse responses[17] estimated via an external instrument approach (e.g., a proxy-VAR) or a VARX consistently identify the true contemporaneous impulse response functions. In fact, even in small samples the contemporaneous response is the same. Consistency

[16] Plots of the impulse response functions for the remaining portfolios are available upon request.

[17] Throughout the paper we refer to the relative impulse response, that is the impulse response for each variable normalized so that the oil news shock generates a 10% increase in oil prices, as the impulse response functions.

of the VARX for subsequent impulse responses hinges on the assumption that z_t is uncorrelated with lags of y_t. In our context, this assumption amounts to lags of y_t having no predictive content for oil surprise, an assumption we test and cannot reject in in the pre-COVID-19 sample.

Nevertheless, to assuage concerns that results may be biased because of correlation between the oil surprise and the rest of the explanatory variables, we also report estimation results for a VARX model where we project the news on all other regressors and use the residuals from this projection in the estimation. Panel A of Table 2 reports the cumulative impulse response functions for three horizons of interest ($h = 0, 6, 12$) obtained with the VARX. Differences in the estimated responses stem from two sources. First, estimates for the proxy-VAR are obtained without purging the instrument. Second, to obtain the proxy-VAR (IV-VAR) responses, we estimate the reduced form VAR using the full 1973:1-2019:12 sample. Then we project the residual of the oil price equation on the oil surprise instrument. Because data for the latter is only available starting in 1983:4, we use a shorter sample in this step. Using the longer sample in the first step allows us to obtain more precise estimates of the reduced-form VAR coefficients. However, the results are robust to using the shorter sample in both steps.

Note that, while the estimates of the proxy-VAR and the VARX are not numerically identical, the magnitudes are very similar and qualitatively they provide identical answers to key questions such as: What are the implications of oil news shocks for investors' portfolio choices? How are these shocks transmitted to aggregate stock returns? How do the estimated responses of stock returns to oil news shocks compare to previous estimates investigating the effect of oil price shocks (e.g., Alsalman and Herrera (2015) and Kilian and Park (2009))?

Our estimates indicate, as the reader might expect, that portfolio losses/gains are related to an industry's exposure to oil price fluctuations. For instance, shares in the automobile, shipping containers, and rubber and plastic industries –which are intensive in the use of oil in consumption or production– experience depreciation in response to news of future oil supply cuts. This depreciation resembles the response to an oil-specific demand shock estimated by Kilian and Park (2009) as well as that obtained by Alsalman and Herrera (2015).

A result that stands in contrast with earlier research is the large cumulative decline in returns faced by the IT industries (computer hardware and software, and electronic equipment) in response to an oil news shock. The loss for shares in IT industries suggests oil news shocks operate in a manner similar to uncertainty shocks in that they result in curtailed demand for investment in IT and, hence, a loss for shares in computer and electronic industries. An important implication for investors is that oil price increases related to OPEC's announcements of future cuts in oil supply pose a risk for the value of computer hardware and software as well as for electronic equipment stocks.

On the contrary, in response to the same oil price increase, the share price of energy companies (coal, petroleum, and natural gas) experience significant gains. Such appreciation is consistent with the heightened storage demand that follows in the wake of a negative oil news shock. This is also the case for shares of precious metals. These results are in line with previous studies suggesting that shares of companies that own crude oil, or that produce goods that could be considered useful in hedging against risk in times of political turmoil (i.e., precious metals), appreciate in the face of unexpected oil price increases. However, in contrast with earlier investigations, we find significant appreciation for the shares of healthcare, defense, insurance and pharmaceutical products. Such response suggests an increase in the demand for these industries' products, and a perceived increase in inflation risk, constitute important transmission mechanisms for oil news shocks.

We note that results for the pre-COVID period obtained with both estimation methods are largely consistent with Känzig (2021) who uses a sample that ends in December 2017. He estimates the response of stock returns for six industries and finds that the shares of oil and gas, precious metals, and electricity appreciate in the short run, whereas those of automobiles and parts, retail, and travel and leisure industries show an immediate and persistent fall.

5 When pandemics hit and wars break out

If "wars are the realm of uncertainty" (Von Clausewitz, 1950), pandemics are its kingdom. The massive shock of the COVID-19 pandemic and the shutdown measures aimed at containing it, plunged the global economy in the deepest recession since World War II. The uncertainty associated with the pandemic amplified the recessionary impact of the demand and supply disruptions that ensued the global shutdown. The effect on oil markets was unprecedented: travel stopped, the demand for oil plummeted, and the oil industry could not adjust production fast enough to match it. On April 20 of 2020, the price of the WTI dropped almost 300% and traded at about -$37 per barrel. At the onset of the pandemic, concerns that oil demand would take a long time to recover were expressed by several oil companies. Yet, less than two years later, Russia invaded Ukraine and oil prices soared to more than $120 a barrel amidst concerns of a global shortfall in oil supply.

Studying the effects of oil news shocks in such an environment is a difficult endeavour. In particular, part of the variation in WTI oil futures prices around OPEC meetings might be accounted for by the way in which OPEC responded to the unparalleled drop in oil demand and not solely reflect news of oil supply changes. Moreover, the reader might worry that during these unprecedented times, foresight by economic agents and increased uncertainty could muddle the estimation of the impulse responses. [18] For these reasons, in the previous section we abstracted from the years of the pandemic and the war in Ukraine. In this section, we extend the data to December of 2022 and re-estimate our models to explore if and how our estimates are affected by extending the sample.

5.1 Proxy-VAR results assuming business as usual

Extending the sample to include data from 2020 until December 2022 somewhat alters the results for the macroeconomic aggregates. More specifically, while the response of the real oil price is virtually unchanged, the negative responses of world oil production and world industrial production become less persistent. Not surprisingly, the response of the oil market variables is estimated with a lower degree of precision when we include the more uncertain 2020-2022 years.

What is more striking is the positive and significant response of aggregate stock returns on impact and in the short-run, which stands in contrast with the insignificant response found in the pre-COVID sample. For instance, estimation results obtained for the extended sample with the Proxy-VAR (VARX) –reported in Panel B of Table 1 (Table 2)– indicate a 2.85% (2.17%) increase in aggregate stock returns on impact, which contrasts with the insignificant 0.60% (0.66%) increase estimated with the pre-COVID data. Moreover, the positive response persists for about six months.

This increased sensitivity of aggregate stock returns begs the question of what portfolios drove the change in the estimated aggregate response in the extended sample. To investigate this issue, let us first focus on the estimates obtained with the Proxy-VAR. On the one hand, estimation results reported in Panel B of Table 1 indicate that, on impact, an oil news shock that results in a 10% increase in oil prices causes a decline in stock returns for only two of the portfolios (computer hardware, and electronic equipment). Twelve months after the shock,

[18] To assuage concerns regarding the role of foresight and non-invertibility, we pre-test for invertibility using the extended sample as suggested by Plagborg-Møller and Wolf (2022). Here, again, we cannot reject the null of invertibility.

a significant decline is still evident in these two sectors (-4.87%, and -2.60%, respectively) as well as in shipping containers (-1.51%). However, the magnitude of the decline is estimated to be about one-half of the magnitude of the pre-COVID sample. On the other hand, industries commonly perceived as being sensitive to oil price fluctuations and whose stock returns were estimated to decline 12 months after the oil news shocks in the 1973:1-2019:12 sample, do not appear to have suffered losses when the sample is extended (e.g., automobiles, consumer goods, rubber and plastics, and textiles). A possible explanation for the positive response of aggregate stock returns could be a shift in the relationship between oil surprises and industry-level stock returns during the times of the pandemic. While the Proxy-VAR estimates seem to provide some support for this hypothesis, such result seem to run against while the macroeconomic theory would suggest. One would expect industries for which production costs increase or demand declines when oil prices rise to suffer losses.

Nevertheless, an alternative conjecture is that the shift was due to increased correlation between the oil surprise measure and the macroeconomic or oil market variables. In particular, asymmetric information issues could have been more prevalent during the years of the pandemic than during the rest of the sample. For instance, both OPEC and private agents could have received noisier signals about the fundamentals of the economy during 2020-2022. Hence, the dynamics of the instrument could have been very different during the last three years of the extended sample. If that were the case, estimates that do not account for this change could be misleading. The next section investigates this issue.

5.2 Did the dynamics of the instrument change?

Fluctuations in the oil surprise variable, which is computed as shifts in oil futures prices around OPEC announcements, were more erratic in the pre-COVID era than in the months spanning the 2020-2022 years (see Figure 4). More specifically, during the COVID-19 pandemic and the onset of the war in Ukraine, OPEC appears to have conducted policy in a more systematic manner. This would suggest that: (a) the dynamics of the oil surprise measure changed in the extended sample; (b) oil price shocks identified via movements in oil futures prices might reflect true oil news shocks in the pre-COVID period, but not in the extended sample.

Indeed, as Figure 4 illustrates, the oil surprise measure and the measure obtained as the residual of a projection of the oil surprise measure on the macro and oil market variables (hereafter, the purged oil surprise) roughly coincide before 2020. Yet, larger differences between the two series are observed during 2020-2022 and the original series appears to be serially correlated. For instance, a large decline in oil surprise measure is observed in June 2020, when OPEC+ reached an agreement to extend production cuts until August, shut-downs started to lift, and economic activity started to recover. Yet, no such drop is evident in the purged oil surprise. Furthermore, for most of the first half of 2021, the purged measure systematically exceeds the raw oil surprise measure, which suggests the former has an endogenous component. To evaluate this hypothesis, we test for Granger causality in the pre-COVID and the extended sample. Our test results indicate that none of the macro or oil market variables have predictive content for the oil surprise measure in the pre-COVID sample, but the real oil price has predictive power for the oil surprise variable in the extended sample. We reject the null of no Granger causality from oil prices to the oil surprise measure (p-value=0.0118). That is, during the COVID period, oil surprises may comprise superior information on the part of OPEC regarding the state of the oil market and the global economy. In other words, when extending the sample to include the months spanning from January 2020 to December 2022, we find that the instrument is contaminated by lags of the real oil prices: the lag exogeneity condition needed for consistency of the Proxy-VAR is violated in the extended sample (see Miranda-Agrippino and Ricco (2023)). Therefore, in the next section we discuss the results obtained when we first project the oil surprise variable on lags of the macroeconomic and oil market variables, and use the residual as our measure of oil news shocks in a VARX framework.

5.3 VARX results taking into account changes in the dynamics of the instrument

What happens when the instrument is purged from the contamination stemming from the oil market variables? Estimation results for the VARX, reported in Panel B of Table 2 and Figures 5-7, suggest the traditional transmission mechanisms are still at play in the medium and long-run, albeit the impact response changes slightly. On impact, the estimated responses for the VARX are close to the Proxy-VAR estimates, as the theory would suggest. That is, some of the industries that traditionally experienced losses in returns exhibit gains in the extended sample. Yet, the subsequent impulse responses reveal a decline in returns for several portfolios. Moreover, this decline is roughly consistent with the drop found for the pre-COVID sample, albeit estimated with a larger degree of uncertainty.

What explains the differences between the proxy-VAR and the VARX estimates for the extended sample? The answer to this question appears to be that the oil surprise variable (i.e., the instrument in the proxy-VAR) contains an endogenous component, which is possibly linked to greater information asymmetries and/or more noise during the COVID years. Once we purge the endogenous component from the news series, we find the VARX estimates a negative effect on key industry-level portfolios such as automobiles, shipping containers, transportation, and rubber and plastics. Moreover, the negative and significant effect on stock returns for the IT industries persists. Lastly, the magnitude of the response for industries that experience gains is similar across samples when estimated with the VARX.

All in all, estimation results suggest care must be taken when using high-frequency identification methods to estimate the effect of oil price shocks during periods of turmoil. In particular, not taking into account the possible endogeneity of the oil surprise measure during the 2020-2022 years could lead the researcher to estimate effects of oil price shocks that run against the theoretical underpinnings of the transmission channels. When this issue has been taken into account, we find that oil news shocks affect industry-level portfolios in a manner similar to that of other uncertainty shocks and oil market-specific shocks. That is to say, oil news shocks pose a risk for high-tech industries (e.g., computer hardware and software, electronic equipment), as well as industries that are intensive in the use of oil in production (e.g., rubber and plastics) or consumption (e.g., transportation).

6 Robustness Checks

This section reports the results of a series of robustness checks carried out on the pre-COVID sample. More specifically, we consider a model where the oil variables are expressed in growth rates rather than levels, we employ an alternative measure for real economic activity, we consider a different measure of oil surprises and a longer lag order. For the sake of brevity, we describe the results here and plot some of the impulse responses in Figures 8-10, but relegate the complete set of figures plotting the impulse responses to an Online Appendix.

Alternative instrument. As noted earlier, we construct the oil surprise instrument by taking the principal component of the changes in the oil futures prices, from one up to twelve months, around the OPEC meetings. As a robustness check, we compute the instrument using only the changes in the 6-month futures. Figures 8-10 demonstrate that the results are robust to this change in the instrument. (See Figures A.1-A.3 of the online appendix).

Growth rates. Earlier literature on the effect of oil price shocks (see, e.g., Kilian (2009), Herrera *et al.* (2019)) estimates SVARs, where the real price of oil, the world oil production, and the measure of global economic activity are expressed in rates of growth, and oil inventories are measured as the change from the previous period. While we estimate our model in levels to facilitate comparison with earlier work by Känzig (2021), we note that impulse response functions (Figures 8-10) estimated using an alternative specification in rates of growth produce very similar results. (See Figures A.4-A.6 of the online appendix).

Kilian's index of real economic activity. Kilian (2009)'s index of global economic activity indicator has been commonly employed in the oil price-macroeconomy literature as an indicator of global economic activity.[19] Figures A.7-A.9 illustrate the responses of the oil market variables, aggregate stock returns, and the 24 portfolios of interest to oil news shocks usingKilian (2009)'s index as the global economic activity indicator. Our estimation results are robust to using this alternative indicator.

Different lag order. Using a longer lag length of 24 months has been common in the literature. Thus, we re-estimate our model using 24 lags instead of 12.[20] Figures A.10-A.12 reveal very similar patterns when we increase the lag length.

Alternative measure of inventories in Känzig (2021) specification. We estimate the model using the same specification of the oil market data as in Känzig (2021), but replacing the level of oil inventories with its first differences. The results are robust to employing this alternative measure of inventories (See Figures A.13-A.15 of the online appendix).

Alternative oil price indicator. In this paper, we use the real refiner acquisition cost of imported crude oil, deflated by the U.S. CPI as a measure of real oil prices. Alternatively, one can employ the real West Texas Intermediate (WTI) spot price. Estimation results reported in Figures A.16 - A.18 of the online appendix indicate our findings are virtually unchanged when we replace the refiner's acquisition cost with the WTI.

Only ordinary announcement. Announcements that follow extraordinary OPEC meetings may have a different information content than those stemming from ordinary meetings. To evaluate whether these extraordinary announcements drive the response of U.S. stock prices we re-estimate the proxy-VAR model after excluding the 39 extraordinary meetings from the

[19]See, among others, Kilian and Zhou (2018), Kilian (2019), Funashima (2020), and Nonejad (2020) about the prominence of the Kilian index as a leading indicator for the world industrial production index.

[20]A large amount of literature in the oil market models applies 24 lags (see, among others, Kilian and Park (2009), Kilian and Murphy (2014), Güntner and Linsbauer (2018), Herrera *et al.* (2019), and Alsalman and Karaki (2019)).

150 total announcements (i.e., 26% of the announcements) between 1983:4 and 2022:12 (see Table A.2 of the online appendix). Figures A.19-A.21 of the online appendix plot these impulse response functions and yield results very similar to the baseline estimates obtained using both ordinary and extraordinary announcements.

7 Conclusions

Using a high-frequency identification scheme, we investigated the effect of oil market expectations on U.S. aggregate and industry-level stock returns. We started our study by estimating a proxy-VAR model on a sample spanning the period between January 1973 and December 2019 where surprises in oil futures prices around OPEC meetings were used as an instrumental variable. We complemented our analysis by estimating a VARX model where a ”purged” measure of the oil surprises was used as an exogenous variable. We then extended the data until December 2022 so as to cover the period of the COVID-19 pandemic and the first months of the war in Ukraine.

We derived several key insights. First, we found oil news to have a statistically insignificant effect on aggregate stock returns in the pre-COVID-19 sample. Yet, this insignificance masks considerable heterogeneity in the response across industry-level portfolios. More precisely, while IT industries and industries that are intensive in the use of energy (e.g., automobiles, shipping containers, rubber and plastic) suffered losses, other portfolios such as petroleum and gas, coal, healthcare, defense, and precious metals experienced gains.

Second, while sectors commonly-thought as being sensitive to oil price shocks (e.g., automobiles, trucks, rubber and plastics, shipping containers) behaved in a manner that resembled the response to an oil-specific demand shock (Kilian and Park, 2009)), the responses of other industries suggested an alternative transmission channel. Namely, losses in the returns of the IT sector and gains in precious metals pointed towards an uncertainty (precautionary) channel. In particular, we concluded that our results are consistent with increased expectations of higher oil prices resulting in curtailed demand for the IT industry and increased demand for primary metals.

Lastly, we showed that during periods of turmoil –such as the 2020-2022 years–, the measure of oil surprise news proposed by Känzig (2021) could be contaminated by lags of the oil market variables. As a result, we found that the use of this instrument in a proxy-VAR setting may bias the estimates of the impulse response functions. We then showed that the residual obtained after projecting the oil surprise measure on the macro and oil market variables produces a measure of true shocks to expectations about oil prices (i.e., it is serially uncorrelated and unforcastable). Indeed, the VARX estimated with the purged news variable produced similar responses in the pre-COVID-19 and the extended sample with one noticeable exception: aggregate stock returns responded positively on impact, but the effect vanished after ten months. The positive aggregate response on impact appears to be accounted for by the rapid recovery of industries that suffered losses during the shut-downs. The estimated response of stock returns for restaurants, recreation, beer and liquor, and real estates for the extended sample more than doubled the estimate for the pre-COVID-19 period, suggesting news of oil price increases conveyed information of future increases in oil demand.

In brief, this paper provided empirical evidence suggesting oil news shocks affect stock returns mainly through three channels. First, oil news shocks translate into higher costs and, thus lower stock returns, for industries that are energy intensive. Second, oil news operate as oil-specific demand shocks, hence resulting in gains for energy related portfolios. Third, expectations about higher future oil prices act as uncertainty shocks negatively affecting returns in IT sectors. We conclude by underlining the importance of controlling for possible contamination of the instrument when considering periods of turmoil.

References

Alsalman, Z. and Herrera, A. M. (2015) Oil price shocks and the us stock market: do sign and size matter?, *The Energy Journal*, pp. 171–188.

Alsalman, Z. N. and Karaki, M. B. (2019) Oil prices and personal consumption expenditures: does the source of the shock matter?, *Oxford Bulletin of Economics and Statistics*, **81**, 250–270.

Barsky, R. B. and Kilian, L. (2002) Do we really know that oil caused the great.

Bloom, N. (2009) The impact of uncertainty shocks, *econometrica*, **77**, 623–685.

Bretscher, L., Hsu, A. and Tamoni, A. (2023) The real response to uncertainty shocks: The risk premium channel, *Management Science*, **69**, 119–140.

Caldara, D. and Kamps, C. (2017) The analytics of svars: a unified framework to measure fiscal multipliers, *The Review of Economic Studies*, **84**, 1015–1040.

Carriero, A., Mumtaz, H., Theodoridis, K. and Theophilopoulou, A. (2015) The impact of uncertainty shocks under measurement error: A proxy svar approach, *Journal of Money, Credit and Banking*, **47**, 1223–1238.

Degasperi, R. (2021) Identification of expectational shocks in the oil market using opec announcements, *manuscript, University of Warwick*.

Editorial (2022) Biden's oil price machinations, *Wall Street Journal*.

Funashima, Y. (2020) Global economic activity indexes revisited, *Economics Letters*, **193**, 109269.

Gertler, M. and Karadi, P. (2015) Monetary policy surprises, credit costs, and economic activity, *American Economic Journal: Macroeconomics*, **7**, 44–76.

Güntner, J. H. and Linsbauer, K. (2018) The effects of oil supply and demand shocks on us consumer sentiment, *Journal of Money, Credit and Banking*, **50**, 1617–1644.

Hamilton, J. D. (2009) Causes and consequences of the oil shock of 2007-08, Tech. rep., National Bureau of Economic Research.

Herrera, A. M., Karaki, M. B. and Rangaraju, S. K. (2019) Oil price shocks and us economic activity, *Energy policy*, **129**, 89–99.

Herrera, A. M. and Rangaraju, S. K. (2020) The effect of oil supply shocks on us economic activity: What have we learned?, *Journal of Applied Econometrics*, **35**, 141–159.

Hooper, K. (2022) Biden pressures oil companies on gas prices, *Politico*.

Jentsch, C. and Lunsford, K. G. (2019) The dynamic effects of personal and corporate income tax changes in the united states: Comment, *American Economic Review*, **109**, 2655–78.

Känzig, D. R. (2021) The macroeconomic effects of oil supply news: Evidence from opec announcements, *American Economic Review*, **111**, 1092–1125.

Kilian, L. (2009) Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review*, **99**, 1053–69.

Kilian, L. (2019) Measuring global real economic activity: Do recent critiques hold up to scrutiny?, *Economics Letters*, **178**, 106–110.

Kilian, L. and Murphy, D. P. (2014) The role of inventories and speculative trading in the global market for crude oil, *Journal of Applied econometrics*, **29**, 454–478.

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

Kilian, L. and Zhou, X. (2018) Modeling fluctuations in the global demand for commodities, *Journal of International Money and Finance*, **88**, 54–78.

Kilian, L. and Zhou, X. (2023) The econometrics of oil market var models, in *Essays in Honor of Joon Y. Park: Econometric Methodology in Empirical Applications*, Emerald Publishing Limited, pp. 65–95.

Mertens, K. and Ravn, M. O. (2013) The dynamic effects of personal and corporate income tax changes in the united states, *American economic review*, **103**, 1212–47.

Miranda-Agrippino, S. and Ricco, G. (2021) The transmission of monetary policy shocks, *American Economic Journal: Macroeconomics*, **13**, 74–107.

Miranda-Agrippino, S. and Ricco, G. (2023) Identification with external instruments in structural vars, *Journal of Monetary Economics*.

Nonejad, N. (2020) An observation regarding hamilton's recent criticisms of kilian's global real economic activity index, *Economics Letters*, **196**, 109582.

OPEC (1990) *OPEC: official resolutions and press releases 1960-1990*, Organization of the Petroleum Exporting Countries.

Paul, P. (2020) The time-varying effect of monetary policy on asset prices, *Review of Economics and Statistics*, **102**, 690–704.

Plagborg-Møller, M. and Wolf, C. K. (2021) Local projections and vars estimate the same impulse responses, *Econometrica*, **89**, 955–980.

Plagborg-Møller, M. and Wolf, C. K. (2022) Instrumental variable identification of dynamic variance decompositions, *Journal of Political Economy*, **130**, 2164–2202.

Stock, J. H. and Watson, M. W. (2012) Disentangling the channels of the 2007-2009 recession, Tech. rep., National Bureau of Economic Research.

Stock, J. H. and Watson, M. W. (2018) Identification and estimation of dynamic causal effects in macroeconomics using external instruments, *The Economic Journal*, **128**, 917–948.

Von Clausewitz, C. (1950) *On war*, vol. 1, Jazzybee Verlag.

Table 1: Cumulative responses estimated using a proxy-VAR

	Panel A: 1973:1-2019:12			**Panel B: 1973:1-2022:12**		
Industry Portfolios / Months after the shock	0	6	12	0	6	12
Aggregate	0.603	-0.393	-0.946	***2.846***	***1.905***	0.979
Computer Hardware	***-5.328***	***-7.123***	***-8.733***	***-2.133***	***-3.196***	***-4.878***
Electronic Equipment	***-4.769***	***-6.859***	***-6.780***	**-1.190**	**-2.218**	**-2.602**
Computer Software	***-3.937***	***-3.752***	***-4.404***	0.112	0.466	-0.702
Automobiles and Trucks	-0.752	***-3.609***	***-4.034***	***3.395***	0.681	0.393
Shipping containers	***-1.609***	***-3.277***	***-3.561***	**1.270**	0.438	**-1.510**
Retail	***-1.852***	***-2.556***	***-2.809***	**1.186**	0.438	-0.133
Entertainment	-0.381	-1.508	-2.534	***2.068***	1.026	-0.416
Consumer Goods	-0.651	-0.975	***-2.357***	**0.599**	0.335	-0.964
Rubber and Plastic	0.088	-0.972	**-2.355**	***1.803***	0.680	-0.984
Textiles	-0.683	-1.286	**-2.135**	***3.433***	1.420	0.660
Electrical Equipment	0.157	-1.011	-1.428	***2.538***	1.247	-0.046
Restaurants	0.391	-1.235	-1.418	***2.697***	1.0503	0.421
Business Services	-0.052	-1.028	-1.324	***2.696***	**1.596**	0.352
Candy and Soda	***1.719***	-0.812	-1.210	***2.659***	1.797	0.854
Transportation	-0.447	-0.557	-1.150	***2.659***	**1.797**	0.854
Machinery	0.823	-0.394	-1.084	***2.817***	1.191	0.137
Recreation	0.486	-0.562	-1.074	***3.449***	**2.104**	0.787
Communication	**0.763**	-0.775	-1.045	***2.429***	0.828	0.491
Food Products	0.388	-0.647	-0.974	***2.116***	0.755	0.390
Chemicals	0.549	-0.237	-0.528	***3.453***	**2.074**	1.624
Printing and Publishing	***1.225***	0.545	-0.420	***4.292***	***3.396***	1.987
Beer and Liquor	0.585	0.510	-0.392	***2.329***	***2.099***	1.106
Trading	***1.776***	-0.284	-0.253	***3.295***	***2.340***	0.853
Business supplies	***1.790***	**1.149**	-0.225	***3.313***	***2.552***	1.148
Agriculture	***1.332***	0.043	-0.222	***2.345***	0.570	-0.459
Aircraft	**1.964**	0.768	-0.184	***5.088***	***3.507***	**2.135**
Measuring and control	-0.454	-0.505	-0.144	**1.112**	1.142	1.008
Apparel	0.384	0.526	0.023	***3.023***	***3.048***	**2.388**
Other	**1.320**	0.879	0.075	***3.847***	***2.831***	1.933
Personal Services	**1.245**	0.544	0.096	***3.995***	***3.007***	**1.891**
Tobacco Products	***1.034***	0.422	0.393	***2.432***	***0.922***	0.737
Medical Equipment	***1.859***	**1.155**	0.646	***2.627***	**1.856**	1.047
Construction Materials	***3.480***	**1.638**	0.800	***5.840***	***3.271***	**2.287**
Real Estate	***2.264***	1.638	0.996	***4.826***	***3.607***	**2.639**
Construction	***2.932***	**2.084**	1.001	***6.350***	***4.297***	**2.684**
Mines	**5.246**	**2.844**	1.608	***6.203***	**2.585**	1.038
Banking	***3.713***	***2.850***	1.882	***6.171***	***4.597***	**3.142**
Steel works	***3.791***	2.208	2.097	***5.735***	***3.605***	**2.773**
Shipbuilding and Railroad	**4.989**	***4.161***	**2.139**	***6.541***	***4.613***	**2.496**
Pharmaceutical Products	***2.864***	***2.440***	***2.349***	***3.013***	***2.564***	**2.116**
Insurance	***4.789***	***3.568***	**2.501**	***5.869***	***4.356***	***3.246***
Wholesale	***2.511***	***2.130***	***2.529***	***4.142***	***3.262***	***3.401***
Utilities	***4.260***	***3.531***	***3.601***	***4.868***	***3.616***	***3.483***
Fabricated Products	***4.421***	***4.171***	***4.292***	***6.417***	***5.369***	***4.617***
Defense	**6.129**	***4.751***	***5.125***	***6.337***	***4.160***	***4.813***
Petroleum and Natural Gas	**8.625**	***7.176***	***6.812***	***12.851***	***10.555***	***10.191***
Healthcare	**4.808**	***6.107***	***6.851***	***6.013***	***6.610***	***6.891***
Precious Metal	**9.345**	***6.847***	***7.268***	***8.492***	***5.223***	***6.161***
Coal	**6.688**	***6.688***	***8.454***	***8.498***	***8.717***	***9.532***

Note: This table reports the 0-, 6-, 12-month cumulative effect of oil news shock. The reported estimates are in percent. **Bold** and *italics* represent significance at the 68 and 90 percentiles, respectively.

Table 2: Cumulative responses estimated using a VARX

	Panel A: 1973:1-2019:12			**Panel B: 1973:1-2022:12**		
Industry Portfolios / Months after the shock	0	6	12	0	6	12
Aggregate	0.662	0.112	-0.965	*2.170*	0.922	-0.393
Computer Hardware	*-4.148*	*-4.820*	*-8.071*	*-2.577*	*-3.652*	*-6.063*
Electronic Equipment	*-3.510*	*-4.262*	*-5.883*	**-1.435**	**-2.344**	**-3.683**
Computer Software	*-2.792*	*-2.182*	*-5.456*	**-1.031**	**-1.903**	*-4.753*
Automobiles and Trucks	-0.230	**-2.359**	**-2.815**	*2.633*	-0.423	-1.116
Shipping containers	**-1.230**	**-2.117**	**-2.445**	**0.756**	-1.151	**-2.021**
Retail	**-1.516**	**-1.645**	*-2.595*	**0.894**	-0.221	**-1.565**
Entertainment	0.476	-0.587	-2.619	*1.645*	0.714	-1.094
Consumer Goods	-0.693	-0.091	*-2.218*	0.089	0.185	-1.888
Rubber and Plastic	0.596	0.259	-1.113	**1.094**	-0.479	*-2.167*
Textiles	-0.589	-0.675	-2.327	*2.172*	-0.666	-2.121
Electrical Equipment	0.603	-0.222	-1.370	*1.869*	0.435	-1.378
Restaurants	0.284	-0.347	**-1.460**	*1.804*	-0.017	**-1.656**
Business Services	0.020	-0.104	-1.303	*1.794*	0.399	-1.076
Candy and Soda	**1.839**	0.650	-0.694	*1.686*	0.038	-1.020
Transportation	-0.663	0.051	-1.246	*1.767*	0.684	-0.871
Machinery	**1.235**	-0.062	-1.372	*2.182*	0.220	-1.241
Recreation	**1.569**	1.416	0.099	*3.298*	1.545	0.418
Communication	0.652	0.226	-0.391	*2.305*	0.791	0.172
Food Products	-0.206	-0.406	**-1.461**	*1.562*	0.494	-0.835
Chemicals	0.803	-0.046	-1.474	*2.533*	-0.057	-1.285
Printing and Publishing	**1.188**	0.654	-0.600	*3.502*	**1.959**	-0.198
Beer and Liquor	0.157	0.661	**-1.068**	*1.815*	**1.388**	-0.372
Trading	**1.041**	-0.021	-0.565	*3.249*	1.480	-0.247
Business supplies	*1.546*	**1.730**	-0.533	*2.263*	**1.346**	-0.669
Agriculture	0.736	-0.910	**-2.163**	*1.884*	-0.328	**-2.289**
Aircraft	*2.293*	0.938	0.122	*4.191*	**1.862**	0.272
Measuring and control	0.293	0.718	-0.785	0.783	0.149	-0.813
Apparel	0.812	0.835	-0.416	*2.491*	**1.478**	-0.058
Other	**1.234**	2.307	0.164	*3.006*	**1.749**	0.165
Personal Services	**0.996**	**1.295**	-0.215	*3.052*	*2.099*	-0.062
Tobacco Products	**1.352**	**2.822**	1.704	*2.268*	0.795	0.159
Medical Equipment	*1.582*	1.256	-0.395	*1.971*	1.012	-0.544
Construction Materials	*2.827*	**2.030**	0.630	*4.727*	**1.648**	0.513
Real Estate	*1.515*	2.340	0.882	*4.026*	**2.667**	1.199
Construction	*3.062*	**2.225**	-0.119	*5.459*	**1.835**	0.262
Mines	*4.629*	1.147	-0.143	*4.827*	-0.430	-2.147
Banking	*2.506*	*3.433*	1.807	*4.918*	*3.823*	1.807
Steel works	*2.823*	0.884	-0.612	*4.417*	0.792	-1.281
Shipbuilding and Railroad	*5.684*	*4.447*	**3.052**	*6.066*	*3.857*	**2.021**
Pharmaceutical Products	*2.224*	*2.113*	1.509	*2.695*	*2.680*	**1.334**
Insurance	*3.478*	*4.052*	**2.327**	*5.273*	*4.656*	*3.175*
Wholesale	*2.688*	**2.248**	**1.847**	*3.610*	**1.637**	1.152
Utilities	*3.445*	*3.396*	*3.225*	*4.292*	*3.205*	*2.481*
Fabricated Products	*3.783*	*3.888*	**2.730**	*5.050*	**2.318**	1.181
Defense	*5.656*	*5.860*	*5.842*	*5.512*	*3.505*	*4.392*
Petroleum and Natural Gas	*8.263*	*5.194*	*4.747*	*11.651*	*6.729*	*5.221*
Healthcare	*3.586*	*5.280*	*5.442*	*5.201*	*4.623*	*4.919*
Precious Metal	*8.894*	2.394	**2.897**	*7.274*	1.266	**3.137**
Coal	*5.187*	**5.242**	**5.484**	*6.561*	**5.042**	3.791

Note: This table reports the 0-, 6-, 12-month cumulative effect of oil news shock. The reported estimates are in percent. **Bold** and *italics* represent significance at the 68 and 90 percentiles, respectively.

Figure 1: Response of macroeconomic aggregates to an oil news shock, pre-COVID-19 sample

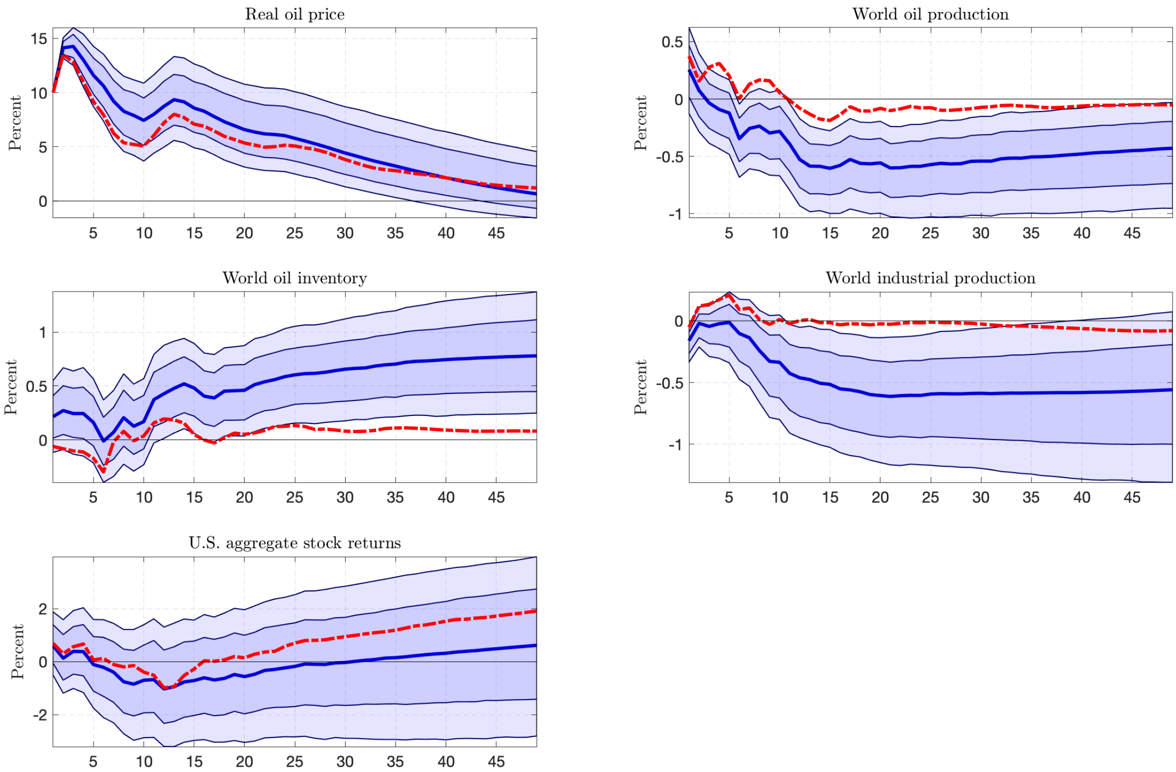

Note: This figure reports impulse responses for the aggregate variables to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line is the response estimated via the proxy-VAR on the 1973:1-2019:12 sample with the IV spanning from 1984:4-2019:12. The shaded regions indicate 68% and 90% confidence intervals. The dash-dotted line indicates the response estimated using a VARX on the 1983:4-2019:12 sample.

Figure 2: Response of industry-level portfolios to an oil news shock, pre-COVID-19 sample

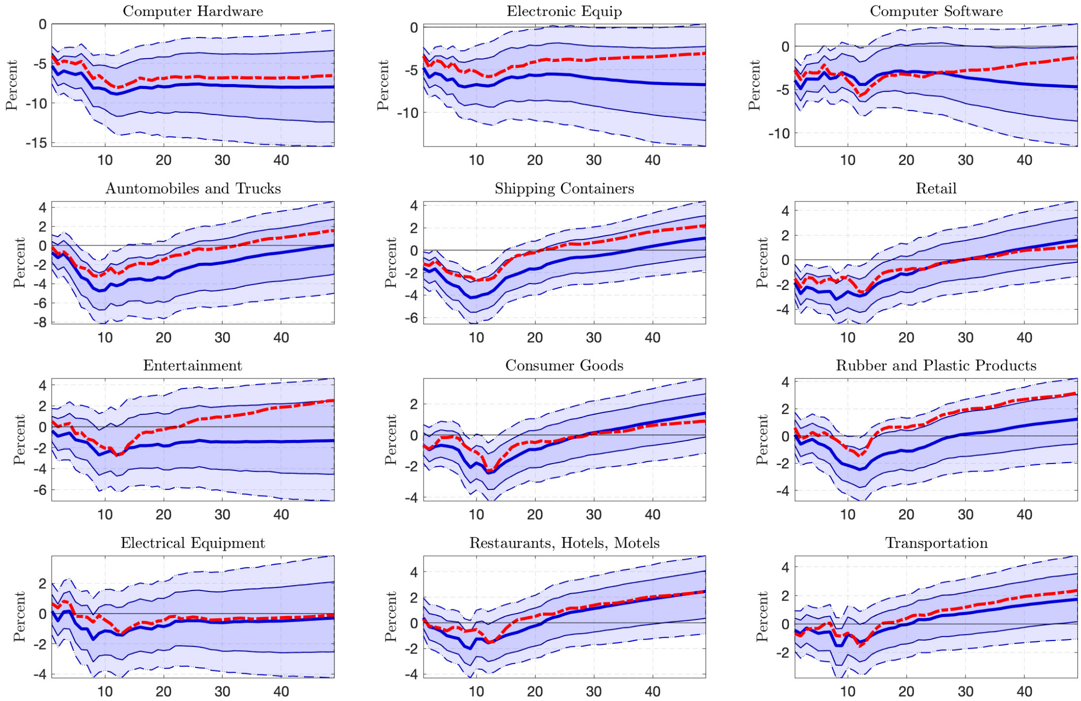

Note: This figure reports impulse responses for industry-level portfolios to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line is the response estimated via the proxy-VAR on the 1973:1-2019:12 sample with the IV spanning from 1984:4-2019:12. The shaded regions indicate 68% and 90% confidence intervals. The dash-dotted line indicates the response estimated using a VARX on the 1983:4-2019:12 sample.

Figure 3: Response of industry-level portfolios to an oil news shock, pre-COVID-19 sample

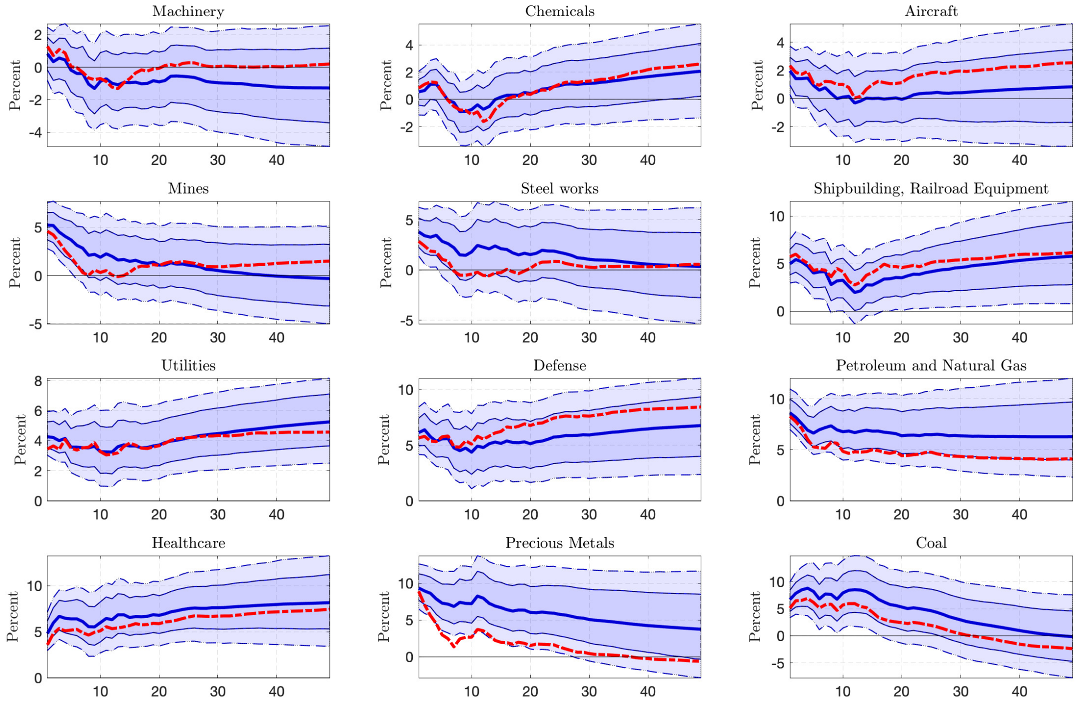

Note: This figure reports impulse responses for industry-level portfolios to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line is the response estimated via the proxy-VAR on the 1973:1-2019:12 sample with the IV spanning from 1984:4-2019:12. The shaded regions indicate 68% and 90% confidence intervals. The dash-dotted line indicates the response estimated using a VARX on the 1983:4-2019:12 sample.

Figure 4: Oil Surprises and Purged Oil Surprises

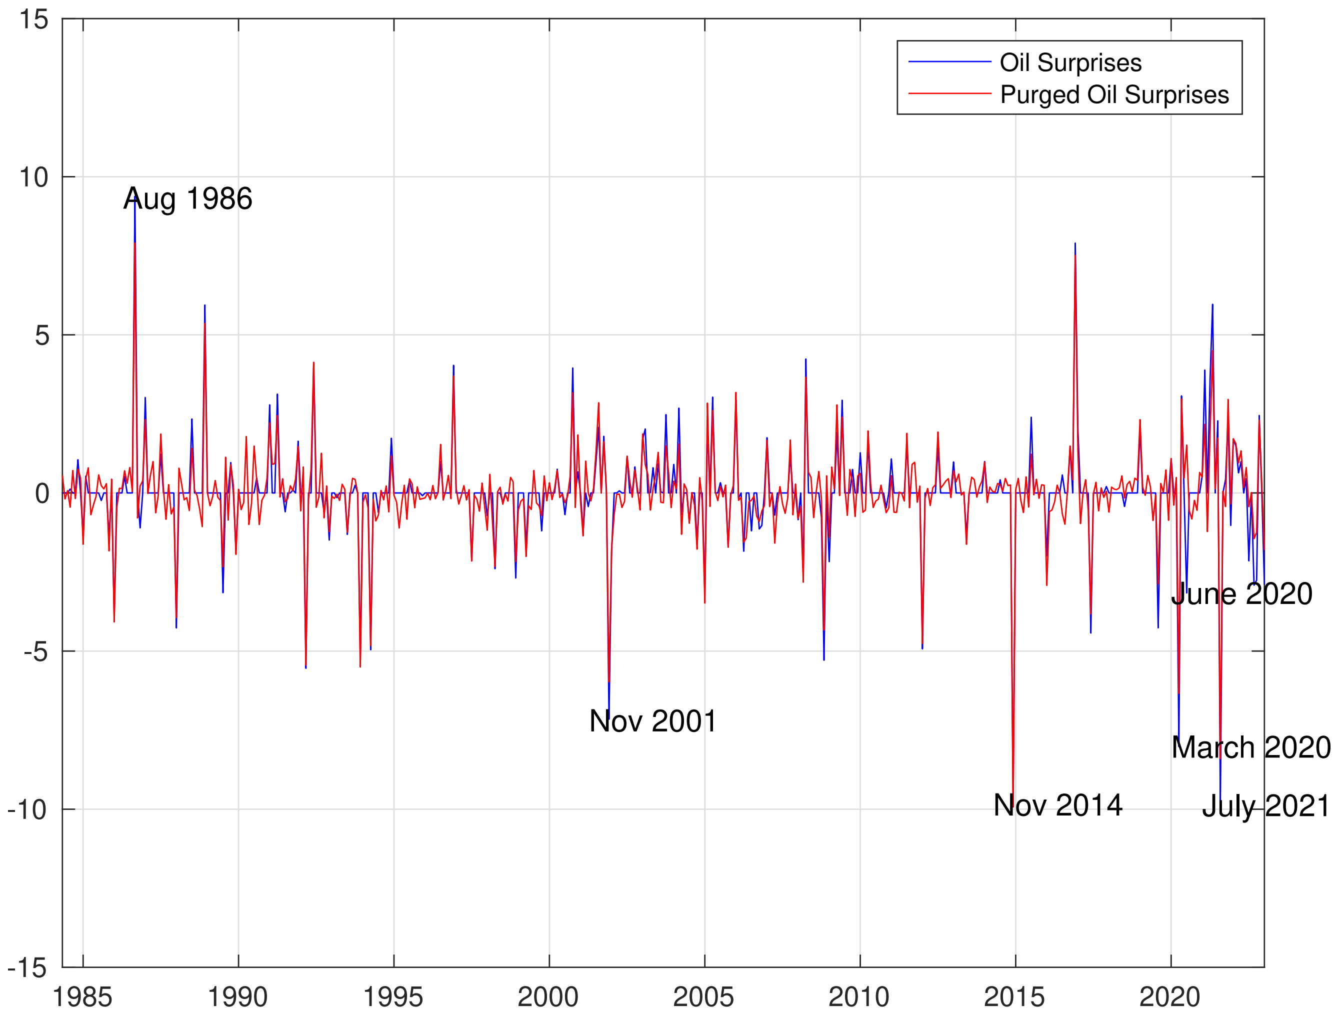

Figure 5: Response of macroeconomic aggregates to an oil news shock, full sample

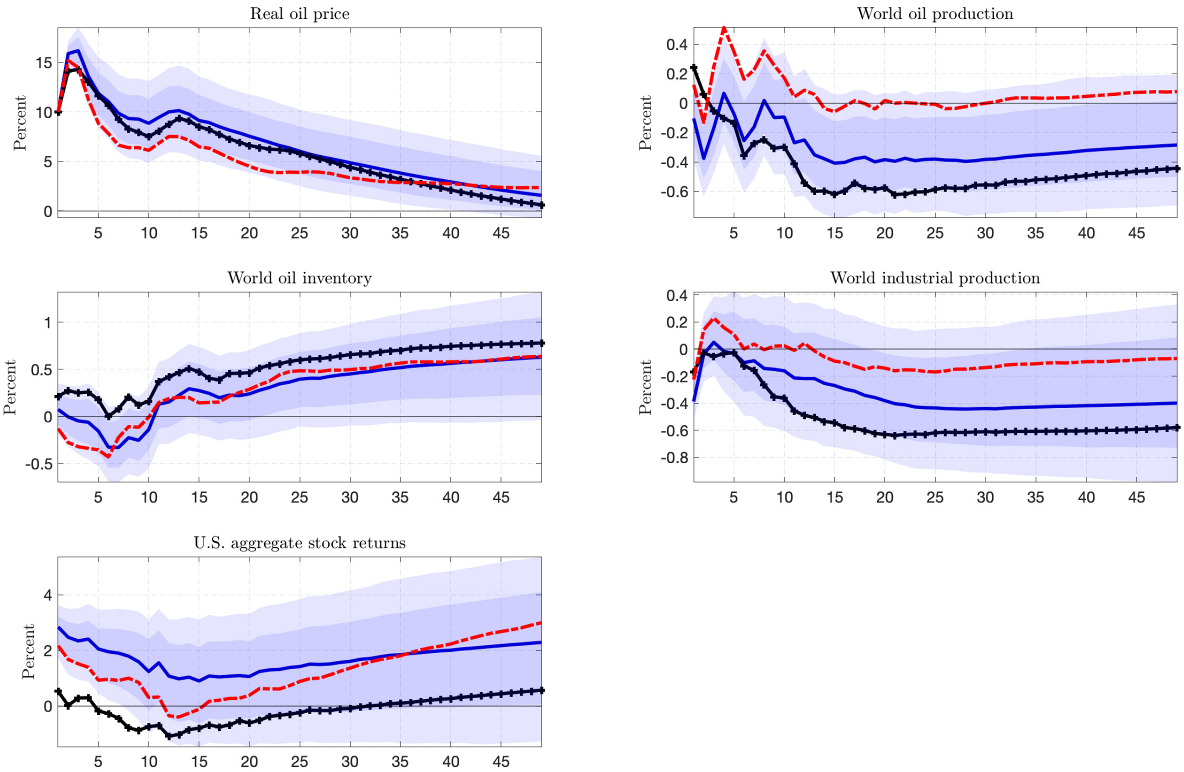

Note: This figure reports the impulse responses for aggregate variables to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line is the response estimated via the proxy-VAR on the 1973:1-2022:12 sample with the IV spanning from 1983:4-2022:12. The shaded regions indicate 68% and 90% confidence intervals. The dash-dotted red line indicates the response estimated using a VARX on the 1983:4-2022:12 sample. The solid black line with a marker indicates the baseline Proxy-VAR model estimates on the 1973:1-2019:12 sample.

Figure 6: Response of industry-level portfolios to an oil news shock, full sample

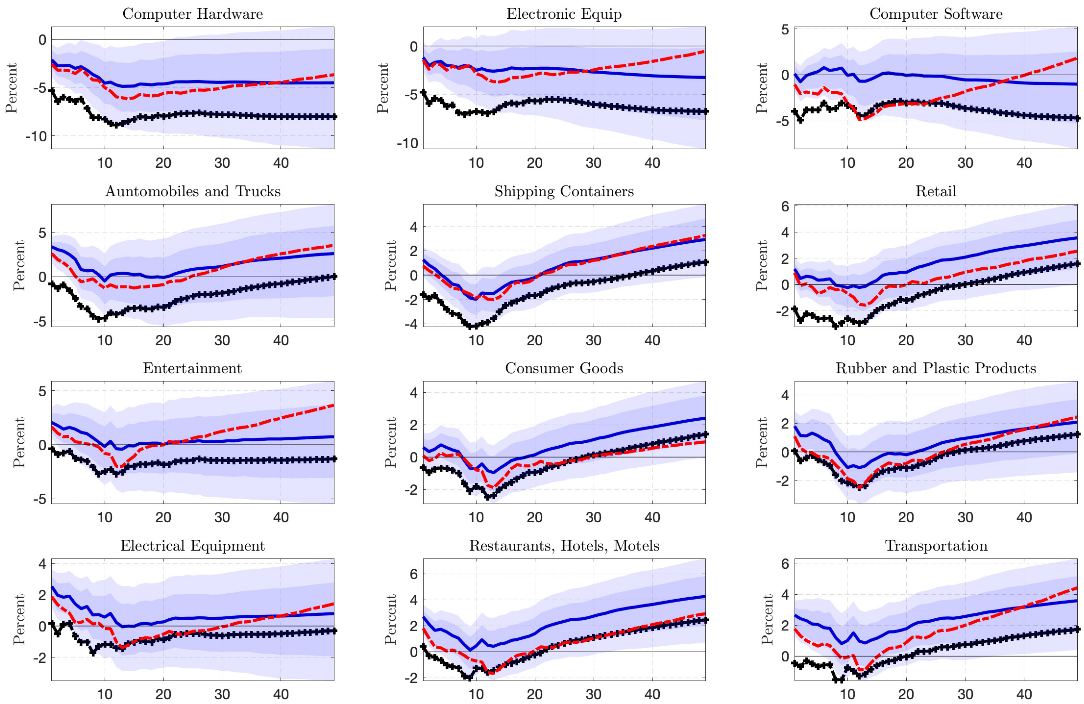

Note: This figure reports the impulse responses for industry portfolio variable to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line is the response estimated via the proxy-VAR on the 1973:1-2022:12 sample with the IV spanning from 1983:4-2022:12. The shaded regions indicate 68% and 90% confidence intervals. The dash-dotted red line indicates the response estimated using a VARX on the 1983:4-2022:12 sample. The solid black line with a marker indicates the baseline proxy-VAR estimates on the 1973:1-2019:12 sample.

Figure 7: Response of industry-level portfolios to an oil news shock, full sample

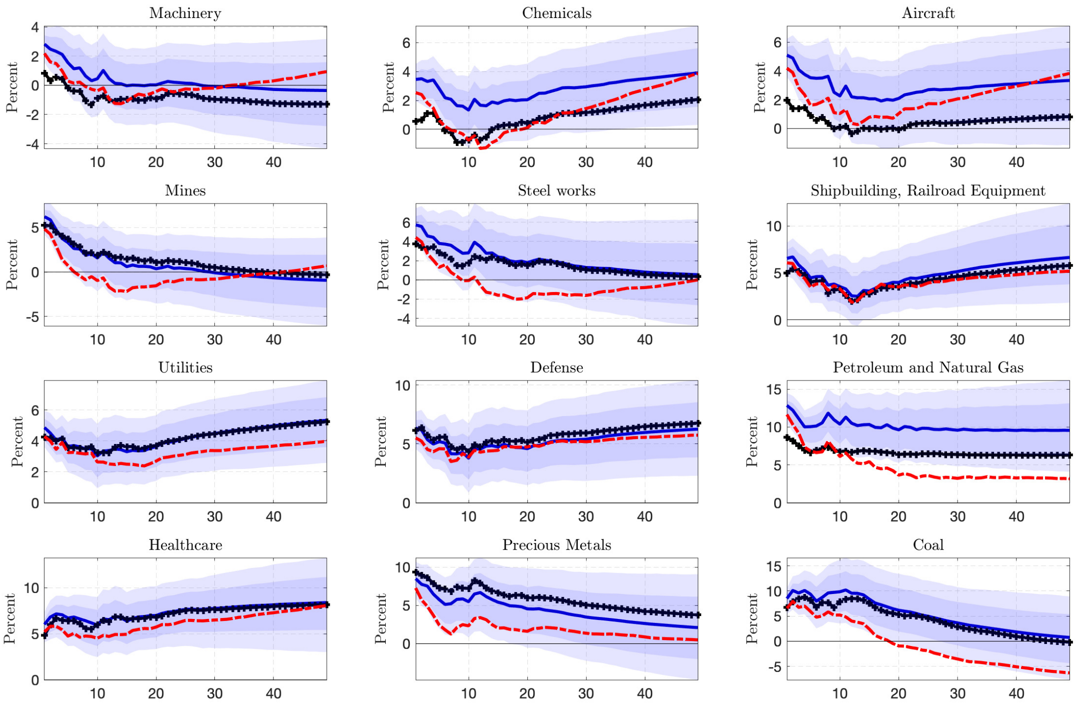

Note: This figure reports the impulse responses for industry portfolio variable to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line is the response estimated via the proxy-VAR on the 1973:1-2022:12 sample with the IV spanning from 1983:4-2022:12. The shaded regions indicate 68% and 90% confidence intervals. The dash-dotted red line indicates the response estimated using a VARX on the 1983:4-2022:12 sample. The solid black line with a marker indicates the baseline proxy-VAR estimates on the 1973:1-2019:12 sample.

Figure 8: Response of macroeconomic variables to an oil news shock - Proxy-VAR specified in rates of growth and alternative instrumental variable, pre-COVID-19 sample

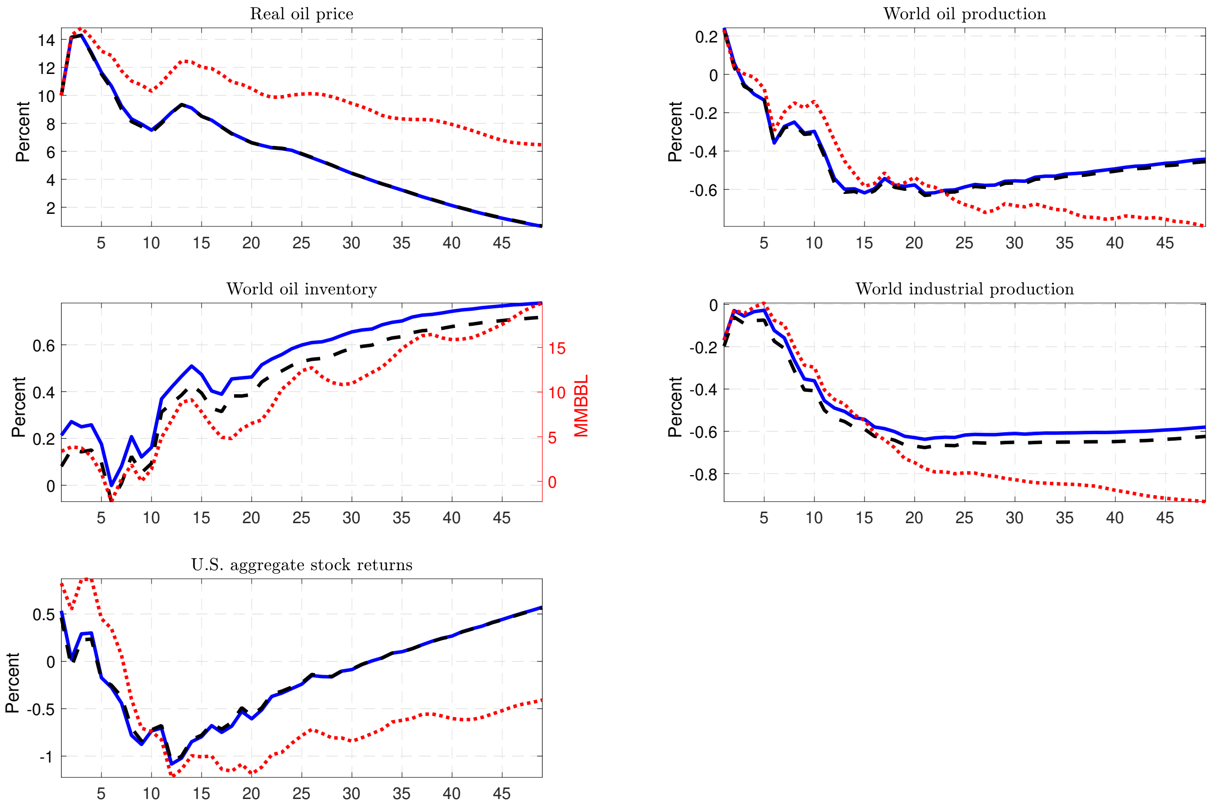

Note: This figure reports the impulse responses for the aggregate variables to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. The solid blue line denotes estimates obtained from baseline specification. The dash black line denotes estimates obtained using a proxy-VAR where the instrument is computed using changes in the six-month oil futures price around OPEC meetings. The dotted red line denotes estimates obtained using a proxy-VAR model specified in rates of growth. The shaded regions indicate 68% and 90% confidence intervals.

Figure 9: Response of industry-level portfolios to an oil news shock - Proxy-VAR specified in rates of growth and alternative instrumental variable, pre-COVID-19 sample

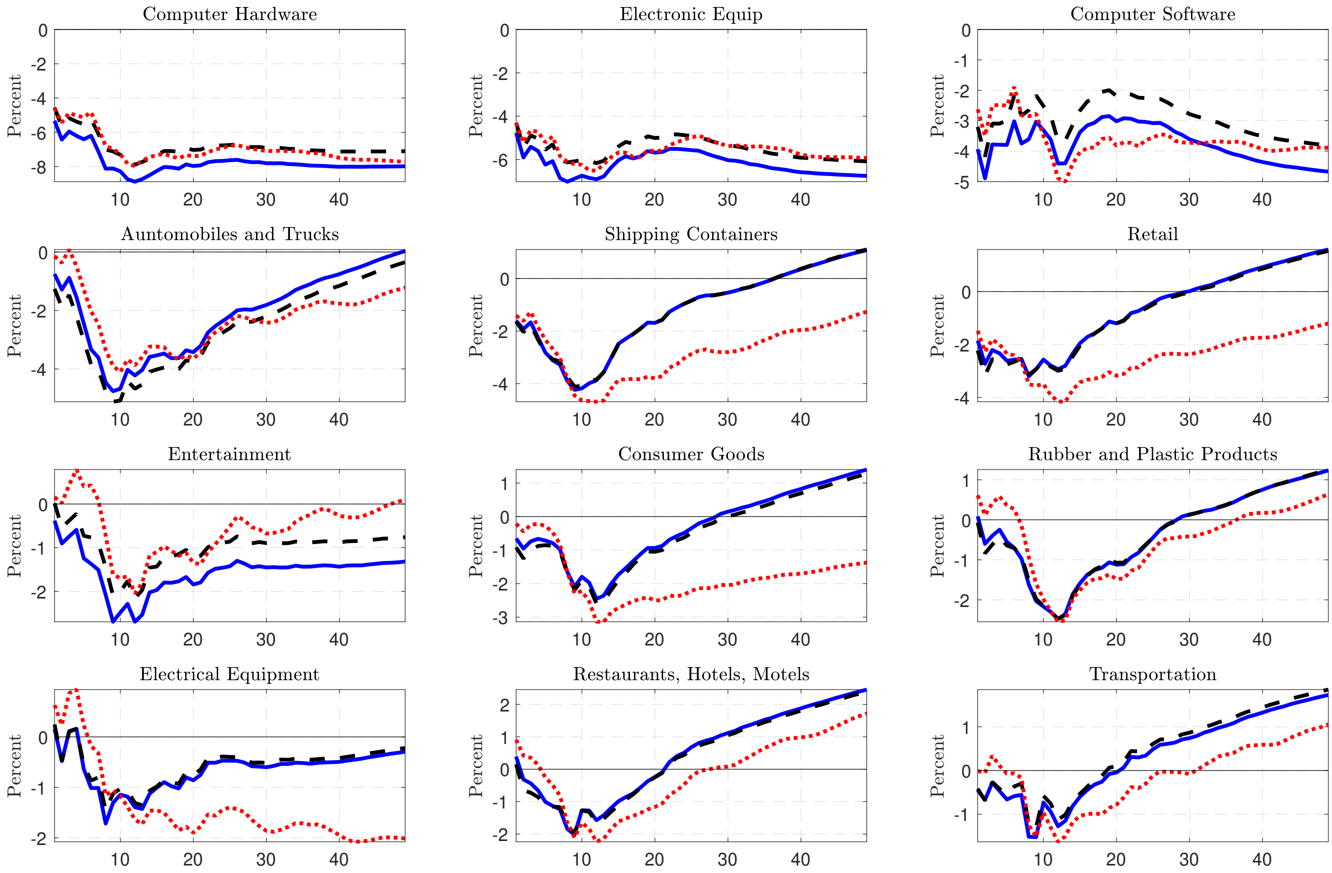

Note: This figure reports the impulse responses for industry portfolio variables to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. The solid blue line denotes estimates from baseline specification. The dash black line denotes estimates obtained using a proxy-VAR where the instrument is computed using changes in the six-month oil futures price around OPEC meetings. The dotted red line denotes estimates obtained using a proxy-VAR model specified in rates of growth. The shaded regions indicate 68% and 90% confidence intervals.

Figure 10: Response of industry-level portfolios to an oil news shock - Proxy-VAR specified in rates of growth and alternative instrumental variable, pre-COVID-19 sample

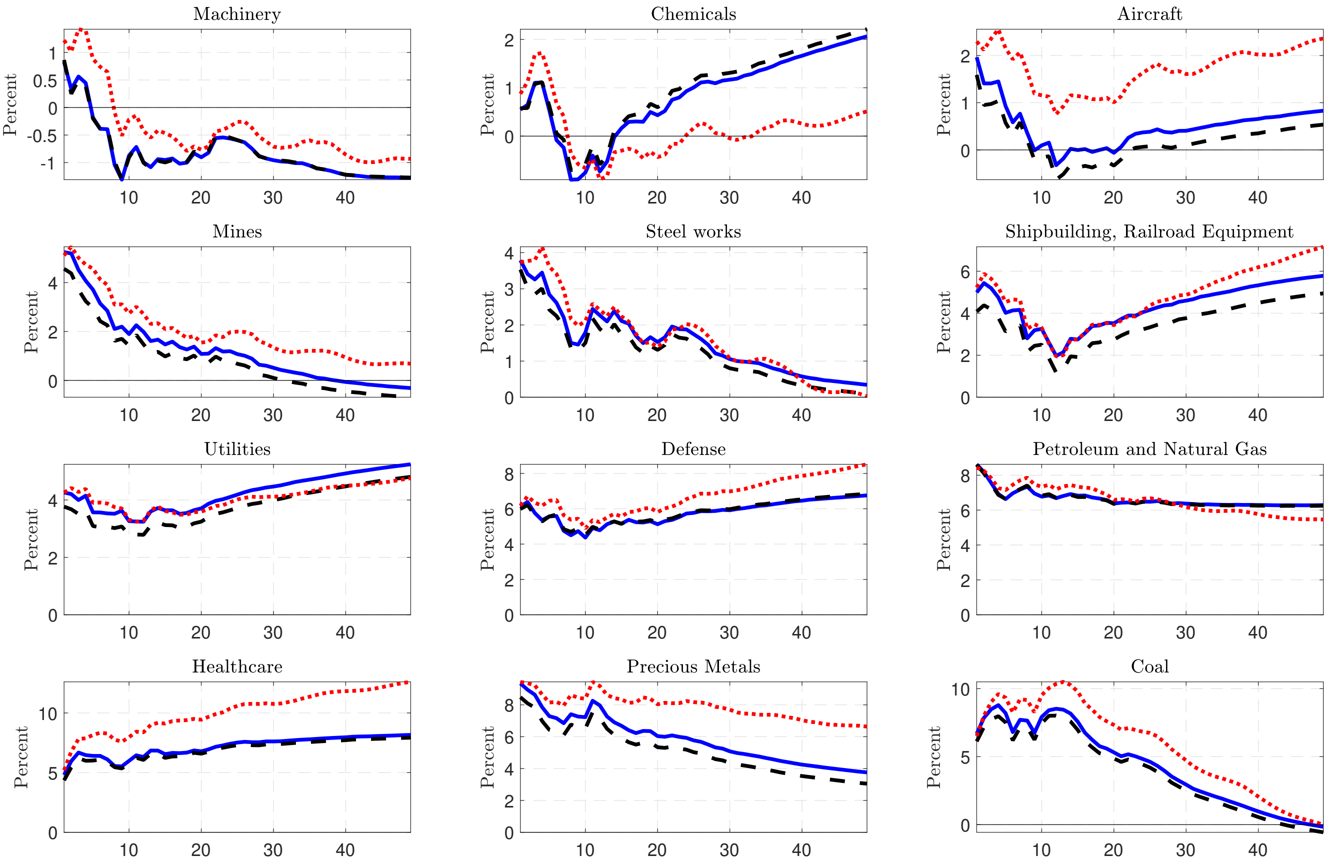

Note: This figure reports impulse responses for industry portfolio variables to an oil news shock, normalized to increase the real price of oil by 10 percent on impact. All responses are reported in percentages. The solid blue line denotes estimates from baseline specification. The dash black line denotes estimates obtained using a proxy-VAR where the instrument is computed using changes in the six-month oil futures price around OPEC meetings. The dotted red line denotes estimates obtained using a proxy-VAR model specified in rates of growth. The shaded regions indicate 68% and 90% confidence intervals.