Asymmetric Transmission of Oil Supply News*

Mario Forni
Università di Modena e Reggio Emilia, CEPR and RECent

Alessandro Franconi†
Università di Pavia

Luca Gambetti‡
Universitat Autònoma de Barcelona, BSE, Università di Torino, CCA

Luca Sala
Università Bocconi, IGIER, Baffi CAREFIN

Abstract

We provide new evidence on the asymmetries in the transmission of oil supply news shocks in the US using a nonlinear Proxy-SVAR. A shock that increases oil prices has large and persistent effects on real activity and relatively small effects on prices. On the contrary, a shock that reduces oil prices has smaller real effects and large effects on prices. We rationalize these findings through the behavior of uncertainty: uncertainty increases independently of the sign of the shock, amplifying the contractionary real effects of a positive shock and dampening the expansionary real effects of a negative shock. The opposite holds for prices. We find little evidence of an asymmetric response of monetary policy.

JEL codes: C32, E31, E32, Q43.
Keywords: Oil Supply News, Nonlinear Proxy-SVAR, Asymmetry.

---

*The authors acknowledge the financial support of the Italian Ministry of Research and University, PRIN 2017, grant J44I20000180001. This paper has benefited from discussions with participants at the 2nd Annual"Non-Linearities in Macro" Workshop (Bank of England) and the 3rd Sailing the Macro Workshop.

†The author has started working on the project at the Università di Modena e Reggio Emilia.

‡Luca Gambetti acknowledges the financial support from the Spanish Ministry of Science and Innovation, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), the financial support from the Spanish Ministry of Science, Innovation and Universities and FEDER through grant PGC2018- 094364-B-I00 and the BSE Research Network.

1 Introduction

Recent developments in the global economy have led to unprecedented decisions by the largest oil producing countries. In April 2020, Russia, Saudi Arabia and the United States jointly cut production by 9.7 million barrels per day (bpd) to counter the negative effects of the COVID-19 pandemic. By mid-2021, a strong recovery in global demand led the Organization of the Petroleum Exporting Countries (OPEC) to ease earlier cuts by 2 million bpd. But in 2022, Russia, the world's largest oil exporter to global markets, caused supply disruptions in the oil market following its invasion of Ukraine. Later that year, OPEC reversed its stance and announced oil supply cuts of 2 million bpd for all of 2023, which were increased to 3.66 million bpd in April 2023, or about 3.7% of global demand. More recently, in an effort to support prices, OPEC and its partners (OPEC+) decided to extend oil cuts through 2024, and Saudi Arabia and Russia unilaterally cut oil production by an additional 1.3 to 1.5 million bpd until the end of December.[1] From a policy perspective, these large oil supply cuts and easings pose new challenges for stabilization policies, motivating a renewed interest in understanding the transmission of oil supply shocks in order to promote better-informed policy decisions.[2]

Our understanding of the relationship between oil prices and the macroeconomy goes back to Hamilton (1983), who initiated a long-standing debate by arguing that oil supply shocks are a major driver of economic fluctuations. More recently, the debate has focused on potential nonlinearities in the transmission of oil price changes; in particular, whether oil price increases have a greater impact on real activity than oil price decreases. The debate can be traced back to the 1980s, when Hamilton (1983) and Mork (1989) observed that oil price increases seemed to be more important for US business cycles than price decreases. More recently, using US data, Hamilton (2003, 2011) finds evidence in favor of the asymmetry; in contrast, Kilian and Vigfusson (2011a, 2017) criticize Hamilton's empirical approach and provide evidence against such nonlinearities. Other contributions in support of a linear relationship between aggregate activity and oil price shocks are Herrera et al.

[1] OPEC+ controls about 90% of global crude oil reserves and 40% of the global daily production. Therefore, OPEC+ decisions can have a significant impact on oil prices.

[2] In our analysis, we end the sample in 2019M12 to avoid including the COVID period in the sample.

(2011) for the US and Herrera et al. (2015) for 18 OECD countries.[3] Lastly, a recent work by Caravello and Martinez-Bruera (2024) finds no evidence of asymmetric effects, suggesting instead the existence of size effects.

In this paper, we contribute to this debate by exploring whether oil supply news shocks have different effects on US output and prices, depending on the sign of the shock. To investigate potential asymmetries we use the nonlinear Proxy-SVAR approach developed by Debortoli et al. (2023) (DFGS henceforth), and identify an oil supply news shock. The instrument we use is the series of surprise changes in oil futures prices around OPEC announcements, developed by Känzig (2021), who show that this variable affects oil prices and the macroeconomy.

The underlying economy is represented by a structural Vector Moving Average which includes nonlinear terms of the shock of interest; here we use the absolute value of the oil shock as the relevant nonlinear function. Under suitable conditions, the macroeconomic variables have a VARX representation, where the shock and its absolute value represent the exogenous variables. By combining the effects of the shock and its absolute value we can estimate the effects of positive and negative shocks. This is a basic difference with respect to the previous literature on this subject: we focus on the nonlinear effects of unobservable oil supply news shocks, rather than observable oil price *changes*. Since the exogenous variables are unobserved, our VARX cannot be estimated directly. The key result of DFGS solves this problem. DFGS shows that, when a suitable instrument is available and the observed variables are informationally sufficient for the shock of interest, the shock can be estimated consistently as the fitted value of the regression of the instrument onto the residuals of a standard linear VAR. Once an estimate of the shock is available, the VARX and the implied nonlinear impulse response functions can be estimated.

In comparison with alternative nonlinear models we find our method particularly appealing for a number of reasons. First, the alternative nonlinear Proxy-SVAR implemented in Pellegrino et al. (2023) allows the estimation of impulse responses across states of the economy but does not have the flexibility of dealing with shock sign or size nonlinearities. Second, in comparison with the FAIR approach of Barnichon and Matthes (2018a), we do not need to assume a specific distribution of the

[3] While both provide evidence against nonlinearity, the former finds a strong asymmetric transmission of oil price shocks in the US at the disaggregated industry level.

shock. Compared to a nonlinear local projection, such as the one in Tenreyro and Thwaites (2016), the benefit of our model is twofold. First, the more parsimonious parameterization and consequently the smaller estimation uncertainty. A second advantage is *internal consistency* as our approach allows the shock identification and the estimation of nonlinear impulse responses in a unique model.

We find that the transmission of oil supply news shocks is asymmetric. A shock raising oil prices produces a large and immediate decline in real activity and a small increase in prices. On the other hand, a shock reducing oil prices has a modest effect on real activity and a large effect on prices. This evidence confirms the findings of Hamilton (2011), but also suggests an additional nonlinearity, previously unexplored, concerning prices. These results are robust to an alternative identification strategy of oil supply shocks (Baumeister and Hamilton, 2019) and to various changes in the model specification. Our findings, however, contrast with Caravello and Martinez-Bruera (2024), who find evidence against asymmetric effects of oil supply news shocks using Känzig's proxy variable in a local projection. In Section 2.4 we argue that these contrasting results arise from the different methods used, with their approach potentially underestimating the asymmetry we document. We further elaborate on this issue in Appendix A using a theoretical argument and a simulation exercise.

The existing literature has suggested two possible explanations for the asymmetric effects of oil shocks. The first is related to uncertainty. We find that oil supply shocks, either positive or negative, increase uncertainty. Higher uncertainty, in turn, increases the returns to waiting for information, causing firms to delay their investment plans, i.e., the "real option" effect (Bernanke, 1983; Bloom, 2009); moreover, it raises the risk of investment and therefore the cost of financing, especially for risky firms, i.e., the "risk premium" effect, (Christiano et al., 2014; Gilchrist et al., 2014). This channel is in line with Elder and Serletis (2010) and Kilian and Vigfusson (2011b). Specifically, the latter points out that *"Because any unexpected change in the real price of oil may be associated with higher expected volatility, whether the real price of oil goes up or down, this uncertainty effect may amplify the effects of unexpected oil price increases and offset the effects of unexpected oil price declines"* (p. 340). The second explanation focuses on the response of monetary policy. The central bank may react to oil price increases by raising the interest rate to contain inflationary pressures, but may decide not to respond to price reductions, therefore generating real asymmetries. Bernanke et al. (1997) suggest that *"...the endogenous monetary policy response can account for a very substantial portion of the depressing effects of oil price shocks on the real economy"* (p. 94).

To assess these potential channels, we extend the model to include three additional variables: the federal funds rate, the VXO uncertainty index used by Bloom (2009) and the excess bond premium of Gilchrist and Zakrajšek (2012). Consistent with the uncertainty channel above, we find an increase in uncertainty after both positive and negative shocks. Moreover, these effects are associated with an increase in the excess bond premium, suggesting that the "risk premium" effect is at work. These results provide strong evidence in favor of the uncertainty channel. As for the second channel, our estimates show that the federal funds rate rises after both shocks. However, these effects are small and not significant, suggesting that monetary policy cannot be the main source of asymmetry.

The remainder of this paper is organized as follows. Section 2 outlines the econometric model, including the identification of the shock and the estimation of nonlinear impulse responses. Section 3 discusses the main empirical results of the baseline specification and presents several robustness checks. Section 4 explores potential channels responsible for our findings. Finally, section 5 concludes.

2 Econometric methodology

In this section we discuss the main features of the econometric methodology of DFGS, adapted to study the nonlinear effects of oil supply shocks.

2.1 Nonlinear representation

Let x_t be a n-dimensional vector of stationary macroeconomic variables with the following structural representation

$$x_t = \nu + \alpha(L)u_t^s + \beta(L)g(u_t^s) + \Gamma(L)\xi_t, \quad (1)$$

where ν is a vector of constants, u_t^s is the oil supply news shock with impulse response functions $\alpha(L) = \alpha_0 + \alpha_1 L + \alpha_2 L^2 ...$ and $g(u_t^s)$ is a nonlinear function of the oil supply news shock with impulse response functions $\beta(L) = \beta_0 + \beta_1 L + \beta_2 L^2 + ...$ The m-dimensional vector ξ_t includes other structural shocks, collected in vector u_t^{-s}, and possibly nonlinear functions of these shocks. The vector $[u_t^s, u_t^{-s}]'$ is $i.i.d.$ zero mean and has identity covariance matrix.[4] Finally, the matrix $\Gamma(L) = \Gamma_0 + \Gamma_1 L + \Gamma_2 L^2 + ...$ is a $n \times m$ matrix of impulse response functions to the remaining structural shocks and, possibly, their non-linear functions.

The nonlinear impulse response functions to the oil shock are derived by combining the two terms $\alpha(L)$ and $\beta(L)$. More specifically, the total effects of an oil supply news shock $u_t^s = \bar{u}^s$ are given by the sum of the linear and nonlinear terms:

$$IRF(u_t^s = \bar{u}^s) = \alpha(L)\bar{u}^s + \beta(L)g(\bar{u}^s). \tag{2}$$

The total responses defined in equation (2) simply correspond, in this nonlinear context, to the Generalized Impulse Response Functions defined as $E(x_{t+h}|u_t^s = \bar{u}^s) - E(x_{t+h}|u_t^s = 0)$, $h = 0, 1, ...$ We discuss below how to estimate the model and the implied impulse response functions.

In our baseline empirical specification we use the absolute value as the relevant nonlinear function, i.e. $g(u_t^s) = |u_t^s|$, since our interest is in the potential sign asymmetries of the shock. The impulse response functions to a positive shock equal to 1 and a negative shock equal to -1 are, respectively

$$\begin{aligned} IRF(u_t^s = 1) &= \alpha(L) + \beta(L) \\ IRF(u_t^s = -1) &= -\alpha(L) + \beta(L). \end{aligned} \tag{3}$$

If $\beta(L) \neq 0$ the shocks of opposite signs have different effects.

Stationarity of the term $\Gamma(L)\xi_t$ implies the existence of the Wold representation $\Psi(L)e_t$. Under the assumption of invertibility of the Wold representation, model (1)

[4] Notice that shock serial and mutual independence implies that all structural shocks, including u_t^s, are uncorrelated with the lags of $g(u_t^s)$ and x_t.

implies the existence of the following VARX representation:[5]

$$A(L)x_t = \mu + \tilde{\alpha}(L)u_t^s + \tilde{\beta}(L)g(u_t^s) + e_t, \tag{5}$$

where $A(L) = \Psi(L)^{-1}$, $\tilde{\alpha}(L) = \Psi(L)^{-1}\alpha(L)$ and $\tilde{\beta}(L) = \Psi(L)^{-1}\beta(L)$.[6]

Notice also that stationarity of x_t implies the existence of its Wold representation. Under the assumption that such representation is invertible, the vector x_t admits the VAR representation

$$x_t = \vartheta + B(L)x_{t-1} + \varepsilon_t \tag{6}$$

where ε_t is orthogonal to x_{t-j}, $j = 1, \ldots, \infty$.[7]

Next we derive the relation between the VAR representation (6) and the VARX representation (5). Let us start from equation (5) and consider the linear projection of $\tilde{\alpha}(L)u_t^s + \tilde{\beta}(L)g(u_t^s)$ onto the constant and the past history of x_t, i.e.

$$\tilde{\alpha}(L)u_t^s + \tilde{\beta}(L)g(u_t^s) = \theta + C(L)x_{t-1} + w_t.$$

It is easily seen that $\vartheta = \mu + \theta$, $B(L) = \tilde{A}(L) + C(L)$, where $\tilde{A}(L) = -[A(L) - I]/L$ and $\varepsilon_t = e_t + w_t$.

Notice that, if $\tilde{\beta}(L) = 0$, the structural representation (1) reduces to a linear model and standard SVAR analysis can be conducted using representation (6). Hence the linear model is nested in our model. We can test for linearity by testing either for the null $\tilde{\beta}(L) = 0$ in equation (5) or for the null $\beta(L) = 0$ in the impulse response functions in (2). In the empirical application below we test for linearity by adopting the null hypothesis $H_0 : \beta(L) = 0$.

[5]This is simply obtained from

$$x_t \quad = \quad \nu + \alpha(L)u_t^s + \beta(L)g(u_t^s) + \Psi(L)e_t, \tag{4}$$

and then multiplying by $\Psi(L)^{-1}$

[6]We assume that all the matrices of polynomials in L can be approximated by finite order matrix polynomials, as it is standard in the VAR literature.

[7]The cointegration case can be treated as usual by considering a VAR in the levels of the variables, rather than in first differences.

2.2 Identification

In the previous subsection we have shown that our nonlinear economy admits a VARX representation (5). Unfortunately, direct estimation of (5) is unfeasible, because the exogenous variables are not observable.[8] We discuss below how to obtain a consistent estimate of the exogenous shocks that can be used to estimate the VARX.

The identification procedure relies on two assumptions. The first assumption is standard in the proxy-SVAR literature and requires the existence of a valid instrument, as specified below.

Assumption A1 (*Proxy*). The proxy z_t is given by

$$z_t = a + bu_t^s + \delta(L)' x_{t-1} + v_t, \tag{7}$$

where $\delta(L)$ is a vector of polynomials of degree p in the lag operator L, $b \neq 0$ and v_t is an error independent of the structural shocks at all leads and lags. Notice that under Assumption A1, the standard conditions for a valid instrument,[9] i.e. $\text{cov}(z_t, u_t^s) = b \neq 0$ (relevance) and $\text{cov}(z_t, \xi_t) = 0$ (exogeneity), are satisfied.

The second assumption ensures that the oil supply news shock can be estimated as a combination of current and past data.

Assumption A2 (*Informational sufficiency*). The oil supply news shock is a linear combination of the current and past values of x_t.

Assumption A2 postulates "partial invertibility" of u_t^s, i.e. that the variables in x_t are informationally sufficient for the oil shock.[10] In other words, the nonlinear term is not needed to estimate the oil supply news shock. Notice that the same assumption has to hold also in the linear case in order for the standard procedure to be valid. Fortunately this is a testable assumption. In the empirical section we will assess whether it holds.

[8]If the oil shock were perfectly observable, then eq. (5) or a local projection version of it could be estimated by OLS. In section 2.4 we discuss why such a procedure could be problematic if only an imperfect measure of the shock is available.

[9]See Mertens and Ravn (2013) and Stock and Watson (2018a).

[10]On the concept of informational sufficiency see Forni and Gambetti (2014) and Forni et al. (2019).

Under assumptions A1 and A2, DFGS show that the shock of interest can be obtained as the fitted value of the linear projection of z_t onto the VAR innovations ε_t. This is the basic result underlying the proposed procedure. First, we get an estimate of the oil shock by using a standard proxy SVAR method; having an estimate of the shock, we can use it to estimate our VARX representation (5).

2.3 Estimation

More in detail, the estimation procedure is the following.

I. Estimate the VAR in (6) with OLS to obtain consistent estimates of the residuals ε_t, call them $\hat{\varepsilon}_t$.

II. Estimate the linear projection

$$z_t = \lambda' \hat{\varepsilon}_t + \eta_t. \tag{8}$$

Following Forni et al. (2023), an estimate of the normalized shock is obtained as follows

$$\hat{u}_t^s = \frac{\hat{\lambda}' \hat{\varepsilon}_t}{\sqrt{\hat{\lambda}' \widehat{\Sigma}_\varepsilon \hat{\lambda}}} \tag{9}$$

where $\widehat{\Sigma}_\varepsilon$ is the variance covariance matrix of the residuals $\hat{\varepsilon}_t$.[11]

III. Estimate equation (5) using as regressors the current value and the lags of the estimated shock $\hat{u}_t^s$ and its nonlinear function $g(\hat{u}_t^s)$. This gives the estimates of $A(L)$, $\tilde{\beta}(L)$ and $\tilde{\alpha}(L)$. Finally, one can obtain estimates of $\alpha(L) = A(L)^{-1}\tilde{\alpha}(L)$ and $\beta(L) = A(L)^{-1}\tilde{\beta}(L)$.

IV. Compute the impulse response functions according to equation (3).

2.4 Discussion

To verify the existence of nonlinearities, Hamilton (2011) uses a censored measure of oil price changes, named *net oil price increases* (see also Hamilton, 1996). This

[11] In equation (9), the covariance matrix Σ_ε is estimated over the full sample. The covariances between z_t and $\hat{\varepsilon}_t$ is of course estimated over the sample for which the instrument is available.

measure distinguishes between oil price increases that set new highs relative to recent experience and those that simply reverse recent declines, reflecting the apparently greater importance of oil price increases for business cycles. Kilian and Vigfusson (2011a) argue that if the conditional forecasting model, e.g., the VAR, involves a nonlinear function such as the *net oil price increases*, then we cannot calculate a multi-period impulse response by iterating as if it was a linear model. In other words, if one includes in a VAR price increases but not decreases, the usual inference will be biased: the econometrician should ensure that the model used nests the linear case when testing for asymmetries. Notice that our approach is not subject to the criticism above, since in our VARX specification we have both the relevant shock and its absolute value, which is equivalent to having both positive and negative shocks (as observed above, the linear case is a special case of our model).

A basic difference of our approach with respect to the previous literature is that we focus on possible nonlinear effects of oil *shocks* rather than oil price *changes*. From this point of view, our approach is perfectly in line with the structural VAR literature, where the observed variables are *endogenous* and are driven by *unobserved* exogenous forces.[12] Our first step is needed precisely because the oil shock is unobserved and therefore must be estimated. Once the shock has been estimated, a valid alternative delivering the same asymptotic result is to use local projections in place of the VARX to estimate the impulse response functions.

Discrepancies relative to other work. Caravello and Martinez-Bruera (2024) finds no evidence of asymmetric effects of oil supply news shocks. In that paper, the authors skip the first step of our procedure and use the proxy of Känzig (2021) and its absolute value directly in a series of local projections in place of the estimated shock and its absolute value. To verify whether the conflicting results are due to this, we applied their methodology to our dataset, and found little evidence of asymmetry, in line with their result (see Appendix B). Below, we provide the intuition for the different results obtained by these authors.

The methodology of Caravello and Martinez-Bruera (2024) is appealing in that it is simple and does not require invertibility. However, their approach in a nonlinear

[12] In this respect, our nonlinear approach is similar to the one of Barnichon and Matthes (2018b), even if the methods are different.

framework delivers the correct result only when the instrument is an exact measure of the shock, $z_t = u_t^s$, or when the difference $z_t - u_t^s$ has the same effects as u_t^s on all variables. For instance, consider the case $u_t^s = u_{1t}^s + u_{2t}^s$, where u_t^s is the sum of two structural shocks, u_{1t}^s and u_{2t}^s and z_t captures only the first term, $z_t = u_{1t}^s$. If the effects of u_{1t}^s and u_{2t}^s are the same, no problems arise.[13] By contrast, if the effects of u_{1t}^s and u_{2t}^s are different, then using z_t in place of u_t^s will produce a bias.

A bias will also arise if z_t is a noisy measure of the shock with an error which is unrelated to the variables, as in equation (7). Consider the simple case $z_t = u_t^s + v_t$, where v_t is independent of the variables at all leads and lags. In the linear case, v_t produces an attenuation bias, which is proportional for all variables and lags, and therefore can be corrected by a suitable normalization (see Stock and Watson, 2018b). In the nonlinear case, however, things are more complicated. Assume that negative shocks, say u_t^{s-}, do not affect output, whereas positive shock, u_t^{s+}, have effect 1, so that the effects are asymmetric. Clearly z_t can be positive even if u_t^s is negative, owing to a positive v_t greater than $|u_t^s|$; hence z_t^+ is a mixture, including both positive and negative shocks. As a consequence, using z_t^+ in place of u_t^{s+} produces a downward bias (in addition to the attenuation bias) and the estimated effect on output will be somewhere in the interval (0 1). For the same reason, we have an upward bias when using z_t^- in place of u_t^{s-}, so that the asymmetry will be underestimated. This confounding effect could in principle explain why Caravello and Martinez-Bruera (2024)'s finding of no asymmetric effects conflicts with ours. In Appendix A we present a simple theoretical example together with a simulation exercise to further elaborate on this point.

3 Results

The baseline monthly linear VAR(12) follows the specification used by Känzig (2021) and includes the real oil price, world oil production, world oil inventories, world industrial production, US industrial production and the US consumer price index (CPI) from 1975M1 to 2019M12.[14] All variables enter in logs. After estimating the model, we identify the structural oil supply news shock over the shorter sample

[13] We thank an anonymous referee for pointing this out.

[14] For the data sources, see Känzig (2021).

1983M4-2019M12, using the proxy constructed by Känzig (2021). This variable captures high-frequency surprise changes in oil futures prices in a tight window around OPEC announcements and can only be considered a noisy measure of the true shock. The choice of estimating the reduced-form model using the longer sample is done to improve the estimation of the parameters of the model, as typically done in the Proxy-SVAR literature. Regarding the strength of the instrument, the first-stage F-statistic in this model is 17.17, which is safely above the threshold of 10.

3.1 Testing for invertibility

To begin, we test whether assumption A2, i.e. invertibility of u_t^s, holds. To do this, we use the invertibility test recently proposed by Forni et al. (2023). The test is based on the theoretical result that, if the shock is non-invertible, then it is a function of current *and future* VAR residuals, instead of being a combination of current residuals only. More specifically, the test is based on regressing the instrument on the current value and the first r leads of the Wold residuals $\hat{\varepsilon}_t$.[15] Formally:

$$z_t = \sum_{k=0}^{r} \lambda_k' \hat{\varepsilon}_{t+k} + \eta_t \tag{10}$$

The invertibility test is an F-test for the significance of the r leads, the null hypothesis being $H_0 : \lambda_1 = \lambda_2 = \cdots = \lambda_r = 0$ against the alternative that at least one of the coefficients is nonzero. We estimate the regression in equation (10) using different numbers of leads ($6 \leq r \leq 12$).

The p-values, reported in Table 1, are very large. Therefore, we cannot reject the null of invertibility for all values of r. In the robustness section 3.6 we briefly discuss the conflicting results of Plagborg-Møller and Wolf (2022, Online Appendix B.7).

[15] An alternative with standard VAR identification schemes can be found in Forni and Gambetti (2014).

	Number of leads r						
	$r = 6$	$r = 7$	$r = 8$	$r = 9$	$r = 10$	$r = 11$	$r = 12$
p-value	0.98	0.99	0.95	0.96	0.91	0.88	0.85

Table 1: Invertibility test. The table shows the p-values for each regression including the current value and up to r leads of the Wold residuals. The null hypothesis is invertibility, i.e., $H_0 : \lambda_1 = \lambda_2 = \cdots = \lambda_r = 0$.

3.2 Linear model

We first present the results for the identification of the oil supply news shock in the linear Proxy-SVAR. Figure 1 plots the impulse responses to an oil supply news shock normalized to increase the real oil price by 10% on impact. The black solid lines are the point estimates and the grey shaded areas are the 68% and 90% confidence bands. A negative oil supply news shock leads to an immediate increase in real oil prices. World oil production falls persistently only after few months and world oil inventories increase significantly at impact and continue to grow sluggishly. World industrial production does not change much in the first year after the shock, but then begins to fall significantly and persistently. For US variables, the shock leads to a delayed and persistent decline in industrial production and an immediate increase in the consumer price index, which continues to rise for a year before returning to its initial level. These results are consistent with those in Känzig (2021).

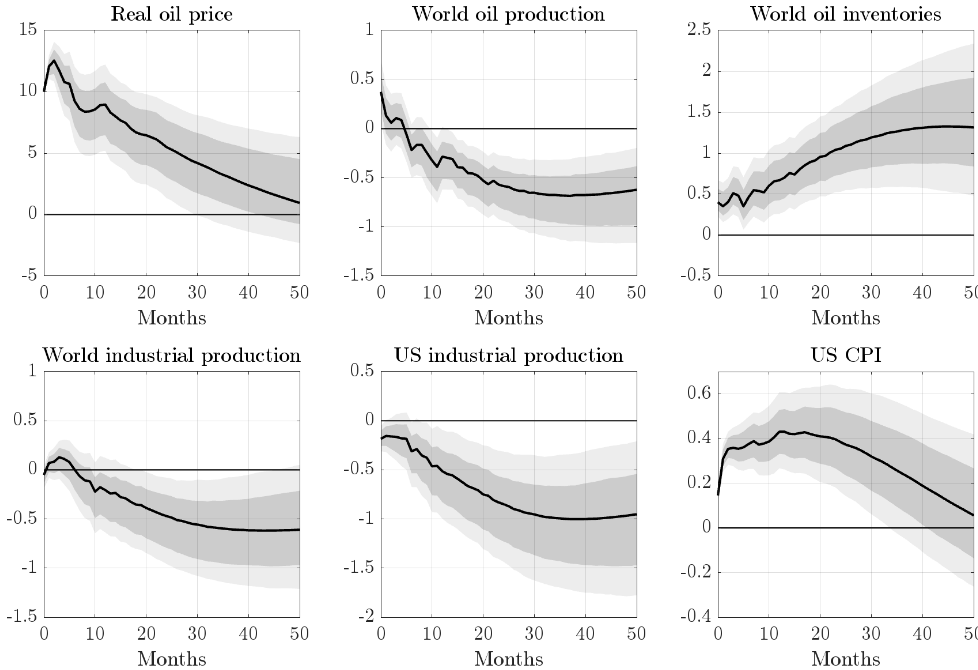

Figure 1: Linear impulse responses to an oil supply news shock using a linear Proxy-SVAR. Solid lines are the point estimates and the shaded areas are 68% and 90% confidence bands.

3.3 Testing asymmetries

In this section, we estimate the VARX and formally test for the assumption of linearity by testing the null hypothesis $H_0 : \beta(L) = 0$ in equation (2).[16]

Figure 2 presents the results of the test, with the responses to the term $|u_t^s|$, $\beta(L)$, plotted as dashed orange lines. The shaded areas are 68% and 90% bootstrapped confidence bands.[17] The nonlinear function of the shock has a significant and long-lasting effect on US variables while being small, temporary and hardly statistically different from zero for world variables. More specifically, industrial production and particularly prices respond negatively to the absolute value of the shock, suggesting

[16]The VARX in (5) includes 6 lags for both endogeneous and exogenous variables. To measure the instrument validity we follow Forni et al. (2023) and estimate the correlation coefficient between the estimated shock and the instrument, which is 0.23.

[17]The confidence bands take into account also the estimation uncertainty on the estimated structural shock from the first step projection in (8). Both steps of the estimation are included in the bootstrapping procedure.

that the overall recessionary effect of a positive shock on real activity (prices) is amplified (dampened) by the nonlinear term. Conversely, the overall expansionary effect of a negative shock on real activity (prices) is dampened (amplified) by the nonlinear term. Overall, the results indicate that the null hypothesis of linearity is rejected.

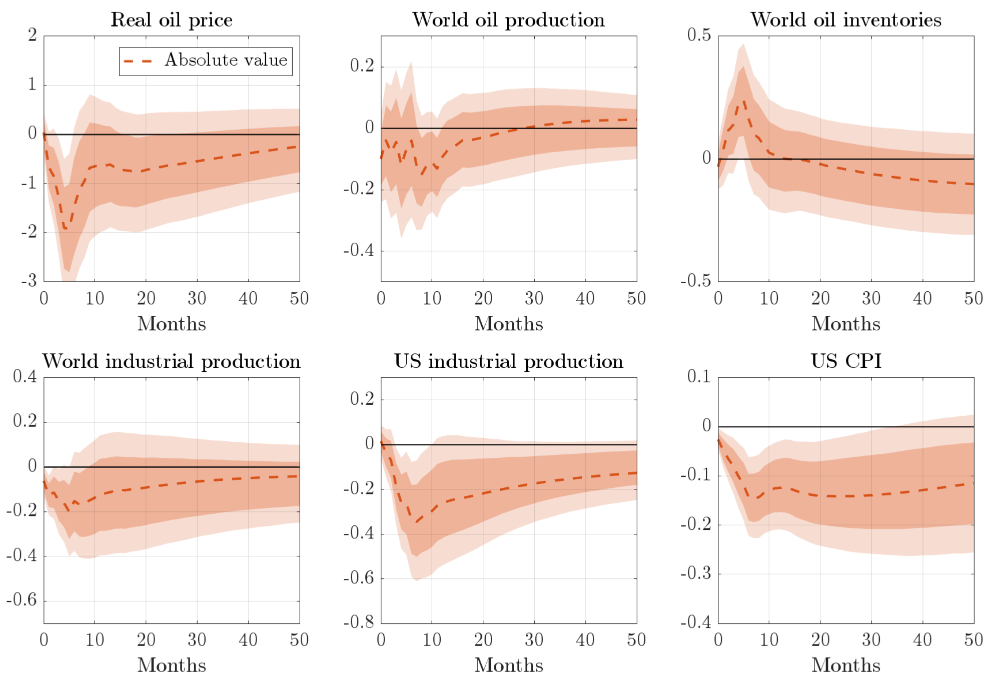

Figure 2: Nonlinear impulse responses to the absolute value of the shock estimated using a nonlinear Proxy-SVAR. The shaded areas represent the 68% and 90% confidence bands.

3.4 Asymmetric transmission of oil supply shocks

Figure 3 plots the IRFs of positive and negative oil supply shocks according to equation (3). The impulse responses are normalized to increase or decrease the real oil price by 10% on impact. The red and blue solid lines are the point estimates for the positive and negative shock, respectively, and the shaded areas are the 68% and 90% confidence bands. For ease of comparison, we multiply the impulse responses to a negative shock by minus one. Our results show no significant asymmetries in the dynamic response of global variables.

Turning our attention to the US economy, we see that a positive shock has larger effects on industrial production than a negative shock. While the nonlinear response of real activity to oil price shocks has been the focus of a heated debate in the literature, well described by Hamilton (2011) and Kilian and Vigfusson (2011b), our results point to an additional nonlinearity not previously explored: a positive shock leads to a modest effect on prices, while a negative shock leads to a much larger price response. Overall, our results suggest that following a positive oil supply shock, the US economy experiences a large and fast decline in real activity and a small increase in prices, while the opposite is true for a negative shock: a small decline in output and a large increase in prices.

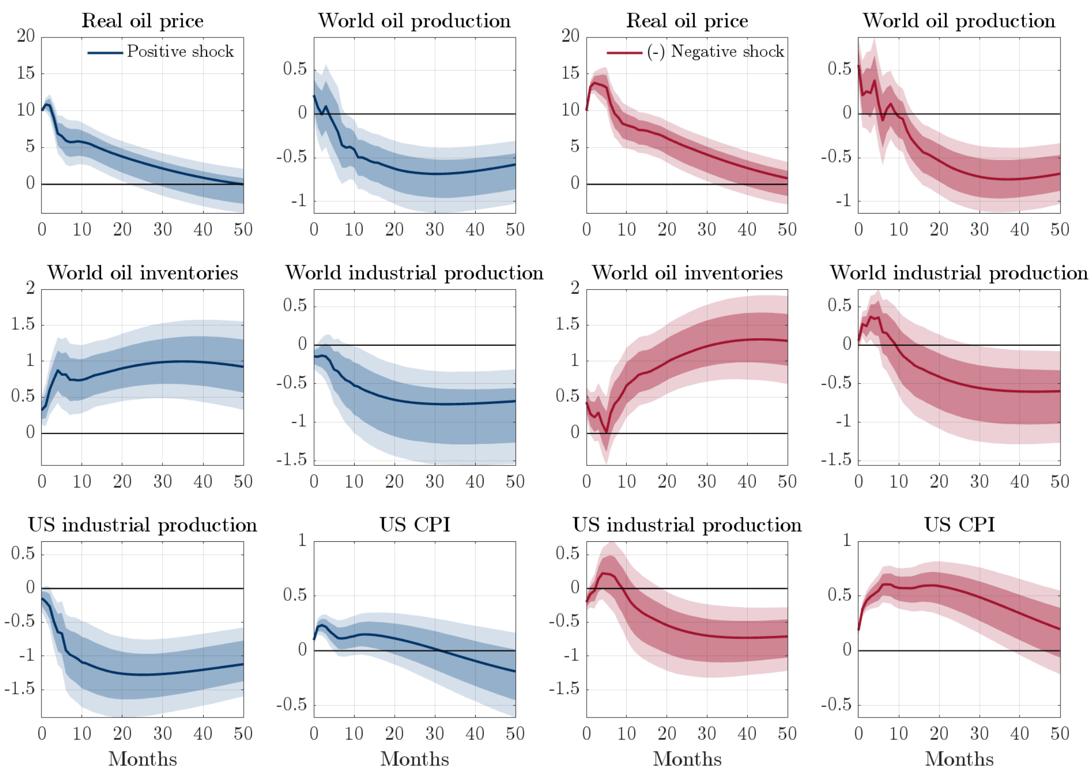

Figure 3: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock estimated using a nonlinear Proxy-SVAR. Impulse responses to a negative shock are multiplied by minus one. The functional form for the nonlinear function of the shock is $g(u_t^s) = |u_t^s|$. The shaded areas represent the 68% and 90% confidence bands.

3.5 Variance decomposition

Given the long-standing debate on whether oil supply shocks are major drivers of economic fluctuations (see Hamilton, 1983), we present estimates of the variance decomposition in our nonlinear framework.

The oil supply news shock u_t^s and its nonlinear function $g(u_t^s)$ are not necessarily orthogonal and, as a result, standard formulas to compute the variance decomposition are not appropriate in this setting. Therefore, following Forni et al. (2024), we address the problem by computing at each horizon the prediction error due to the linear and nonlinear oil supply shock, respectively u_t^s and $g(u_t^s)$, along with the total prediction error. The h-step ahead prediction error implied by equation (1) is

$$e_{t+h} = \sum_{k=0}^{h-1} \alpha_k u_t^s + \sum_{k=0}^{h-1} \beta_k g(u_t^s) + \sum_{k=0}^{h-1} \Gamma_k \xi_t$$

whereas the prediction error driven by the oil supply shock is composed by the part relative to the shock of interest, which is

$$e_{t+h}^s = \sum_{k=0}^{h-1} \alpha_k u_t^s + \sum_{k=0}^{h-1} \beta_k g(u_t^s)$$

We compute the prediction errors according to the above formulas; then, we compute the ratio of their sample variances. Table 2 presents the total contribution of the oil shock to the volatility of our variables in the right panel. In the left panel we also report the variance decomposition obtained by ignoring the nonlinear term (i.e., by assuming $\beta(L) = 0$). The oil supply news shock accounts for a large part of the variation of oil prices in the short run. Comparing our results with the literature, our estimates from the linear shock are in line with those in Känzig (2021). Once we add the contribution of the nonlinear function, the oil price shock explains a very large portion of the variance of US industrial production and prices, especially at longer horizons. Our conclusion is that oil shocks are major drivers of business cycle fluctuations.

	Linear				Total			
	$h = 0$	$h = 12$	$h = 24$	$h = 36$	$h = 0$	$h = 12$	$h = 24$	$h = 36$
Oil price	72.4	53.9	44.2	40.8	72.4	55.6	46.4	43.5
Oil production	3.0	3.7	18.9	31.1	3.2	5.3	20.7	32.7
Oil inventories	6.1	17.4	34.3	41.3	6.1	18.8	35.4	42.5
World IP	0.3	2.1	9.3	18.3	0.9	4.8	12.1	22.0
US IP	3.2	7.2	17.8	25.6	3.3	12.9	25.7	36.2
US CPI	18.8	36.0	24.0	17.4	19.5	44.9	36.4	34.6

Table 2: Variance Decomposition

3.6 Robustness checks

In this section, we assess the robustness of our results by: (i) implementing a different shock identification; (ii) adopting an alternative approach to test for invertibility; (iii) estimating impulse responses using local projections; (iv) running a horse-race against size nonlinearities; (v) changing various lag lengths in the model and (vi) ending the sample at the onset of the Global Financial Crisis.

Shock identification. In the first robustness check, we use the oil supply shock identified by Baumeister and Hamilton (2019) as a proxy in our nonlinear Proxy-SVAR.[18] Figure 4 plots the impulse responses, with the solid black lines and the shaded areas representing the point estimates and confidence bands for our baseline model, respectively. The dashed black lines are the point estimates using the alternative oil supply shock. In terms of instrument strength, the first-stage F-statistic in our model is 49.58, which is well above the recommended safe level of 10. The results indicate that there are no significant differences among the two proxies in a linear Proxy-SVAR framework. Figure 5, instead, plots the baseline impulse responses estimated using the nonlinear Proxy-SVAR and the alternative shock (dashed black lines) together with our baseline shock (solid lines). The results suggest that the asymmetries we document with our preferred instrument are robust to the use of the alternative identification strategy.

[18] We took the oil supply shock from Baumeister's webpage.

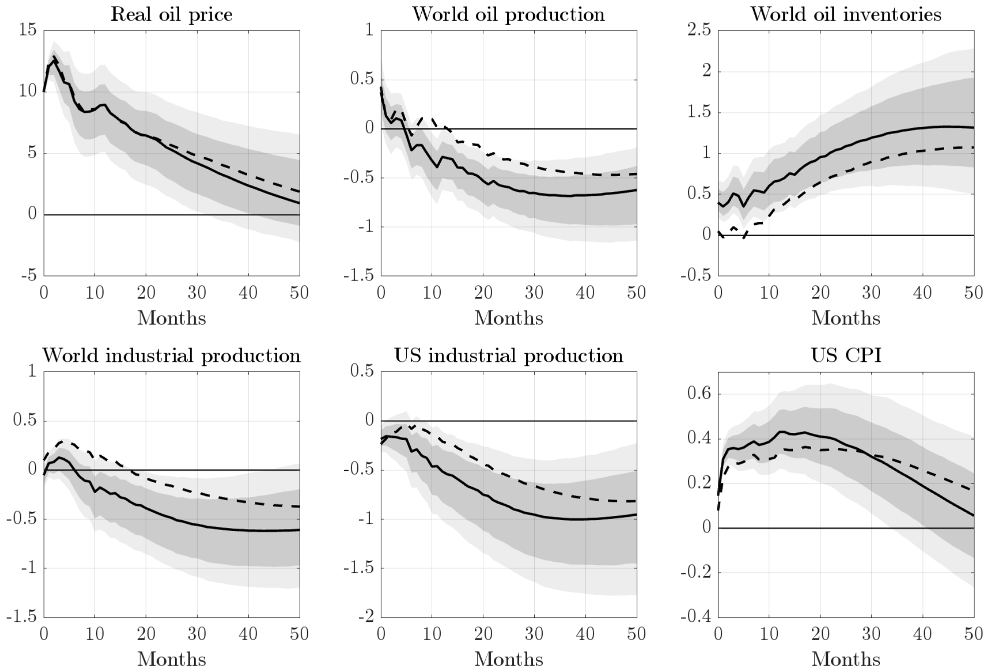

Figure 4: Linear impulse responses to our baseline oil supply news shock (black solid lines) and to the alternative identification strategy of Baumeister and Hamilton (2019) (black dashed lines) estimated using a linear Proxy-SVAR. The shaded areas are 68% and 90% confidence bands for our baseline impulse responses.

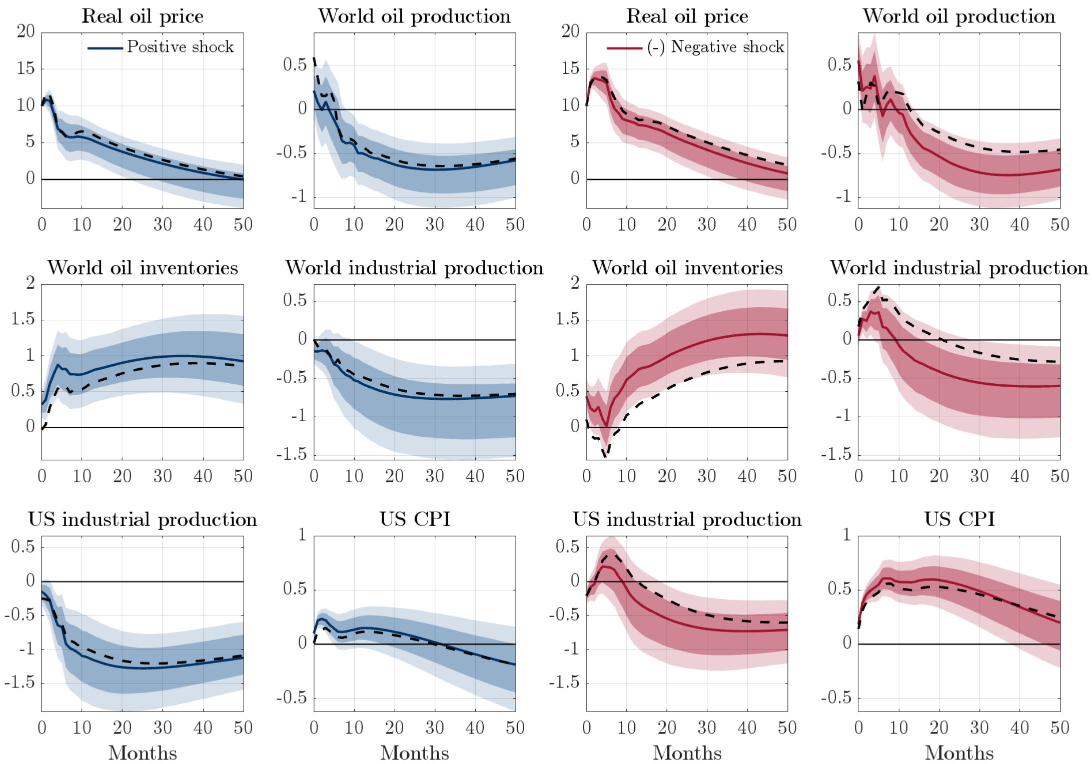

Figure 5: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock and the alternative identification strategy of Baumeister and Hamilton (2019) (black dashed lines) estimated using a nonlinear Proxy-SVAR. Impulse responses to a negative shock are multiplied by minus one. The functional form for the nonlinear function of the shock is $g(u_t^s) = |u_t^s|$. The shaded areas represent the 68% and 90% confidence bands for our baseline impulse responses.

Invertibility. In the second robustness check, we address potential concerns regarding the invertibility assumption, in view of the conflicting results with Plagborg-Møller and Wolf (2022, Online Appendix B.7).

We follow Forni and Gambetti (2014) and assess whether our baseline *shock*, estimated using the Känzig (2021) specification, can be predicted by one, three, six and twelve lags of the first five and the first eight principal components, obtained from the McCracken and Ng (2016) dataset. The results are reported in Table 3. In all cases we cannot reject the null hypothesis of orthogonality (lack of predictability). This result points to partial invertibility of the shock.

	First 5 PCs, k lags				First 8 PCs, k lags			
	$k=1$	$k=3$	$k=6$	$k=12$	$k=1$	$k=3$	$k=6$	$k=12$
F-stat	0.706	0.968	0.897	0.907	0.555	0.954	0.897	0.821
p-value	0.588	0.486	0.624	0.672	0.793	0.524	0.669	0.879

Table 3: Structuralness test. The table shows the F-statistics and the p-values for each regression including up to k lags of the first 5 and 8 principal components (PCs). The null hypothesis is orthogonality.

To further address the conflicting results with Plagborg-Møller and Wolf (2022), we perform our test described in Section 3.1 using the same set of variables included in Plagborg-Møller and Wolf (2022) and we confirm the rejection of invertibility. When we perform their Granger causality test using our set of variables, we cannot reject invertibility consistently with the results in Section 3.1. The finding suggests that the difference is attributable to the different sets of variables included in the test.

Local projections. We estimate the nonlinear impulse responses using local projections and the shock estimated from the first-step VAR. Figure 6 plots the results of the baseline VAR responses together with the local projection responses.[19] Reassuringly, the two models produce comparable results.

One might wonder why we use the shock estimated in the first stage, instead of the proxy itself. This is because in this nonlinear framework the use of a noisy measure of the shock, i.e., the proxy of Känzig (2021), in place of the true shock, would deliver biased estimates. A theoretical explanation and a simulation in support of this argument is available in Appendix A.

[19]In the local projection we control for p lags of all the variables in our baseline VAR model and compute heteroscedasticity-robust standard errors (see Montiel Olea and Plagborg-Møller, 2021).

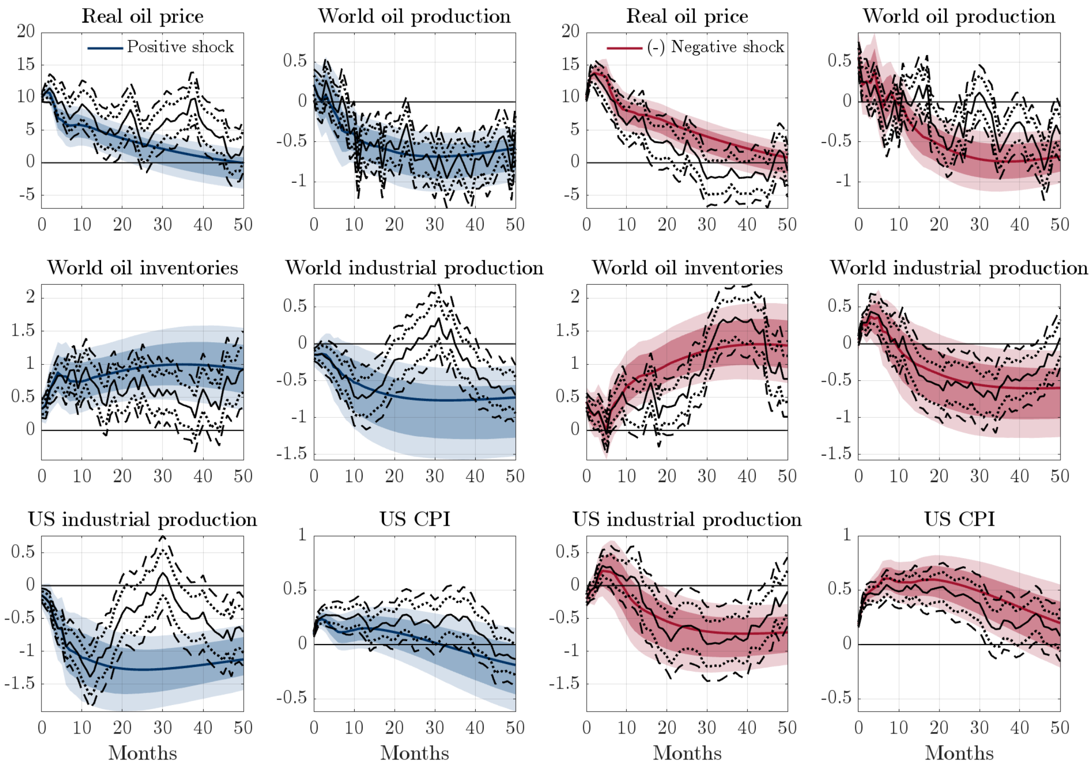

Figure 6: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock estimated using a nonlinear Proxy-SVAR and the nonlinear estimates from local projections (black dashed lines). Impulse responses to a negative shock are multiplied by minus one. The functional form for the nonlinear function of the shock is $g(u_t^s) = |u_t^s|$. The shaded areas represent the 68% and 90% confidence bands for our baseline impulse responses, while dotted and dashed black lines represent the 68% and 90% confidence bands for local projections.

Size nonlinearities. In this paper we use the absolute value as the functional form for $g(u_t^s)$. However, it is important to explore whether our results are influenced by the presence of size nonlinearities. We therefore extend our model to capture both sign and size nonlinearities.

First, we check whether controlling for size nonlinearities would alter our results. We proceed as follows. We extend our structural model in equation (1) as follows:

$$x_t \quad = \quad \nu + \Gamma(L)\xi_t + \alpha(L)u_t^s + \beta(L)|u_t^s| + \psi(L)(u_t^s)^3 \qquad (11)$$

To control for size nonlinearities we introduce a cubic term of the shock, as in Tenreyro and Thwaites (2016). Figure 7, Panel (b) presents a test of linearity for real oil price and US variables in the extended model (11), together with the linearity test in the baseline model (Panel (a)). In Panel (b), the first row reports the responses associated to the absolute value of the shock, $\beta(L)$, and the second row those associated to the cubic term, $\psi(L)$. The test indicates that (i) sign nonlinearities are still significant and (ii) size nonlinearities on US industrial production and CPI are hardly significant at the 90% confidence level.

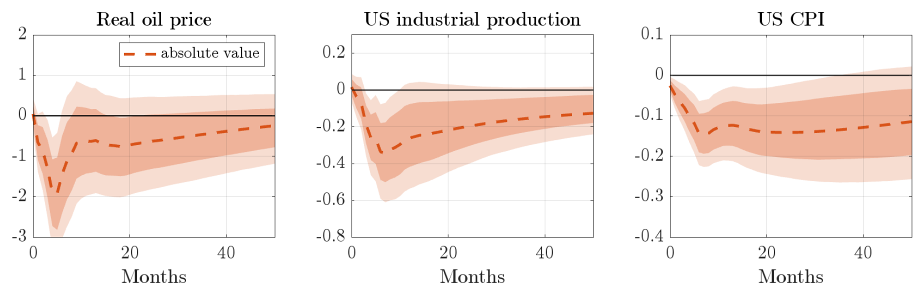

(a) *Baseline model*

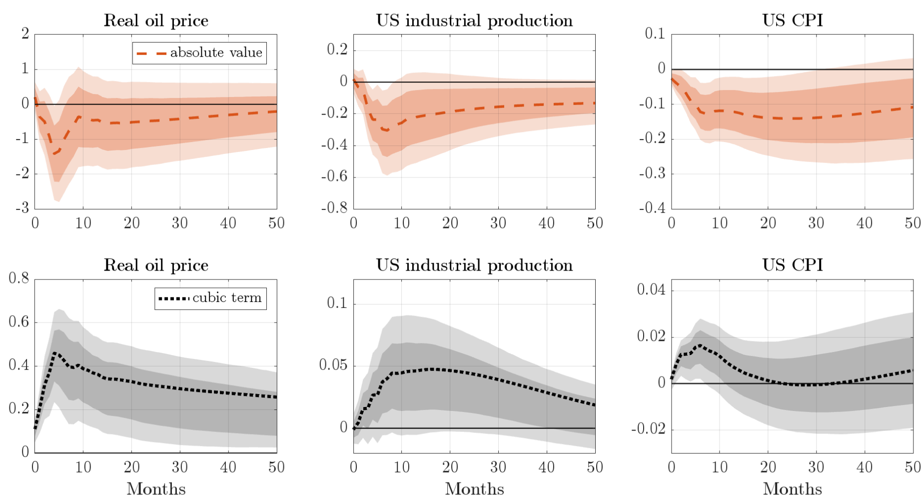

(b) *Extended model*

Figure 7: Panel (a): Nonlinear impulse responses for the coefficients associated with the absolute value (orange dashed line) in the baseline model (1). Panel (b): Nonlinear impulse responses for the coefficients associated with the absolute value (orange dashed line) and the cubic term (black dotted line) in model (11). The shaded areas represent the 68% and 90% confidence bands for both models.

Next, Figure 8 reports the sign-dependent impulse responses in the extended model (11) in black dashed lines together with our baseline point estimates and confidence bands. These responses indicate that our baseline results are unchanged when controlling for size nonlinearities. Overall, we have shown that (i) the asymmetric transmission of oil supply news shocks holds even when controlling for potential size nonlinearities and (ii) size nonlinearities are *qualitatively* present but not *quantitatively* important for output and prices in comparison with sign nonlinearities.

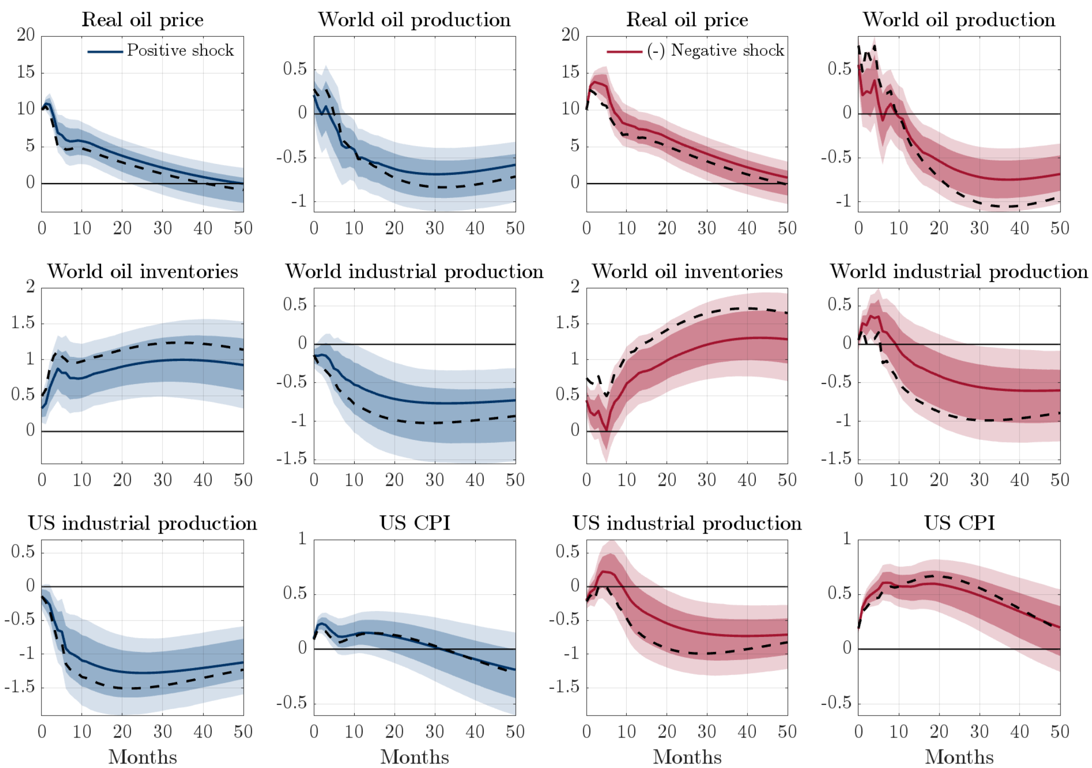

Figure 8: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock and the alternative specification which controls for size nonlinearities (black dashed lines) estimated using a nonlinear Proxy-SVAR. Impulse responses to a negative shock are multiplied by minus one. The functional forms for the nonlinear functions of the shock are $g_1(u_t^s) = |u_t^s|$ and $g_2(u_t^s) = (u_t^s)^3$. The shaded areas represent the 68% and 90% confidence bands for our baseline impulse responses.

Lag length. We now assess the sensitivity of our results to changes in the number of lags. Specifically, we estimate our nonlinear Proxy-SVAR with lags ranging from 3 to 9, which also includes our baseline lag selection ($p = 6$). Figure 9 plots the impulse responses from all these specifications and the 68% and 90% confidence bands for our baseline specification. As the figure shows, the nonlinearity we document is not affected by changes in the lag length of the model.

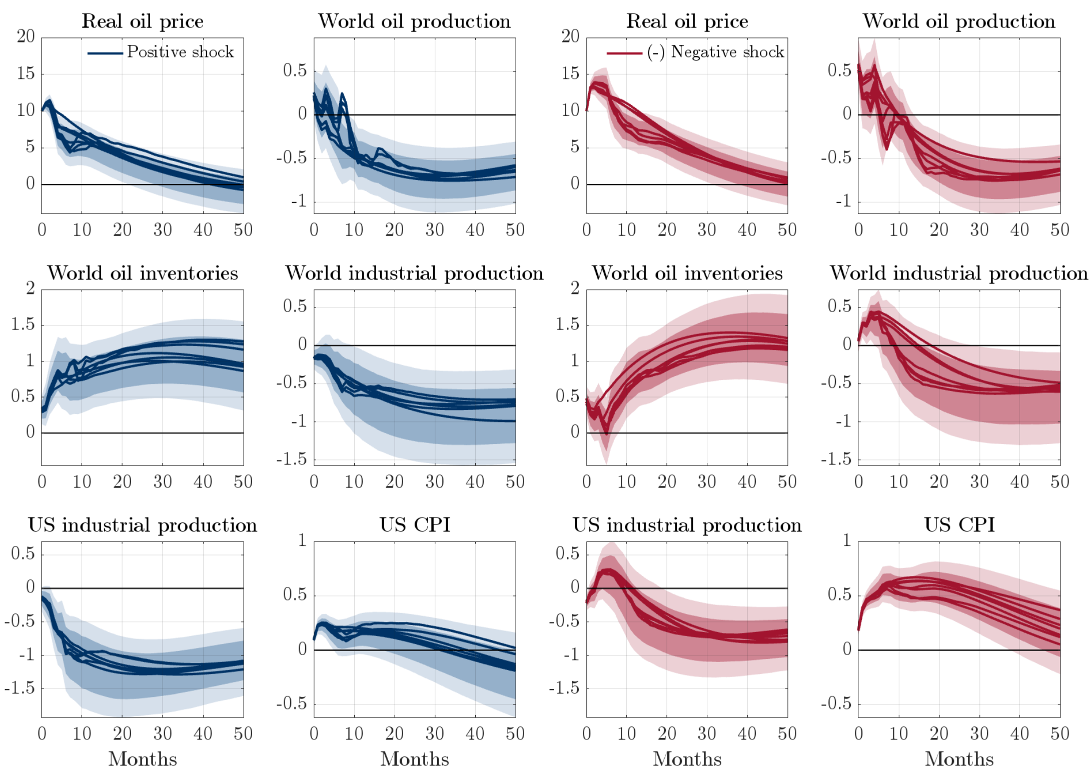

Figure 9: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock. Impulse responses to a negative shock are multiplied by minus one. The functional form for the nonlinear function of the shock is $g(u_t^s) = |u_t^s|$ estimated using a nonlinear Proxy-SVAR. Solid lines are the point estimates for models with 3 to 9 lags and the shaded areas represent the 68% and 90% confidence bands for our baseline specification (6 lags).

Sample specification and the Global Financial Crisis. Lastly, we assess the robustness of our results by ending the sample at the onset of the Global Financial Crisis (GFC) in August 2008. This exercise is also motivated by the results in Baumeister and Peersman (2013), who find greater effects of oil supply shocks on prices and real activity around the GFC.[20] Figure 10 plots the baseline impulse responses and confidence bands together with the alternative sample estimates (black dashed lines). The results indicate that our main findings are not sensitive to the exclusion of the GFC.[21]

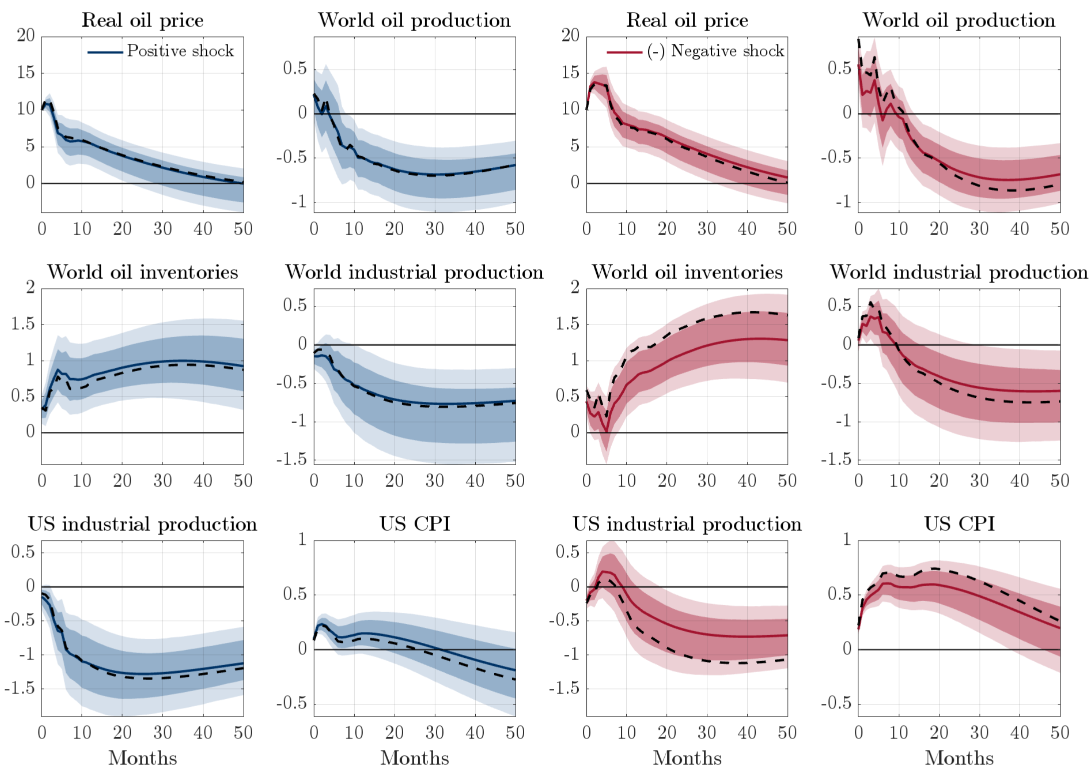

Figure 10: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock for our baseline model, and for the alternative model when the sample ends before the GFC, in August 2008 (black dashed lines), estimated using a nonlinear Proxy-SVAR. Impulse responses to a negative shock are multiplied by minus one. The functional form for the nonlinear function of the shock is $g(u_t^s) = |u_t^s|$. The shaded areas represent the 68% and 90% confidence bands.

[20] See Figure 1 (Panel A) in Baumeister and Peersman (2013).

[21] This sensitivity check also assures that our results are not driven by the large OPEC surprises such as the one in November 2014.

4 Potential indirect channels

The standard theoretical approach in modern dynamic stochastic general equilibrium models is to study surprise changes in the price of *imported* crude oil. Specifically, the literature has highlighted two channels, demand and supply, as the main direct effects of exogenous oil price shocks on real activity (see Kilian, 2014, for a review). The impact of a positive oil price shock through the demand channel focuses on the reduction in the disposable income of domestic consumers, as higher energy prices imply a transfer of income abroad.[22] The supply channel instead emphasizes that a positive oil price shock increases the cost of a factor of production (e.g., Rotemberg and Woodford, 1996; Finn, 2000). In this regard, Lee and Ni (2002) noted that while many industries are affected by oil price shocks through the demand channel, only oil-intensive industries are affected by the supply channel.

In addition to these direct effects, Kilian and Vigfusson (2011b) have highlighted the importance of other indirect channels, which might explain the asymmetric response of real activity to oil price shocks. These channels are related to the role of uncertainty and monetary policy.

In the next two sections, we estimate the response of key variables to explore the relevance of these indirect channels. To do so, we estimate a nine-variables model by adding the following three variables to our baseline specification: the financial uncertainty index used by Bloom (2009), the excess bond premium (EBP) from Gilchrist and Zakrajšek (2012) and the federal funds rate.[23] Figure 11 reports the results.

4.1 Uncertainty channel

A first explanation hints at the role that uncertainty about the future price of oil may play in current investment decisions. Oil supply shocks, both positive and negative, can in principle increase oil price volatility and therefore financial uncertainty. The larger is the size of the shock (in absolute value), the larger is the effect on uncertainty. Uncertainty in turn may reduce private investment both because of

[22] A recent empirical study by Hamilton (2009) explores the causes and consequences of this channel by analyzing the 2007-08 oil shock on consumer spending.

[23] The financial uncertainty index is retrieved from FRED (id: VXOCLS), extended as in Bloom (2009), and the federal funds rate from FRED (id: DFF).

a "real option" effect and a "risk premium" effect. According to the "real option" effect, when individual projects are irreversible, increased uncertainty leads to a reduction in investment expenditures, since higher uncertainty increases the returns to waiting for information, causing firms to delay investment (e.g., Bernanke, 1983; Bloom et al., 2007). According to the "risk premium" effect, higher uncertainty increases the probability of bad outcomes for the firm, raising the risk of investment and therefore the cost of finance (Christiano et al., 2014; Gilchrist et al., 2014).

Given that unexpected changes in the real price of oil are potentially associated with higher expected volatility, this channel may be responsible for amplifying the negative real effects of unexpected oil price increases and dampening the positive real effects of price decreases (see Kilian and Vigfusson, 2011b). The reduction of private investment can in principle affect prices, amplifying the depressing effect of an oil price reduction and reducing the inflationary effect of oil price increases. Along these lines, Baumeister and Kilian (2016) suggest that oil price shocks can reduce future investment projects through their effect on expectations about the future path of oil price.

To shed light on the role of uncertainty, we estimate the impulse response functions of an uncertainty index (the VXO) and of the EBP, an indicator of credit market conditions. Figure 11 shows that both negative and positive shocks raise financial uncertainty and the EBP. These responses are consistent with the uncertainty channel described above (see also Elder and Serletis, 2010). In Appendix C we show that the uncertainty argument works through oil price uncertainty. On the other hand, the responses of the macroeconomic uncertainty indicator of Jurado et al. (2015) are symmetric.

4.2 Monetary policy channel

Bernanke et al. (1997) suggest that the recessionary effects of an oil price shock are largely a result of the response of the Federal Reserve to contain inflationary pressures. The channel they have in mind works as follows: the higher interest rate set by the central bank in an attempt to contain inflationary pressures may exacerbate the negative impact of a positive oil price shock on real economic activity. Later on, this mechanism has been challenged in a number of ways. On one hand,

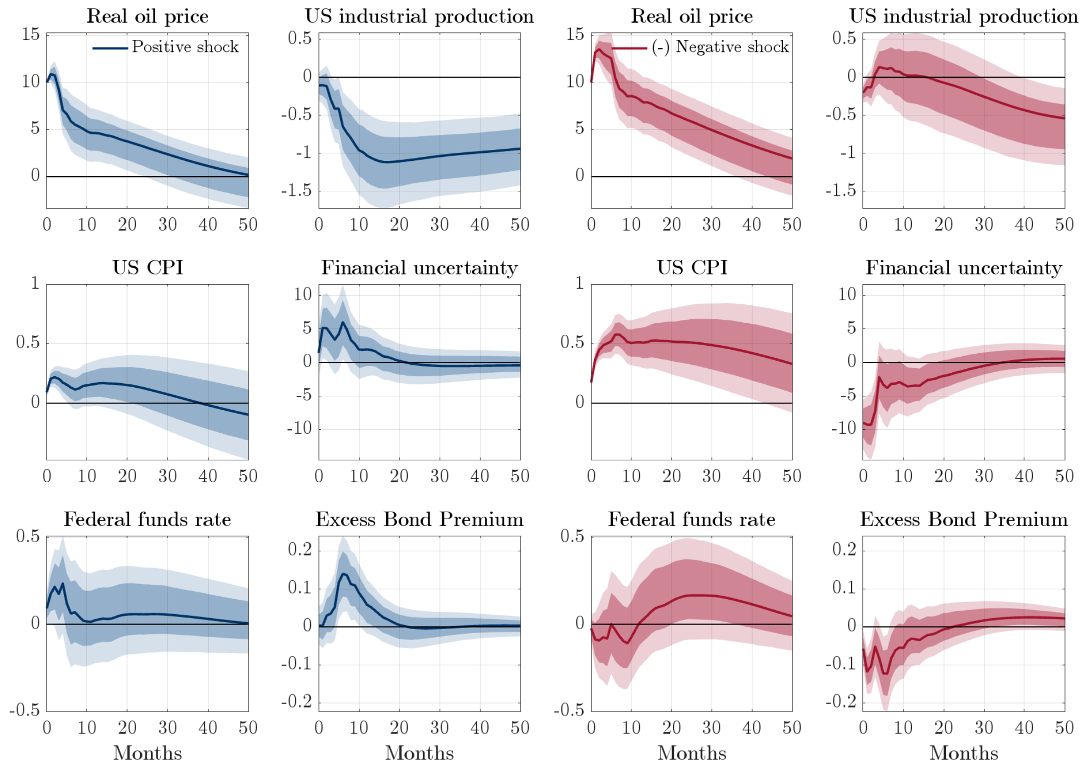

Figure 11: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock. The baseline specification is augmented with three additional variables, giving rise to a nine-variable model. Impulse responses to a negative shock are multiplied by minus one. World variables are excluded from the plot. The functional form for the nonlinear function of the shock is $g(u_t^s) = |u_t^s|$. The shaded areas represent the 68% and 90% confidence bands.

Hamilton and Herrera (2004) show that these results on the response of the interest rate are sensitive to model specification and cast doubts on the feasibility of policies to offset the recessionary effects of an oil price shock. On the other hand, Herrera and Pesavento (2009) show that the systematic response of monetary policy caused fluctuations in economic activity only in the pre-Volcker period, while Kilian and Lewis (2011) find that this was not the case even in the pre-Volcker period.

The estimated impulse responses in Figure 11 show that the Fed Funds Rate increases after both shocks. However, these responses are barely statistically significant after a positive shock and not statistically different from zero after a negative shock. This evidence suggests that the systematic response of the central bank can only explain a residual component of the asymmetries documented in this paper.

4.3 A counterfactual exercise

In this subsection, we report a counterfactual scenario enforced by switching off the uncertainty channel following the method proposed by McKay and Wolf (2023).

The counterfactual impulse responses are estimated by using three types of financial shocks. The first is identified as in Forni et al. (2024), the second with the high-frequency instrument of Piffer and Podstawski (2018), and the third following

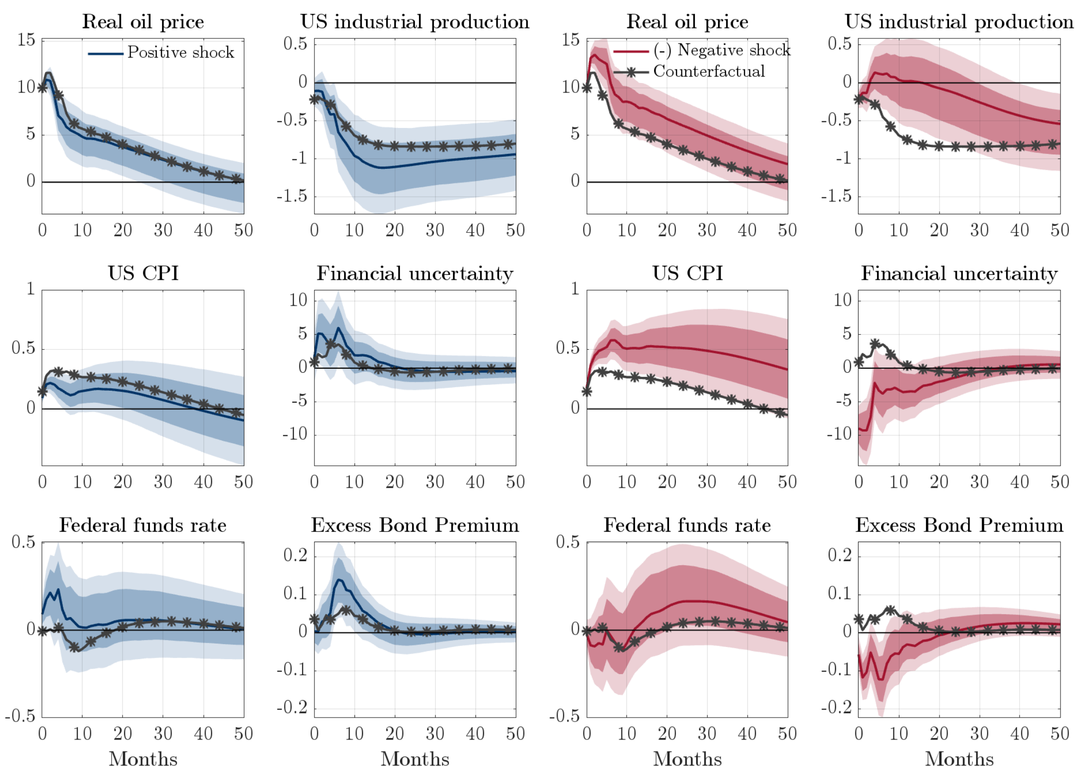

Figure 12: Nonlinear impulse responses to a positive (blue line), a negative (red line) oil supply news shock, and counterfactual responses for a negative shock (grey line with asterisk markers) when shutting down the uncertainty channel. Impulse responses to a negative shock are multiplied by minus one. World variables are excluded from the plot. The shaded areas represent the 68% and 90% confidence bands.

the optimization-based approach of Caldara et al. (2016). Such shocks are used within a VARX as exogenous variables and are linearly combined in such a way to minimize the asymmetric responses of uncertainty. We enforce this counterfactual by focusing on the baseline responses to a negative shock (red lines). Thus, the counterfactual responses to a negative shock are obtained by combining the responses under the baseline scenario and the weighted effects of the three uncertainty shocks. Figure 12 plots the results. The counterfactual responses are the grey lines with asterisk markers.

The figure shows that the counterfactual responses are very close to those following a positive shock (blue lines) and very different from those following a negative shock (red lines), so that the asymmetric responses of industrial production and prices are almost eliminated. Overall, we conclude that the uncertainty channel is sufficient to explain the asymmetric transmission of oil supply news shocks to output and prices.

5 Conclusion

We document important asymmetries in the transmission of oil supply news shocks using the novel nonlinear Proxy-SVAR developed in Debortoli et al. (2023). We find that a positive shock has large real effects and a small impact on prices. In contrast, a negative shock has small real effects and a large impact on prices. While we do find evidence in favor of sign-dependent effects, we do not find any statistically significant size-dependent effects in the transmission of oil supply news shocks.

We rationalize the asymmetric transmission of oil supply news shocks in view of two channels. The first is related to uncertainty, which includes a "real option" and a "risk premium" effect. According to this channel, an oil supply shock, regardless of its sign, increases uncertainty, which contributes to depress economic activity. This channel therefore operates by amplifying the negative real effects of unexpected oil price increases and dampening the positive real effects of unexpected oil price decreases. The opposite holds for prices. The second channel is related to the systematic response of monetary policy, and we find little role for this mechanism. The quantitative relevance of the first channel is confirmed in a counterfactual exercise, where second-round effects from uncertainty explain much of the difference in the estimates between a positive and a negative shock.

References

Abiad, Abdul and Irfan A Qureshi, "The macroeconomic effects of oil price uncertainty," *Energy Economics*, 2023, *125*, 106839.

Barnichon, Regis and Christian Matthes, "Functional approximation of impulse responses," *Journal of Monetary Economics*, 2018, *99*, 41–55.

_ **and** _ , "Functional Approximation of Impulse Responses," *Journal of Monetary Economics*, 2018, *99*, 41–55.

Baumeister, Christiane and Gert Peersman, "Time-varying effects of oil supply shocks on the US economy," *American Economic Journal: Macroeconomics*, 2013, *5* (4), 1–28.

_ **and James D Hamilton**, "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks," *American Economic Review*, 2019, *109* (5), 1873–1910.

_ **and Lutz Kilian**, "Forty years of oil price fluctuations: Why the price of oil may still surprise us," *Journal of Economic Perspectives*, 2016, *30* (1), 139–160.

Bernanke, Ben, Mark Gertler, Mark Watson, and Christopher Sims, "Systematic monetary policy and the effects of oil price shocks," *Brookings papers on economic activity*, 1997, (1), 91–157.

Bernanke, Ben S, "Irreversibility, uncertainty, and cyclical investment," *The quarterly journal of economics*, 1983, *98* (1), 85–106.

Bloom, Nicholas, "The impact of uncertainty shocks," *Econometrica*, 2009, *77* (3), 623–685.

Bloom, Nick, Stephen Bond, and John Van Reenen, "Uncertainty and investment dynamics," *The Review of Economic Studies*, 2007, *74* (2), 391–415.

Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek, "The macroeconomic impact of financial and uncertainty shocks," *European Economic Review*, 2016, *88*, 185–207.

Caravello, Tomás and Pedro Martinez-Bruera, "Disentangling Sign and Size Non-linearities," *Available at SSRN 4704050*, 2024.

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno, "Risk Shocks," *American Economic Review*, January 2014, *104* (1), 27–65.

Debortoli, Davide, Mario Forni, Luca Gambetti, and Luca Sala, "Asymmetric Monetary Policy Tradeoffs," *CEPR Press DP18438*, 2023.

Elder, John and Apostolos Serletis, "Oil price uncertainty," *Journal of Money, Credit and Banking*, 2010, *42* (6), 1137–1159.

Finn, Mary G, "Perfect competition and the effects of energy price increases on economic activity," *Journal of Money, Credit and banking*, 2000, pp. 400–416.

Forni, Mario and Luca Gambetti, "Sufficient information in structural VARs," *Journal of Monetary Economics*, 2014, *66*, 124–136.

_ , _ , and Giovanni Ricco, "External Instrument SVAR Analysis forNoninvertible Shocks," Working Papers 2023-03, Center for Research in Economics and Statistics January 2023.

_ , _ , and Luca Sala, "Structural VARs and noninvertible macroeconomic models," *Journal of Applied Econometrics*, 2019, *34* (2), 221–246.

_ , _ , Nicolò Maffei-Faccioli, and Luca Sala, "Nonlinear transmission of financial shocks: Some new evidence," *Journal of Money, Credit and Banking*, 2024, *56* (1), 5–33.

Gilchrist, Simon and Egon Zakrajšek, "Credit spreads and business cycle fluctuations," *American economic review*, 2012, *102* (4), 1692–1720.

_ , Jae W. Sim, and Egon Zakrajšek, "Uncertainty, Financial Frictions, and Investment Dynamics," NBER Working Papers 20038, National Bureau of Economic Research, Inc April 2014.

Hamilton, James D, "Oil and the macroeconomy since World War II," *Journal of political economy*, 1983, *91* (2), 228–248.

_ , "This is what happened to the oil price-macroeconomy relationship," *Journal of monetary economics*, 1996, *38* (2), 215–220.

_ , "What is an oil shock?," *Journal of econometrics*, 2003, *113* (2), 363–398.

_ , "Causes and consequences of the oil shock of 2007-08," *Brookings Papers on Economic Activity*, 2009, pp. 215–284.

_ , "Nonlinearities and the macroeconomic effects of oil prices," *Macroeconomic dynamics*, 2011, *15* (S3), 364–378.

_ and Ana Maria Herrera, "Comment: oil shocks and aggregate macroeconomic behavior: the role of monetary policy," *Journal of Money, Credit and Banking*, 2004, pp. 265–286.

Herrera, Ana María and Elena Pesavento, "Oil price shocks, systematic monetary policy, and the "Great Moderation"," *Macroeconomic Dynamics*, 2009, *13* (1), 107–137.

_ , Latika Gupta Lagalo, and Tatsuma Wada, "Oil price shocks and industrial production: Is the relationship linear?," *Macroeconomic dynamics*, 2011, *15* (S3), 472–497.

_ , _ , and _ , "Asymmetries in the response of economic activity to oil price increases and decreases?," *Journal of International Money and Finance*, 2015, *50*, 108–133.

Jurado, Kyle, Sydney C Ludvigson, and Serena Ng, "Measuring uncertainty," *American Economic Review*, 2015, *105* (3), 1177–1216.

Känzig, Diego, "The macroeconomic effects of oil supply news: Evidence from OPEC announcements," *American Economic Review*, 2021, *111* (4), 1092–1125.

Kilian, Lutz, "Oil price shocks: Causes and consequences," *Annu. Rev. Resour. Econ.*, 2014, *6* (1), 133–154.

_ and Logan T Lewis, "Does the Fed respond to oil price shocks?," *The Economic Journal*, 2011, *121* (555), 1047–1072.

_ and Robert J Vigfusson, "Are the responses of the US economy asymmetric in energy price increases and decreases?," *Quantitative Economics*, 2011, *2* (3), 419–453.

_ and _ , "Nonlinearities in the oil price–output relationship," *Macroeconomic Dynamics*, 2011, *15* (S3), 337–363.

_ and _ , "The role of oil price shocks in causing US recessions," *Journal of Money, Credit and Banking*, 2017, *49* (8), 1747–1776.

Lee, Kiseok and Shawn Ni, "On the dynamic effects of oil price shocks: a study using industry level data," *Journal of Monetary economics*, 2002, *49* (4), 823–852.

McCracken, Michael W and Serena Ng, "FRED-MD: A monthly database for macroeconomic research," *Journal of Business & Economic Statistics*, 2016, *34* (4), 574–589.

McKay, Alisdair and Christian K Wolf, "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?," *Econometrica*, 2023, *91* (5), 1695–1725.

Mertens, Karel and Morten O. Ravn, "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, June 2013, *103* (4), 1212–47.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The transmission of monetary policy shocks," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 74–107.

Mork, Knut Anton, "Oil and the macroeconomy when prices go up and down: an extension of Hamilton's results," *Journal of political Economy*, 1989, *97* (3), 740–744.

Olea, José Luis Montiel and Mikkel Plagborg-Møller, "Local projection inference is simpler and more robust than you think," *Econometrica*, 2021, *89* (4), 1789–1823.

Pellegrino, Giovanni, Efrem Castelnuovo, and Giovanni Caggiano, "Uncertainty and monetary policy during the great recession," *International Economic Review*, 2023, *64* (2), 577–606.

Piffer, Michele and Maximilian Podstawski, "Identifying uncertainty shocks using the price of gold," *The Economic Journal*, 2018, *128* (616), 3266–3284.

Plagborg-Møller, Mikkel and Christian K Wolf, "Instrumental variable identification of dynamic variance decompositions," *Journal of Political Economy*, 2022, *130* (8), 2164–2202.

Rotemberg, Julio and Michael Woodford, "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit and Banking*, 1996, *28* (4), 550–77.

Stock, James H. and Mark W. Watson, "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, 2018, *128* (610), 917–948.

Stock, James H and Mark W Watson, "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 2018, *128* (610), 917–948.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a string: US monetary policy is less powerful in recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Appendix A

In this Appendix we discuss the potential problems which arise using LP in a non-linear setting when the shock is not perfectly observed but rather is contaminated by noise.

Consider the following simple model:

$$\begin{aligned} x_t &= \alpha_x u_t + \beta_x g(u_t) + \eta_t^x \\ y_t &= \alpha_y u_t + \beta_y g(u_t) + \eta_t^y \end{aligned}$$

where $u_t \sim iid(0,1)$. The impulse response functions to a shock $u_t = u^*$ are $\alpha_x u^* + \beta_x g(u^*)$ and $\alpha_y u^* + \beta_y g(u^*)$. The relative responses are $\frac{\alpha_x u^* + \beta_x g(u^*)}{\alpha_y u^* + \beta_y g(u^*)}$. The shock is not observed but a noisy measure of the shock is available to the econometrican: $z_t = u_t + v_t$, where $v_t \sim iid(0,1)$ is independent from u_t at all lead and lags. Also assume for simplicity that the distribution of u_t and $g(u_t)$ are such that $Cov(z_t, g(z_t)) = Cov(u_t, g(u_t)) = Cov(u_t, g(z_t)) = Cov(z_t, g(u_t)) = 0$. Consider estimating the two regressions using the instrument instead of the shock:

$$\begin{aligned} x_t &= \tilde{\alpha}_x z_t + \tilde{\beta}_x g(z_t) + e_t^x \\ y_t &= \tilde{\alpha}_y z_t + \tilde{\beta}_y g(z_t) + e_t^y. \end{aligned}$$

The OLS population parameters are

$$\begin{aligned} \tilde{\alpha}_x = \alpha_x \delta, \quad \tilde{\beta}_x = \beta_x \gamma \\ \tilde{\alpha}_y = \alpha_y \delta, \quad \tilde{\beta}_y = \beta_y \gamma \end{aligned} \tag{A1}$$

where $\delta = \frac{Cov(z_t, u_t)}{Var(z_t)}$ and $\gamma = \frac{Cov(g(z_t), g(u_t))}{Var(g(z_t))}$ represent the effects of the distortions arising from the presence of measurement error. Of course these parameters will fail to deliver the correct impulse response functions. This is true also in linear LP settings and justifies the use of the LP-IV approach which delivers the impulse response functions to a shock with a unit effect on a given variable, i.e. the rescaled impulse response functions. Is there a way to correct the distorted parameters in (A1) in order to get a reliable estimate of the relative responses? The answer is no, unless the distortions are equal: $\delta = \gamma$. To see this, let us define the bias as the difference between the distorted and true responses:

$$b = \frac{\delta\alpha_x u^* + \gamma\beta_x g(u^*)}{\delta\alpha_y u^* + \gamma\beta_y g(u^*)} - \frac{\alpha_x u^* + \beta_x g(u^*)}{\alpha_y u^* + \beta_y g(u^*)}.$$

Rearranging terms we obtain

$$b = \frac{(\alpha_x\beta_y - \alpha_y\beta_x)u^* g(u^*)}{(\delta\alpha_y u^* + \gamma\beta_y g(u^*))(\alpha_y u^* + \beta_y g(u^*))}(\delta - \gamma)$$

Except in the special case where $\alpha_x\beta_y = \alpha_y\beta_x$, the bias in the relative responses will be zero if and only if $\delta = \gamma$, a condition which of course does not hold in general. This example shows that in a nonlinear context, not even the relative responses in general can be correctly estimated using a noisy measure of the shock. In the case of sign asymmetries, sometimes researchers use positive and negative values of the shock as regressors. Again, if the shock is not available, using positive and negative values of the instrument yields correct results only if the sign of the instrument is equal to the sign of the shock, so that the conditioning set is correct. Otherwise, splitting the instrument in positive and negative values would deliver biased results.

It should be noted that, in this simple example, the parameters of the linear and nonlinear term can be corrected separately as in the LP-IV approach. The intuition is that z_t is a valid instrument for u_t and $g(z_t)$ a valid instrument for $g(u_t)$. Since the bias is constant across equations, the rescaled parameters (those obtained in the LP-IV approach) are the correct ones:

$$\frac{Cov(x_t z_t)}{Cov(y_t z_t)} = \frac{\tilde{\alpha}_x}{\tilde{\alpha}_y} = \frac{\alpha_x}{\alpha_y}, \quad \frac{Cov(x_t g(z_t))}{Cov(y_t g(z_t))} = \frac{\tilde{\beta}_x}{\tilde{\beta}_y} = \frac{\beta_x}{\beta_y}.$$

But of course the combination $\frac{\alpha_x}{\alpha_y}c + \frac{\beta_x}{\beta_y}g(c)$ (where c is the value of the shock) will not deliver, in general, the correct responses. To see this notice that $\frac{\alpha_x}{\alpha_y}c$ corresponds to the response of x_t to a shock of size $u^* = \alpha_y^{-1}c$. This means that the consistent effect of the nonlinear term is $\beta_x g(\alpha_y^{-1}c)$. So the term $\frac{\beta_x}{\beta_y}g(c)$ would correspond to the correct nonlinear effect if and only if $\beta_y^{-1}g(c) = g(\alpha_y^{-1}c)$, which of course will not be true in general. For instance, in the case $g(u_t) = u_t^2$ the correct estimation of the total effects requires $\beta_y = \alpha_y^2$ (from $c^2\beta_y^{-1} = \alpha_y^{-2}c^2$). This shows that the rescaled parameters, despite being non-distorted, are not useful to obtain the IRF simply because the correct weights associated to the two parameters is unknown.

We run a simulation to have a better assessment of the above implications. We generate data from the model

$$y_t = Ay_{t-1} + B\nu_t$$

where

$$A = \begin{pmatrix} 0.9 & 0.4 \\ -0.2 & 0.7 \end{pmatrix} \quad B = \begin{pmatrix} 1 & 0.2 & 0.4 \\ -0.2 & 0.5 & 0.8 \end{pmatrix}$$

and $\nu_t = [u_{1t}\ u_{2t}\ |u_{1t}|]'$ where $[u_{1t}\ u_{2t}]' \sim N(0, I)$ and u_{1t} is the shock of interest. We also generate an instrument $z_t = u_{1t} + v_t$ where $v_t \sim N(0, 1)$. We generate 500 datasets of 500 observations each (large samples) and apply both our procedure (see Section 2) and the nonlinear LP using the noisy instrument.

Results are displayed in Figure A.1. Panel A shows results with our approach, Panel B shows results obtained using local projections with a noisy measure of the shock (z_t). The solid lines are the true IRF, the gray areas are the 5th and 95th percentile of the distribution of the point estimates. When using LP, distortions are evident: solid lines are outside the bands for several horizons, especially for negative shocks. Of course, by setting the variance of the noise to zero, LP works perfectly, suggesting that the problem is not with the model used to obtain the impulse responses, but rather with the shock measure used.

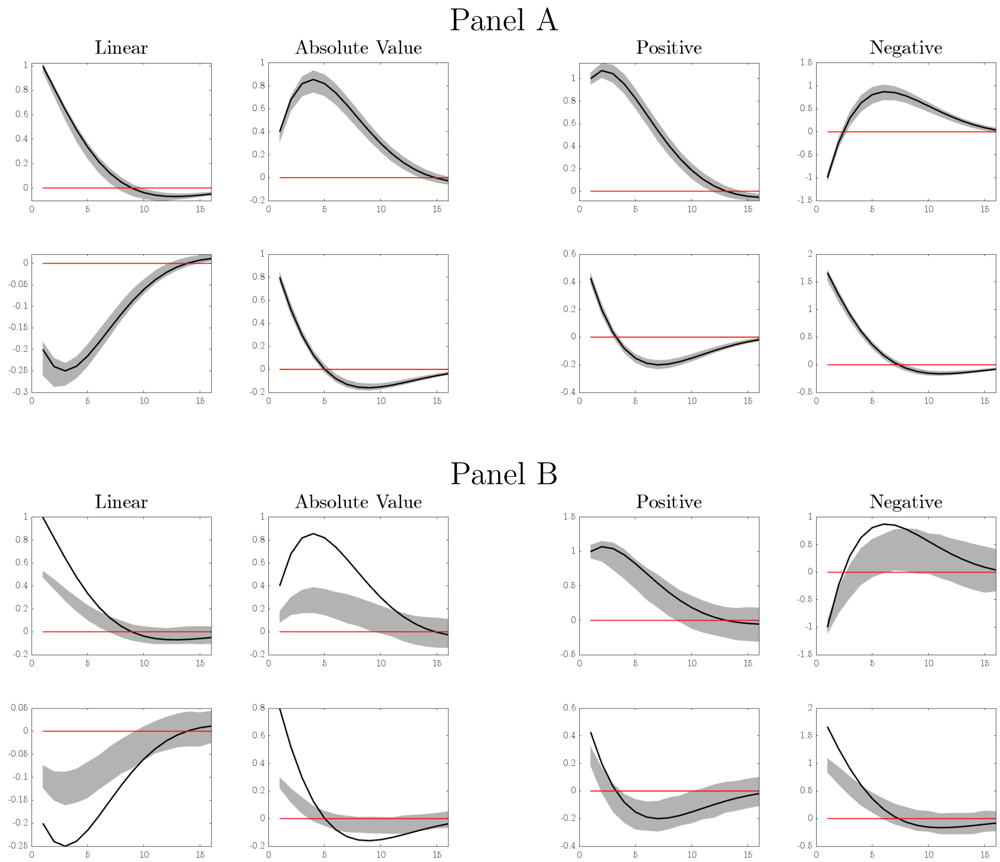

Figure A.1

Appendix B

To investigate discrepancies with the results of Caravello and Martinez-Bruera (2024) we have used their replication codes with our dataset. We use the proxy of Känzig (2021) and its absolute value as the exogenous variables. We include as controls 18 lags for real oil price, oil production, oil stocks, world industrial production, US industrial production and US CPI. Results are reported in Figure B.1. The first row reports the linear coefficients, the second row those associated with the absolute value of the *shock*, and the third row the combined nonlinear impulse responses. The black lines represents the impulse responses to a positive shock and the red lines those to a negative shock (multiplied by minus 1). Since red and black lines are very close to each other, we do not have evidence of asymmetry, in line with

Caravello and Martinez-Bruera (2024). Hence the different results are due to the different method used; more specifically, to the use of the proxy and its absolute value in place of the estimated shock and its absolute value.

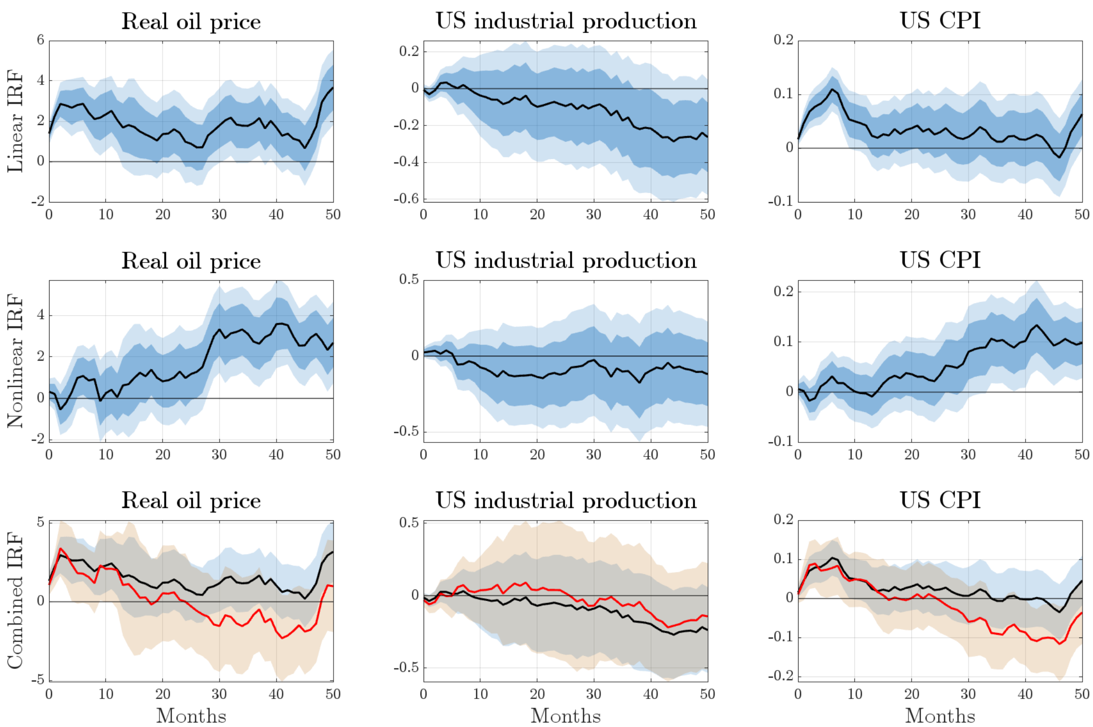

Figure B.1: Linear IRF (first row), Nonlinear IRF (second row), and Combined IRF (third row) estimated using local projections and the proxy measure of Känzig (2021). The black lines represents those to a positive shock and the red lines those to a negative shock (multiplied by minus 1). Shaded areas in the first two plots are 68% and 90% confidence bands. Shaded areas in the third plot are 90% confidence bands.

Appendix C

In this Appendix, we estimate the effects of the oil supply news shock on two additional uncertainty measures.

The former is the measure of oil price uncertainty constructed using textual analysis techniques by Abiad and Qureshi (2023). The result is reported in Figure C.1. The figure shows that both positive and negative oil supply shocks lead to a significant increase in oil price uncertainty (the effect of negative shocks are taken as usual with the minus sign). This result provides an intuition for the uncertainty channel documented in the paper.[24]

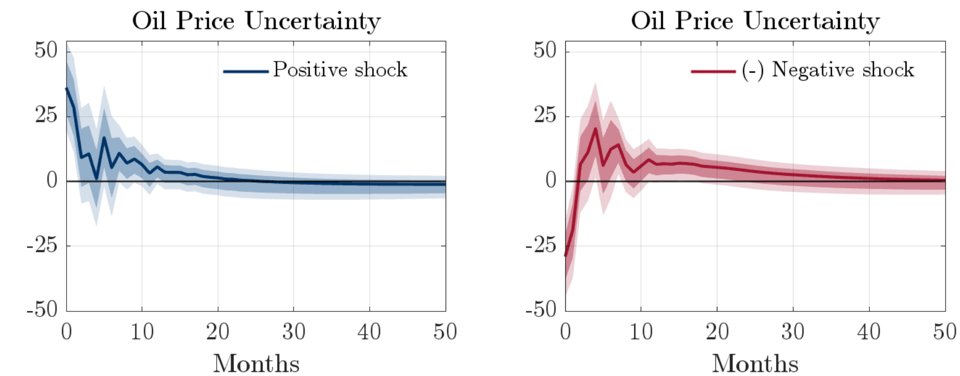

Figure C.1: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock estimated using a nonlinear Proxy-SVAR. The core six-variable model is extended including the oil price uncertainty measure of Abiad and Qureshi (2023). Impulse responses to a negative shock are multiplied by minus one. The shaded areas represent the 68% and 90% confidence bands.

The latter measure is the macroeconomic uncertainty measure of Jurado et al.

[24] The results hold for a different measure of oil price uncertainty calculated using a GARCH model as in Elder and Serletis (2010).

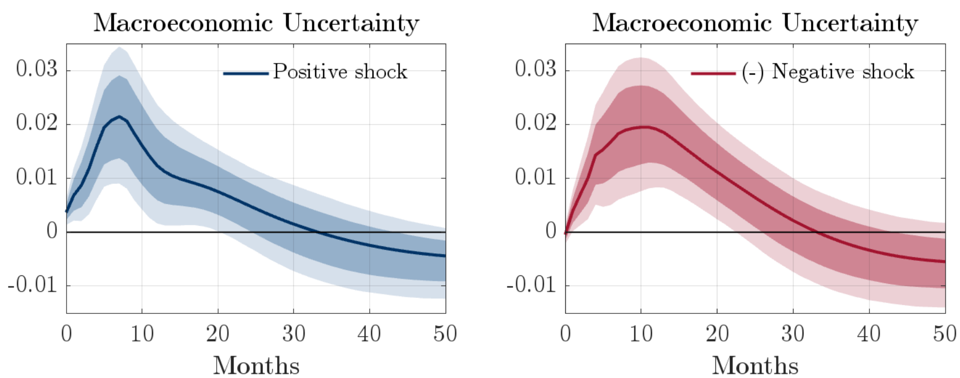

Figure C.2: Nonlinear impulse responses to a positive (blue line) and a negative (red line) oil supply news shock estimated using a nonlinear Proxy-SVAR. The core six-variable model is extended including the macroeconomic uncertainty indicator of Jurado et al. (2015). Impulse responses to a negative shock are multiplied by minus one. The shaded areas represent the 68% and 90% confidence bands.

(2015). The result, reported in Figure C.2, is that the effects on macroeconomic uncertainty are symmetric, suggesting that the second-round effects do not come from macroeconomic uncertainty.