Estimating the Macroeconomic Effects of Oil Supply News*

Lorenzo Mori
University of Padova

Gert Peersman
Ghent University

This version: January 2025

Abstract

A common approach for estimating the macroeconomic effects of oil supply news employs SVAR-IV models identified using changes in oil futures prices around OPEC quota announcements as an instrument. We show that the reduced-form oil price innovations, structural shocks, and the instrumental variable in such estimations are all Granger-caused by financial variables, indicating informational deficiencies in the VAR model and contamination of the instrument. To resolve these issues, we incorporate financial indicators into the econometrician's information set, yielding significantly different results: the shocks contribute less to oil price variation, are less inflationary, and induce a sharper short-term output contraction. The revised results also exhibit greater stability over time and the disappearance of puzzling responses. Notably, we find that oil supply news accounts for a substantial share of S&P 500 volatility. Finally, we identify similar informational deficiencies in other prominent oil-market SVAR models, suggesting this problem is pervasive in oil-market research. Our findings highlight the critical role of financial variables in accurately analyzing the causes and consequences of oil-market shocks.

JEL classification: C32, C36, E31, E32, F31, Q43.

*We thank Eftrem Castelnuovo, Gerdie Everaert, Mario Forni, Diego Känzig and seminar participants at Ghent University for useful comments and suggestions.

1 Introduction

The macroeconomic implications of oil market disturbances have long been an important topic for academic research and policymakers. Känzig (2021) makes a major contribution by exploiting institutional features of OPEC to estimate how changes in oil supply expectations—oil supply news shocks—affect the oil market and the macroeconomy. Specifically, Känzig (2021) quantifies the changes in oil futures prices within a narrow (daily) window around OPEC quota announcements. Given that expectations about global economic conditions should be priced in at the time of the announcement, such high-frequency price changes can plausibly be considered as consequences of revisions in market expectations about future oil supply due to the announcement. These price changes are then used as an external instrumental variable in a monthly oil market structural vector autoregression (SVAR-IV) model to estimate the macroeconomic effects of oil supply news. Subsequently, numerous scholars have used the instrument or the structural shocks series to explore related research questions.[1]

While the novel instrument and the empirical results derived from it have greatly improved our understanding of oil-market dynamics, certain impulse responses to oil supply news shocks remain puzzling and unstable across sample periods (see, e.g., Degasperi, 2021; Castelnuovo et al., 2024). This is illustrated in Figure 1, which replicates Känzig's original VAR model—augmented with the one-year US interest rate—for 1974-2017 and 1982-2017, respectively.[2] In both samples, an oil supply news shock that raises oil prices leads to a temporary increase in world industrial production, which is surprising given that most oil-exporting countries are not part of the index. Whereas global economic activity starts to fall significantly after more than one year in the longer sample, the impact is never significantly negative in the shorter sample. Furthermore, we observe a less persistent impact on US output and inflation in the shorter sample period, and even an opposite monetary policy response. Finally, the responses of world oil production and oil inventories differ markedly between both samples. Although it is plausible that the dynamics of oil market shocks shifted in the early 1980s, such large differences are striking given that more than 80% of the observations in both samples overlap.

[1]Examples are Bruns (2021), Degasperi (2021), Gagliardone and Gertler (2023), Bruns and Lütkepohl (2023), Forni et al. (2023), Nguyen et al. (2024), Castelnuovo et al. (2024), Kilian (2024), Degasperi et al. (2024), Caravello and Martinez-Bruera (2024), Patzelt and Reis (2024). A key advantage of the instrumental variable is the possibility of having a longer sample period for the VAR model than the proxy. It also offers flexibility in estimating, for example, nonlinearities.

[2]The baseline specification in Känzig (2021) is estimated over the sample period 1974-2017 and includes six variables: the real price of crude oil, world oil production, world oil inventories, world industrial production, US industrial production, and the US consumer price index. Due to the availability of the instrument, the contemporaneous impact matrix is estimated over a shorter period, 1983-2017. In the figure, we augment this model with the one-year US interest rate, as the instability of the monetary policy response is one of the puzzling findings (Castelnuovo et al., 2024). We use the one-year rate to also capture information on unconventional monetary policy actions during the zero lower bound period, as pointed out by Gertler and Karadi (2015). It is important to emphasize, however, that these choices do not affect the conclusions. In the appendix (Figure A2), we show the instability in the baseline six-variables specification. Finally, as in Känzig (2021), the confidence intervals are constructed using a block bootstrap with a block size of 24 and 10,000 replications.

Figure 1: Instabilities in the Känzig (2021) VAR-model

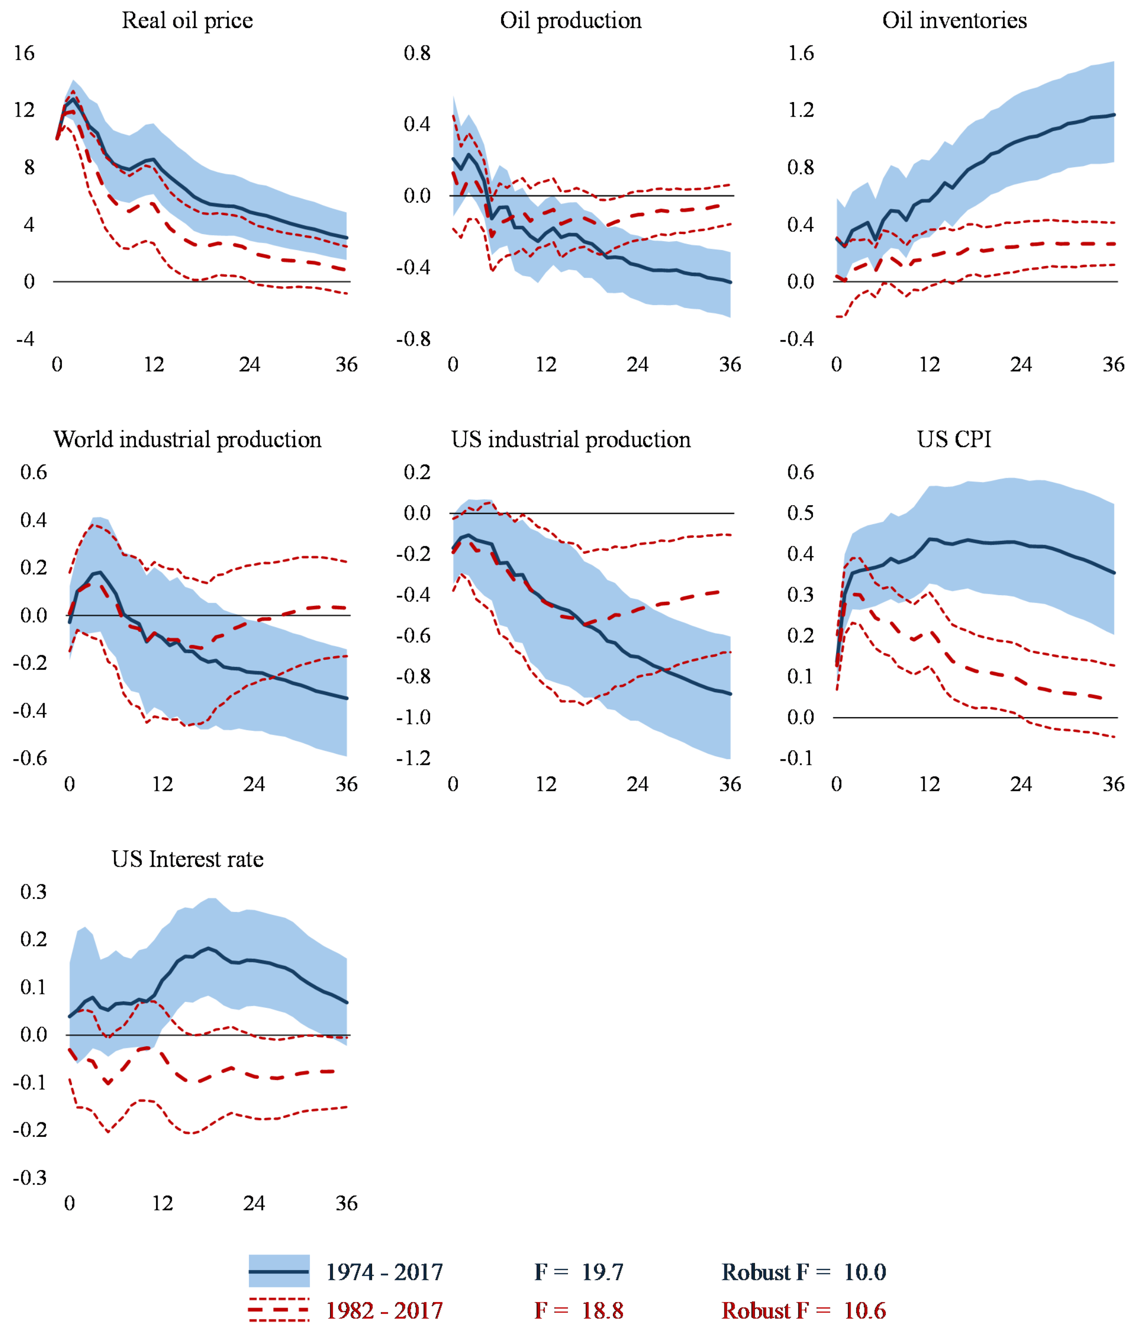

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Känzig (2021) VAR model augmented with the US 1-year interest rate estimated over sample periods 1974-2017 and 1982-2017, respectively. Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

Moreover, time variation cannot explain the instability of the contemporaneous effects, as these coefficients are estimated over the same period in both VARs. As pointed out by Ramey (2016) and Miranda-Agrippino and Ricco (2023) in the context of monetary policy VARs identified with an external instrument, puzzling responses and sample instabilities often stem from model misspecification, such as omitted variables, or contamination of the IV by other shocks.

In this paper, we first demonstrate that there are indeed distortions due to omitted variables, as well as contamination of the instrument, in Känzig (2021) and related studies. Our analysis is based on Forni and Gambetti (2014), Stock and Watson (2018), and particularly Miranda-Agrippino and Ricco (2023), who derive the general conditions for identification in SVAR-IV models when only a subset of structural shocks is of interest. A first requirement is that the shock of interest is partially invertible; that is, it can be derived from a linear combination of the current and past values of the VAR variables. Non-invertibility indicates information insufficiency in the VAR model, meaning that the econometrician's information set does not span the true information of agents to recover the structural shocks. Phenomena such as anticipation and foresight can lead to non-invertible VAR representations.

A necessary condition for partial invertibility is that no other variables Granger-cause the residuals of the relevant VAR equation, and that the structural shocks are orthogonal to past information. We show, however, that the common factors identified in McCracken and Ng (2016) jointly Granger-cause the oil price residuals and the oil supply news shocks. Further investigation reveals that this result is driven by the factor with a strong loading on financial variables. This observation is corroborated when considering a set of financial variables: equity prices (S&P 500), stock price volatility (VXO), and interest rate spreads (excess bond premium) all Granger-cause the innovations and/or shocks. Hence, the condition of partial invertibility is not fulfilled, indicating informational deficiencies of the VAR model and non-fundamental shocks. Omitted variables may also imply that the VAR does not correctly captures the dynamic responses at longer horizons after the shocks.

Miranda-Agrippino and Ricco (2023) show that, under partial invertibility, in addition to the standard relevance and contemporaneous exogeneity conditions, the IV must also satisfy a "limited lead-lag exogeneity condition". This condition ensures that the instrument correlates with the VAR residuals only via the invertible shock of interest. Conversely, the IV must be uncorrelated at any leads and lags with non-invertible shocks. Otherwise, the instrument is contaminated by other shocks, inducing bias in the results. It requires that the component of the instrument that is orthogonal to the VAR information set must be unpredictable by variables not included in the VAR model. We show that the changes in oil future prices induced by the OPEC announcements are also Granger-caused by financial variables, particularly the VXO, implying contamination of the IV by other shocks.[3] Note that

[3]In the appendix, we document a negative relationship between lagged changes in the VXO and the IV. The stability of

this condition must also be satisfied for methods that do not require invertibility, such as SVARs that include the IV as an internal instrument or local projections with controls (see Plagborg-Møller and Wolf, 2021), indicating that such estimates are also distorted. Intuitively, when the internal instrument is correlated with past shocks that are not included in the VAR, its innovations will also capture these shocks. The same applies to local projections using the lags of the VAR variables as controls.

In general, partial invertibility can be restored by incorporating additional variables into the VAR model to align the econometrician's information set with the agents' information set. Similarly, the "limited lead-lag exogeneity condition" can be fulfilled by finding an information set that renders the IV conditionally exogenous. We show that both requirements can be met by including the US one-year interest rate, the S&P 500, and the VXO in the VAR model. By doing this, the oil price residuals, the oil supply news shocks, and the IV are no longer Granger-caused by external variables. These financial variables should also enrich the dynamics of the VAR system at longer horizons.

The VAR model augmented with the financial variables yields results that differ markedly from the original estimates, both statistically and quantitatively. Specifically, the rise in oil prices is less persistent, the shocks contribute considerably less to oil price variation, there is a much sharper short-term decline in world and US industrial production, and the inflationary effects are lower and less persistent. For example, an oil supply news shock that raises real oil prices by 10% leads to a 0.57% decline in world industrial production after one year, compared to 0.11% in the original VAR, and the peak response of US consumer prices is only half. The shocks account for 27% of the historical variance in real oil prices, compared to 73% in the original model. The results also exhibit greater stability over time, while the puzzling responses disappear. Thus, while the informational deficiencies and contamination of the instrument are quantitatively important, they can be resolved by incorporating financial variables into the econometrician's information set. Doing so also renders the instrumental variable less susceptible to critiques related to sample period instability and output puzzles. For instance, Degasperi (2021) interprets the output puzzle as evidence that the instrument also captures revisions in expectations about oil demand, but this puzzle vanishes in the augmented VAR model.

On the other hand, the augmented VAR results reveal that oil supply news shocks account for a substantial share of S&P 500 volatility. Specifically, a shock that raises oil prices by 10% on impact leads to a peak decline of 4.9% in the S&P 500. Overall, oil supply news shocks accounts for about 40% of the S&P 500's variation during the sample period, which is significantly larger than the contributions reported in previous studies examining the relationship between the oil market and equity prices. Furthermore, we show that the financial variables are essential for distinguishing between

the coefficients across subsamples, along with a symmetry between increases and decreases in the VXO, suggests that this relationship is not merely a result of sample correlation. Identifying the precise reasons and mechanisms underlying this relationship is, however, beyond the scope of this paper.

oil price increases due to expected supply shortfalls and those caused by anticipated future demand surges. The Granger-causality tests indicate that financial markets often anticipated oil supply news in advance, but we show that the information contained in financial variables is also crucial to disentangle expected supply and demand at the moment of impact. Broadly, our approach aligns with studies that augment SVAR models with financial variables to identify, for example, technology shocks (Beaudry and Portier, 2006) or monetary policy shocks (Caldara and Herbst, 2018).

Finally, because most oil-market SVAR models in the literature—including those not identified with an external instrument—rely on similar information sets, our findings suggest that informational deficiencies may be a widespread issue in oil-market research. We explore this for two prominent oil-market SVAR models: Baumeister and Hamilton (2019) and Kilian (2009). Our analysis reveals that the structural shocks in both models are indeed Granger-caused by omitted variables, although it remains unclear whether the distortions are substantial. Unlike oil supply news shocks, the shocks in these studies are not solely predictable by financial variables, while expanding these VAR models introduces complexity into the identification strategy. We leave this extension for future research. A possible avenue is the use of FAVAR models, as in Juvenal and Petrella (2015), Aastveit et al. (2015) and Stock and Watson (2018).

Section 2 discusses the conditions for identification in SVAR models with an external instrument. Section 3 reports the predictability of the VAR residuals, structural shocks, and IV. Sections 4 presents the results of the VAR model incorporating financial variables, while section 5 documents the predictability of the shocks in other popular oil-market VARs. Finally, section 6 concludes.

2 Identification with External Instruments in Structural VARs

In the SVAR literature, the dynamics of an $n \times 1$ vector of observed endogenous variables Y_t are described by the following reduced-form VAR(p) model:

$$Y_t = b + B_1 Y_{t-1} + \cdots + B_p Y_{t-p} + \mu_t \tag{1}$$

where p is the lag order, μ_t is an $n \times 1$ vector of reduced-form innovations, b is an $n \times 1$ vector of constants, and $B_1, \ldots, B_p$ are $n \times n$ coefficient matrices. Furthermore, it is assumed that the reduced-form innovations are related to the structural shocks via a linear mapping:

$$\mu_t = S\varepsilon_t \tag{2}$$

where S is an $n \times n$ structural impact matrix and ε_t is an $n \times 1$ vector of structural shocks. Although the structural shocks are not directly observable, they can be recovered from the reduced-form innovations if S is invertible (non-singular): $\varepsilon_t = S^{-1}\mu_t$. Obtaining the elements of S also requires identifying restrictions. Stock and Watson (2012) and Mertens and Ravn (2013) demonstrate that the coefficients of a single column i of S, which is sufficient for retrieving the corresponding structural shock i, can be estimated with an external instrument Z_t if the standard instrument relevance and exogeneity conditions are fulfilled; that is, the instrumental variable must be correlated with the shock of interest and uncorrelated with all other structural shocks on impact.

The assumption of invertibility is crucial but non-trivial. It implies that all structural shocks in the economy can be derived from linear combinations of the VAR residuals; that is, based on current and lagged values of Y_t.[4] This is unlikely the case in small VAR models due to omitted variables and insufficient lag length. The presence of anticipation and foresight can also lead to non-invertible VAR models. Non-invertibility indicates that the VAR model suffers from information insufficiency and has non-fundamental shocks. However, even if the VAR does not contain enough information or variables to retrieve all structural shocks, it can be sufficient to identify a specific shock or a subset of shocks; that is, only the invertibility of the shock(s) of interest—or partial invertibility—is required (Forni and Gambetti, 2014; Miranda-Agrippino and Ricco, 2023). Partial invertibility holds if the structural shock(s) of interest can be accurately recovered as a linear combination of the VAR residuals.

While Stock and Watson (2012) and Mertens and Ravn (2013) assume full invertibility, Miranda-Agrippino and Ricco (2023) formalize the conditions that an instrumental variable must satisfy to achieve correct identification in SVAR-IV models under partial invertibility. Specifically, let $\varepsilon_t^{1:m}$ be the m invertible structural shocks from a VAR in Y_t, and $\varepsilon_t^{m+1:n}$ the remaining $n-m$ non-invertible shocks. Define $Z_t^{\perp} = Z_t - Proj\left(Z_t | \mathscr{H}_{t-1}^{Y}\right)$, where Z_t is a candidate IV for the shock of interest ε_t^1 and $\mathscr{H}_{t-1}^{Y}$ is the Hilbert space generated by all the observations of Y_t up to time t. The impact effects of ε_t^1 on Y_t are identified if Z_t satisfies:

$$E\left[\varepsilon_t^1 Z_t^{\perp}\right] \neq 0 \tag{3}$$

$$E\left[\varepsilon_t^{2:n} Z_t^{\perp}\right] = 0 \tag{4}$$

$$E\left[\varepsilon_{t-j}^{m+1:n} Z_t^{\perp}\right] = 0 \text{ for } j \neq 0 \tag{5}$$

In addition to the standard instrument relevance and exogeneity conditions (i.e., equations 3 and 4), the instrumental variable must also fulfill a limited lead-lag exogeneity condition (equation 5). This condition ensures that the instrument correlates with the VAR residuals only through the invert-

[4]Note that all the structural shocks should not explicitly be identified; they should be captured by the VAR model.

ible shock of interest. In particular, while the instrument may correlate with leads or lags (but not contemporaneously) of other invertible shocks in the system, it must be uncorrelated at any leads and lags with all non-invertible shocks. This condition arises because of the VAR model's dynamics: if non-invertible shocks correlate at any leads or lags with the instrument, the effects of these shocks will also be captured by the instrument, leading to biased estimates (Miranda-Agrippino and Ricco, 2023). Conversely, leads and lags of other partial invertible shocks can affect the instrument without distorting the identification because they do not enter the VAR residuals.

Note that the lead-lag exogeneity condition is automatically fulfilled under full invertibility because the VAR residuals are a linear combination of the contemporaneous structural shocks only. Furthermore, Miranda-Agrippino and Ricco (2023) show that this condition must also be met for methods that do not require (partial) invertibility, such as SVAR models that include the IV as an internal instrument or local projections with the lagged observables Y_t as controls. Intuitively, if the shocks contaminating the instrument are unknown, the only way to achieve identification is to control for all possible shocks. This is equivalent to including these variables in the VAR to ensure invertibility. This equivalence is the so-called "no-free lunch" result of Stock and Watson (2018).

3 Oil Supply News Shocks: Invertibility and Exogeneity of the IV

In this section, we investigate whether the partial invertibility and limited lead-lag exogeneity conditions are fulfilled in SVAR-IV models identified using changes in oil futures prices around OPEC quota announcements as an instrument. For this analysis, we use the original 6-variable VAR model and sample period from Känzig (2021) as the "baseline" VAR model.[5]

3.1 Partial Invertibility

Forni and Gambetti (2014) provide a testing procedure for (partial) invertibility and informational sufficiency. A necessary condition is that no other variables Granger-cause the reduced-form oil price innovations and the oil supply news shocks, respectively. The underlying idea is that if these "omitted variables" do not help to predict the VAR variables (innovations), the VAR model must contain the same information; that is, the variables in the VAR convey all relevant information, and there is no gain from including these other variables. Conversely, if forecasts improve by adding the variables to

[5]Känzig (2021) adheres to the standard practice in the literature by assuming full invertibility. He examines how the results depend on the information in the VAR by adding other variables (one at a time) to the baseline VAR model. The robustness of the results suggested there are no informational deficiencies. Note that the results are similar for the SVAR model augmented with the one-year interest rate shown in Figure 1, as well as for other modifications to the baseline specification used in the literature. An exception is Castelnuovo et al. (2024), who have followed our recommendations.

the regressions, the VAR is missing this information, indicating a failure of the invertibility condition. Furthermore, a necessary condition for a shock to be "structural" is orthogonality to the past of all variables that drive the economy. Hence, once a shock has been identified and estimated, its "structuralness" can be evaluated by testing for orthogonality with respect to the lags of variables omitted from the VAR model.

Because it is not possible to test all macroeconomic variables simultaneously, and testing each of them separately would lead to rejection of informational sufficiency due to Type I error, Forni and Gambetti (2014) propose to use the principal components of a large dataset capturing all relevant macroeconomic information as a starting point. Following Miranda-Agrippino and Ricco (2023), we use the common factors extracted from the large monthly database of McCracken and Ng (2016), which is systematically updated for economic research.[6] Based on the Bai and Ng (2002) criterion, McCracken and Ng (2016) identify eight relevant common factors in the macroeconomic series, which we use for the Granger-causality tests.

The left part of Table 1 summarizes the results of the partial invertibility tests for the baseline six-variables VAR. We report the p-values of the (robust) F-statistic for the null hypothesis that the lagged variables do not Granger-cause the oil price residuals and oil supply news shocks, respectively. Given the sensitivity of such tests for the number of lags (L), we report the results for $L = 6$ and $L = 12$, the latter being the number of lags in the VAR. The first row reveals that the common factors identified in McCracken and Ng (2016) jointly Granger-cause the oil price residuals and the structural shocks, suggesting that there is informational insufficiency in the VAR to retrieve the shocks, and that the estimated shocks are not orthogonal to past macroeconomic data.

A closer inspection of the individual factors, also shown in Table 1, reveals that the joint significance is driven by the factor F7. Notably, according to McCracken and Ng (2016), this factor (as well as F6) explains much of the variation in stock market variables. Specifically, the top 4 series that load most on F7 are S&P 500, S&P Industrials, S&P dividend yields, and the VXO. This indicates that financial variables are the source of the informational insufficiency in the VAR. This observation is corroborated when we perform the Granger-causality tests for several financial indicators used in macroeconomic research. As can be observed in the table, the "Global Financial Cycle", a common factor identified in Miranda-Agrippino and Rey (2020) that explains an important share of the variation of financial aggregates around the world, also Granger-causes the oil price innovations and oil supply news shocks. In contrast, the OECD Global Composite Leading Indicator, reflecting expecta-

[6]We have used the “2023-10” vintage (downloaded in November 2023). Note that the time series in the database are mainly for the US. We are not aware of comparable monthly databases at the global level. However, considering the hegemonic role of the US in international markets, particularly in the oil and financial markets, this database should contain the bulk of relevant information.

Table 1: P-values Granger causality tests - baseline VAR-model

	Oil price residuals		Oil supply news shocks		Z_t		$Z_t^{\perp}$	
	$L=6$	$L=12$	$L=6$	$L=12$	$L=6$	$L=12$	$L=6$	$L=12$
Common factors (all)	0.40	**0.03**	**0.04**	**0.00**	0.34	0.51	0.28	**0.00**
F1	0.99	1.00	0.99	0.98	0.32	0.27	0.68	0.86
F2	0.99	1.00	0.99	1.00	0.14	0.17	0.87	1.00
F3	0.99	1.00	0.98	1.00	0.64	0.73	0.98	1.00
F4	0.31	0.41	0.22	0.43	0.13	0.13	0.40	0.12
F5	0.87	0.49	0.49	0.18	0.89	0.56	0.87	0.61
F6	0.19	0.28	0.08	0.18	0.07	0.14	0.16	0.09
F7	**0.01**	**0.01**	**0.02**	**0.03**	0.43	0.14	0.41	0.55
F8	0.74	0.81	0.74	0.81	0.40	0.53	0.53	0.22
Global Financial Cycle	**0.02**	**0.01**	0.06	**0.01**	0.36	0.56	0.68	0.75
OECD Global CLI	0.52	0.78	0.49	0.89	0.17	0.26	0.55	0.70
S&P 500	**0.04**	**0.03**	**0.05**	**0.04**	0.44	**0.05**	0.37	0.10
MSCI World	**0.01**	**0.02**	**0.05**	0.08	0.72	0.49	0.90	0.65
VXO	0.06	**0.05**	**0.04**	**0.02**	**0.00**	**0.04**	**0.05**	0.24
Financial uncertainty	0.65	0.93	0.91	0.99	**0.00**	**0.00**	**0.01**	**0.02**
Macro uncertainty	0.33	0.13	0.36	0.22	**0.04**	0.09	0.33	0.68
Excess bond premium	0.26	0.16	0.09	**0.03**	**0.04**	0.24	0.41	0.85
BAA-AAA spread	0.81	0.96	0.68	0.74	**0.05**	0.26	0.50	0.76
US 1-year interest rate	0.33	0.35	0.22	0.35	0.26	0.32	0.43	0.56
USD nominal effective exchange rate	0.62	0.94	0.51	0.85	0.77	0.92	0.89	0.70

Note: P-values of (robust) F-statistic for the null hypothesis that the lagged variables do not Granger-cause the oil price residuals, oil supply news shocks, the IV, and the IV orthogonalized to the impact of the lagged VAR variables. The baseline VAR includes six variables: the real price of crude oil, world oil production, world oil inventories, world industrial production, US industrial production, and the US consumer price index. L is the number of lags used in the test. For S&P 500 and MSCI, we consider log differences. The common factors are obtained from McCracken and Ng (2016). Financial and macro uncertainty are collected from Ludvigson et al. (2021). The Global Financial Cycle are obtained from Miranda-Agrippino and Rey (2020), the S&P 500 from Robert Schiller's webpage, and the Excess bond premium from Gilchrist and Zakrajsek (2012). Numbers are in bold when p<0.05.

tions about global real economic conditions, cannot predict the residuals and shocks. Furthermore, the S&P 500, the MSCI World, and the VXO all Granger-cause the oil price innovations and structural shocks. Finally, there is some weak evidence that the excess bond premium of Gilchrist and Zakrajšek (2012), a popular variable in VAR models estimating the effects of US monetary policy shocks, can also predict the shocks. In sum, the condition of partial invertibility is not fulfilled, indicating informational deficiencies of the VAR model and non-fundamental shocks.

3.2 Exogeneity of OPEC Announcements

As discussed in section 2, the instrument must be uncorrelated with all non-invertible shocks at any leads and lags. Conversely, leads and lags of partially invertible shocks in the VAR can be correlated with the instrument, as they do not enter the VAR residuals. This is why the lead-lag condition in equation 5 applies to the orthogonalized instrumental variable $Z_t^{\perp}$. Miranda-Agrippino and Ricco (2023) further show that leads, lags, or even contemporaneous realizations of the non-invertible shocks can contaminate the instrument only via their projectable component in the space spanned by past realizations of the VAR variables Y_t. Granger-causality tests for "omitted variables" can thus also be used to evaluate the limited lead-lag exogeneity condition of the external instrument.

The results are also presented in Table 1. We apply Granger-causality tests to the instrument, as well as the instrument orthogonalized to the impact of the lagged variables included in the VAR. The orthogonalized instrument is equivalent to the structural shocks series identified by ordering the instrument first in a recursive VAR, following the approach of Plagborg-Møller and Wolf (2021). These tests can thus also be viewed as a direct assessment of the "structuralness" of the shocks obtained from such a VAR, as well as local projections that include the lagged VAR variables as controls. Table 1 shows that the principal components jointly Granger-cause $Z_t^{\perp}$ when $L = 12$, but this result is not confirmed for Z_t and $L = 6$, nor for the individual factors. On the other hand, there is strong indication that OPEC announcements are Granger-caused by uncertainty indicators: both the VXO and the financial uncertainty indicator of Ludvigson et al. (2021) Granger-cause the instrument. This implies that the lead-lag exogeneity condition is also not fulfilled. In the appendix, we further explore the relationship between past VXO observations and the instrument. The analysis reveals that increases (decreases) in the VXO tend to be followed by OPEC announcements that lower (raise) oil prices. This negative relationship appears to be remarkably stable across subsamples and is symmetric for positive and negative changes in the VXO, indicating that the instrument's predictability is neither accidental nor the result of outlier observations in the sample period.

Table 2: P-values Granger causality tests - VAR-model augmented with financial variables

	Oil price residuals		Oil supply news shocks		$Z_t^{\perp}$	
	$L=6$	$L=12$	$L=6$	$L=12$	$L=6$	$L=12$
Principal components (all)	0.98	1.00	0.54	0.16	0.61	0.14
F1	1.00	1.00	0.99	0.93	0.97	0.99
F2	1.00	1.00	0.97	0.98	0.88	1.00
F3	1.00	1.00	1.00	1.00	0.99	1.00
F4	0.97	0.99	0.98	0.93	0.92	0.68
F5	0.98	0.96	0.91	0.62	0.97	0.97
F6	0.64	0.88	0.86	0.93	0.81	0.87
F7	0.96	1.00	0.97	0.96	0.97	0.99
F8	0.75	0.89	0.47	0.79	0.52	0.34
Global Financial Cycle	0.19	0.43	0.86	0.87	1.00	0.81
OECD Global CLI	0.96	0.87	0.85	0.92	0.84	0.85
S&P 500	1.00	1.00	1.00	0.99	1.00	0.98
MSCI World	0.07	0.26	0.45	0.85	0.98	0.74
VXO	1.00	1.00	1.00	1.00	1.00	1.00
Financial uncertainty	0.53	0.57	0.14	0.24	0.67	0.77
Macro uncertainty	0.63	0.40	0.94	0.48	0.84	0.92
Excess bond premium	0.24	0.22	0.07	0.19	0.56	0.94
BAA-AAA spread	0.96	1.00	0.92	0.98	0.55	0.77
US 1-year interest rate	1.00	1.00	1.00	1.00	1.00	1.00
USD nominal effective exchange rate	0.82	0.98	0.40	0.75	0.92	0.79

Note: P-values of (robust) F-statistic for the null hypothesis that the lagged variables do not Granger-cause the oil price residuals, oil supply news shocks, and the IV orthogonalized to the impact of the lagged VAR variables. The VAR model contains all variables of the baseline VAR, augmented with US 1-year interest rate, the S&P 500, and the VXO. L is the number of lags used in the test. Financial and macro uncertainty are collected from Ludvigson et al. (2021). The Global Financial Cycle are obtained from Miranda-Agrippino and Rey (2020), the S&P 500 from Robert Schiller's webpage, and the Excess bond premium from Gilchrist and Zakrajsek (2012). Numbers are in bold when p<0.05.

4 Oil-Market SVAR Model Augmented with Financial Variables

This section addresses the informational deficiencies and instrument contamination by expanding the econometrician's information set. It then evaluates the significance of this extension for the estimation results, explores the impact of oil supply news on financial markets, analyzes the contribution of financial variables for identification, and assesses the stability of the results across sample periods.

4.1 Invertibility and Lead-Lag Exogeneity of the Augmented SVAR Model

A standard solution to resolve informational deficiencies is to include the omitted variables in the VAR until invertibility and orthogonality of the shocks are achieved (Forni and Gambetti, 2014; Stock and Watson, 2018). Similarly, contamination of the IV can be mitigated by expanding the information set with variables that render the instrument conditionally exogenous (Miranda-Agrippino and Ricco, 2023). The analysis in section 3 suggests that financial variables are crucial for addressing both issues. We incorporate the S&P 500, the VXO, and the monetary policy rate into the baseline VAR. These variables are forward looking and capture different dimensions of expectations, including those of the central bank. The financial variables and the monetary policy response could also enrich the VAR system's dynamics. Most importantly, the invertibility and lead-lag exogeneity conditions are fulfilled by incorporating these variables. This is demonstrated in Table 2, which presents the Granger-causality tests for the augmented VAR model. Specifically, none of the common factors or other financial variables Granger-cause the oil price innovations, the identified oil supply news shocks, or the external instrument conditional on the past VAR variables.

The results are very similar when we incorporate more or alternative combinations of financial variables into the VAR model. Due to their ready accessibility, unlike for example the Global Financial Cycle or the financial uncertainty index of Ludvigson et al. (2021), the S&P 500, the VXO and the interest rate form the most obvious combination. It is, however, important to include all three variables in the VAR model. For instance, the structural shocks can still be predicted by the common factors and the excess bond premium when the interest rate is excluded from the information set.

4.2 Revisiting the Macroeconomic Effects and Relevance of Oil Supply News

Figure 2 shows the impulse responses of the oil-market SVAR model augmented with financial variables. To evaluate the quantitative impact of incorporating the financial variables, the figure also displays the responses of the baseline VAR. Formal statistical comparisons between the impulse responses of both models are shown in Figure 3. These tests are conducted by nesting the VARs within

Figure 2: Impact of oil supply news shocks in an SVAR augmented with financial variables

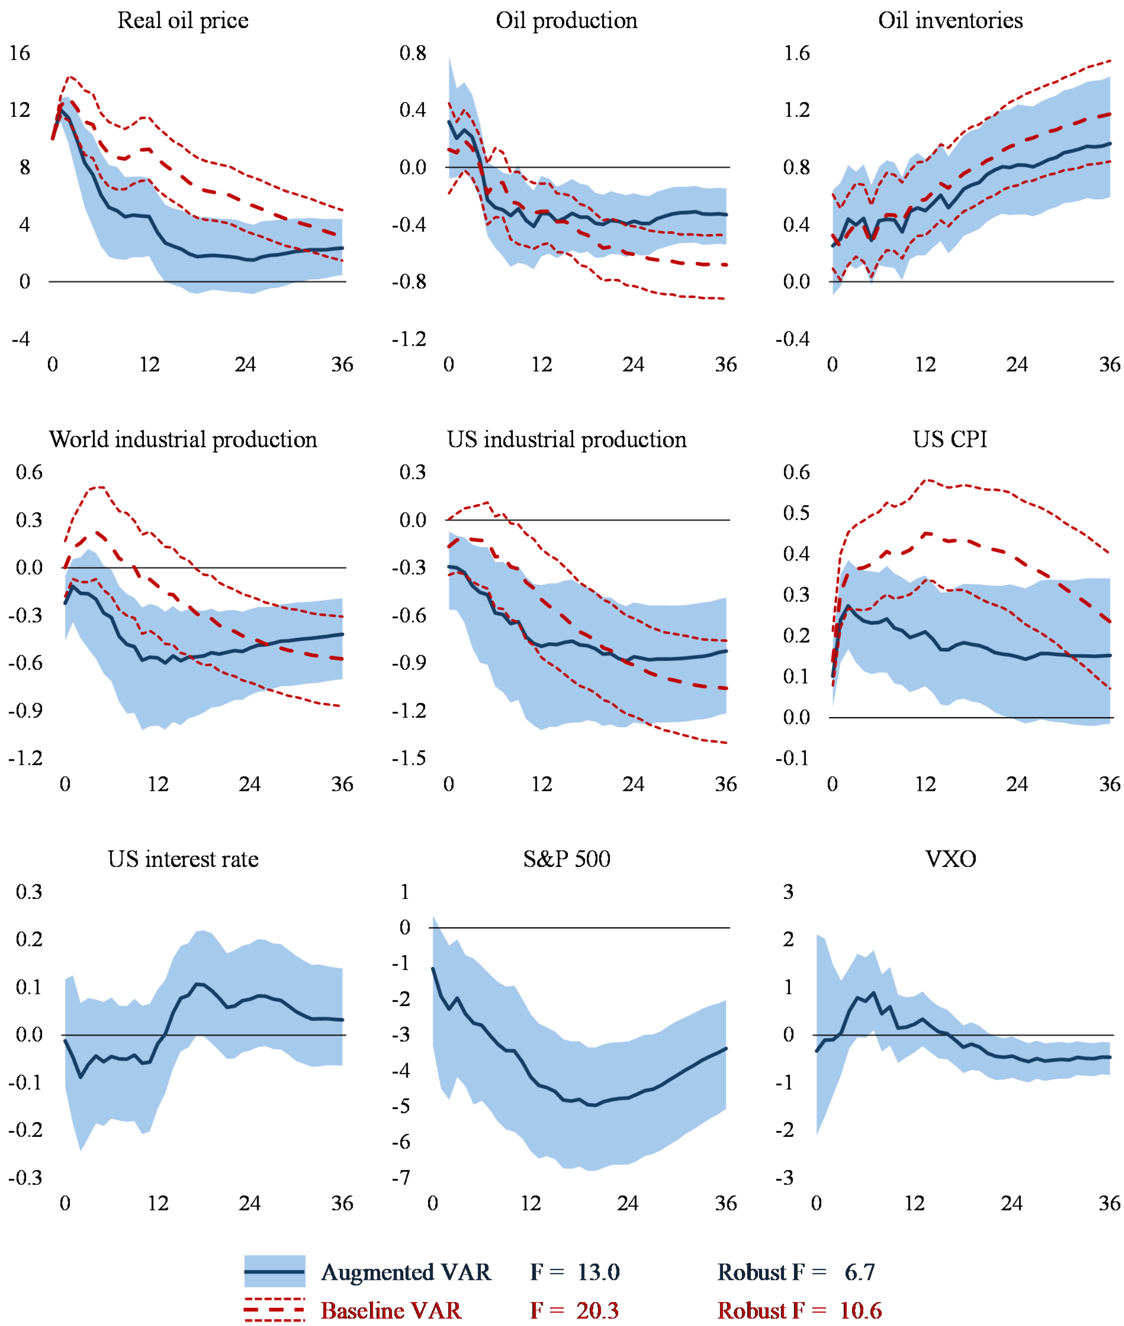

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Baseline VAR (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

Figure 3: Tests for differences between augmented and baseline VAR impulse responses

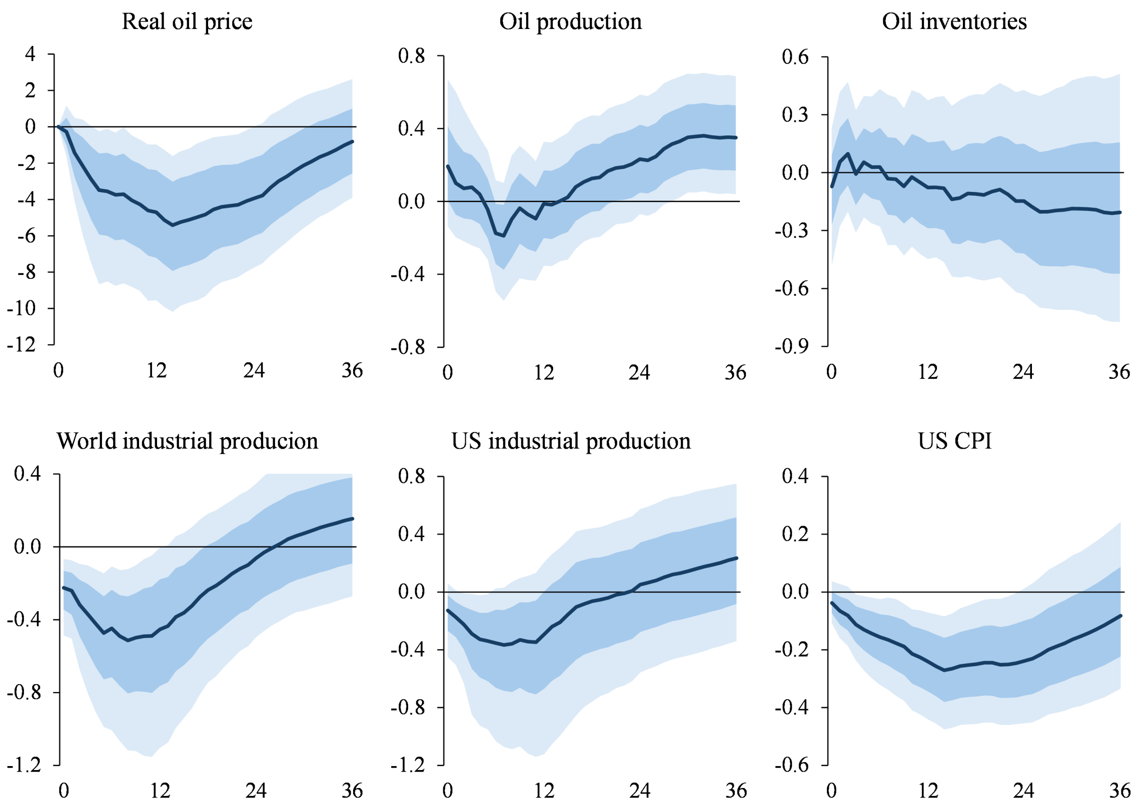

Note: Estimated differences between impulse responses of the VAR model augmented with financial variables and the baseline VAR for oil supply news shocks that raise oil prices by 10% on impact. Monthly horizon. 68% and 90% confidence intervals constructed using a moving block bootstrap.

the bootstrap procedure. Specifically, for each bootstrap sample of the 9-variable VAR model with financial variables, we also estimate the 6-variable baseline VAR and compute the difference between the impulse response functions. The 16-84 and 5-95 percentiles of these differences (i.e., 68% and 90% confidence intervals) are depicted in Figure 3 and used to assess the statistical significance of the omitted variable bias. In the appendix, Figure A3 and A4 display the impulse responses of both VAR models using alternative inference methods: the Bayesian approach of Miranda-Agrippino and Ricco (2021), and the weak-instrument robust inference method of Montiel Olea et al. (2021), respectively. The latter addresses possible concerns about the robust F-statistic dropping below 10 in the augmented VAR. Notably, the confidence (credible) intervals from these alternative methods are somewhat narrower compared to the block bootstrap approach that we use.

Figure 2 and 3 demonstrate that the results differ significantly from the original estimates, both statistically and quantitatively. First, the rise in real oil prices following unfavorable news about future oil supply is less persistent. In the original VAR, oil prices remain 9.3% above the baseline after one year, whereas this is only 4.6% in the augmented VAR. Additionally, the "output puzzle", with world industrial production rising in the first year following an oil supply news shock, is resolved. Instead, we observe an immediate sharp decline in production, reaching its peak after one year. More specifically, world industrial production decreases by -0.57% after one year, compared to -0.11% in the baseline VAR. For policymakers, such a difference at the one-year horizon is substantial. A similar divergence occurs for US industrial production, although the magnitude is somewhat smaller. As shown in Figure 3, the differences between the output effects are also statistically significant.

The impact on inflation also differs markedly between the two specifications. In the VAR with financial variables, the peak response of US consumer prices is nearly half compared to the original VAR and exhibits much less persistence. This difference in the inflationary effects is highly statistically significant (see Figure 3). Lastly, compared to Figure 1—where the interest rate was included in the baseline VAR— there appears to be a (statistically insignificant) reversal in the point estimate of the monetary policy response in the VAR model with the richer information set.

Incorporating financial variables into the information set also revisits the relevance of oil supply news for macroeconomic and oil market fluctuations. Table 3 reports the forecast error variance of the VAR variables attributed to oil supply news shocks at horizons 0, 12, and 36 months. In the baseline VAR, the shocks account for 73% of the historical variance in real oil prices at the three-year horizon, whereas this share drops to 27% in the VAR model with financial variables. This more limited role of oil supply news shocks as a driver of oil price fluctuations aligns more closely with historical narratives and existing studies, such as Kilian (2009), which typically emphasize the significant influence of aggregate demand shocks. For instance, during the onset of the Great Financial Crisis, the real oil price collapsed by 114% between July and December 2008. As shown in the appendix (which includes historical decompositions), the baseline VAR attributes 76 percentage points of this decline to oil supply news, while the augmented VAR attributes only 31 percentage points. As discussed in section 4.3, this discrepancy arises because the shocks identified in the baseline VAR model partially capture news about (future) economic activity and oil demand.

A similar pattern emerges for the variance in US consumer prices. In particular, oil price fluctuations driven by oil supply news account for 38% of consumer price variation at the three-year horizon in the baseline model but only 9% in the financial variable-augmented VAR. However, the reverse holds for output variation. For instance, at the one-year horizon, oil supply news shocks explain 10% and 18% of world and US industrial production fluctuations, respectively, in the VAR with financial variables, compared to just 1% and 4% in the baseline model. In summary, incorporating financial variables is crucial for accurately measuring the macroeconomic effects and historical relevance of oil supply news shocks, as the omitted variables substantially affect the results.

Table 3: Forecast error variance decompositions

		Baseline VAR		VAR with financial variables	
Real oil price	0	0.83	[0.39 0.91]	0.70	[0.12 0.81]
	12	0.75	[0.24 0.81]	0.41	[0.05 0.55]
	36	0.73	[0.20 0.75]	0.27	[0.05 0.43]
Oil production	0	0.00	[0.00 0.10]	0.02	[0.00 0.13]
	12	0.03	[0.02 0.17]	0.05	[0.02 0.17]
	36	0.24	[0.05 0.40]	0.12	[0.03 0.23]
Oil inventories	0	0.05	[0.00 0.21]	0.03	[0.00 0.18]
	12	0.11	[0.01 0.28]	0.08	[0.01 0.26]
	36	0.37	[0.04 0.52]	0.24	[0.02 0.42]
World industrial production	0	0.00	[0.00 0.19]	0.08	[0.00 0.29]
	12	0.01	[0.01 0.23]	0.10	[0.01 0.33]
	36	0.11	[0.04 0.45]	0.19	[0.02 0.38]
US industrial production	0	0.03	[0.00 0.22]	0.10	[0.00 0.32]
	12	0.04	[0.01 0.27]	0.18	[0.02 0.40]
	36	0.32	[0.07 0.52]	0.31	[0.04 0.44]
US CPI	0	0.20	[0.01 0.50]	0.10	[0.00 0.33]
	12	0.48	[0.11 0.68]	0.18	[0.01 0.38]
	36	0.38	[0.06 0.50]	0.09	[0.01 0.25]
US interest rate	0			0.00	[0.00 0.11]
	12			0.01	[0.00 0.11]
	36			0.02	[0.01 0.14]
S&P 500	0			0.04	[0.00 0.36]
	12			0.20	[0.02 0.46]
	36			0.41	[0.04 0.53]
VXO	0			0.00	[0.00 0.27]
	12			0.02	[0.02 0.23]
	36			0.04	[0.03 0.20]

Note: Forecast error variance of the variables explained by oil supply news shocks at horizons 0, 12 and 36 months, with 90% confidence intervals shown in the brackets.

4.3 The Impact of Oil Supply News on Financial Markets

The impact of oil supply news on the financial variables is also presented in Figure 2 and Table 3. Several key observations emerge. First, equity prices experience a significant decline following negative news about oil supply. Specifically, an oil supply news shock that raises oil prices by 10% on impact results in a peak decline of 4.9% in the S&P 500. Notably, oil supply news accounts for 41% of the S&P 500's variation at the three-year horizon during the sample period.[7] A historical decomposition of the S&P 500's evolution, presented in Figure A5 of the appendix, highlights the particular importance of oil supply news shocks for equity price movements during the period 1990–2002.

The relevance of oil supply news shocks in explaining S&P 500 variation is significantly greater than the contributions reported in previous studies examining the relationship between the oil market and equity prices. For example, Kilian and Park (2009) find that oil supply shocks do not significantly affect US stock returns and account for only 6% of their long-run variation. While they also show that shocks to the precautionary demand for oil—potentially driven by news about future supply—have a significant negative impact, these shocks explain only 11% of the long-run variation in stock returns.

On the other hand, the responses of the one-year government bond yield and the VXO are insignificant in the short run, and the contribution of the shocks to their variability is negligible across all horizons. However, the absence of a significant response does not imply that both variables are uninformative for estimating the consequences of oil supply news. For example, these responses may play a critical role in disentangling the shocks from innovations in expected demand, which also lead to an immediate increase in oil prices but are typically accompanied by changes in the VXO and interest rates. The analysis in the next subsection suggests that this is indeed the case.

4.4 The Role of Financial Variables for Identification

The added value of incorporating financial variables to accurately estimate the effects of oil supply news is not surprising, as these variables capture expectations about future economic activity, macroeconomic uncertainty, and how central banks interpret the news. In the baseline VAR, expectations about future oil market developments are implicitly formed based on past and current oil production and economic activity. Apart from the real price of oil, the only forward looking variable is oil inventories. This is, for example, insufficient to distinguish between oil price increases due to expected supply shortfalls and anticipated future demand surges (which both typically increase oil inventories).

[7] Note that the point estimate for the forecast error variance of the S&P 500 explained by oil supply news shocks tends to lean toward the upper bound of the confidence intervals at longer horizons. For instance, at the three-year horizon, the median of all bootstrap draws is 25%, which, while lower, remains considerably.

Financial variables provide precisely this type of forward-looking information. The collapse of oil prices during the onset of the Great Financial Crisis, discussed in section 4.2, illustrates this point.

Figure 4: The role of financial variables for estimating the effects of oil supply news shocks

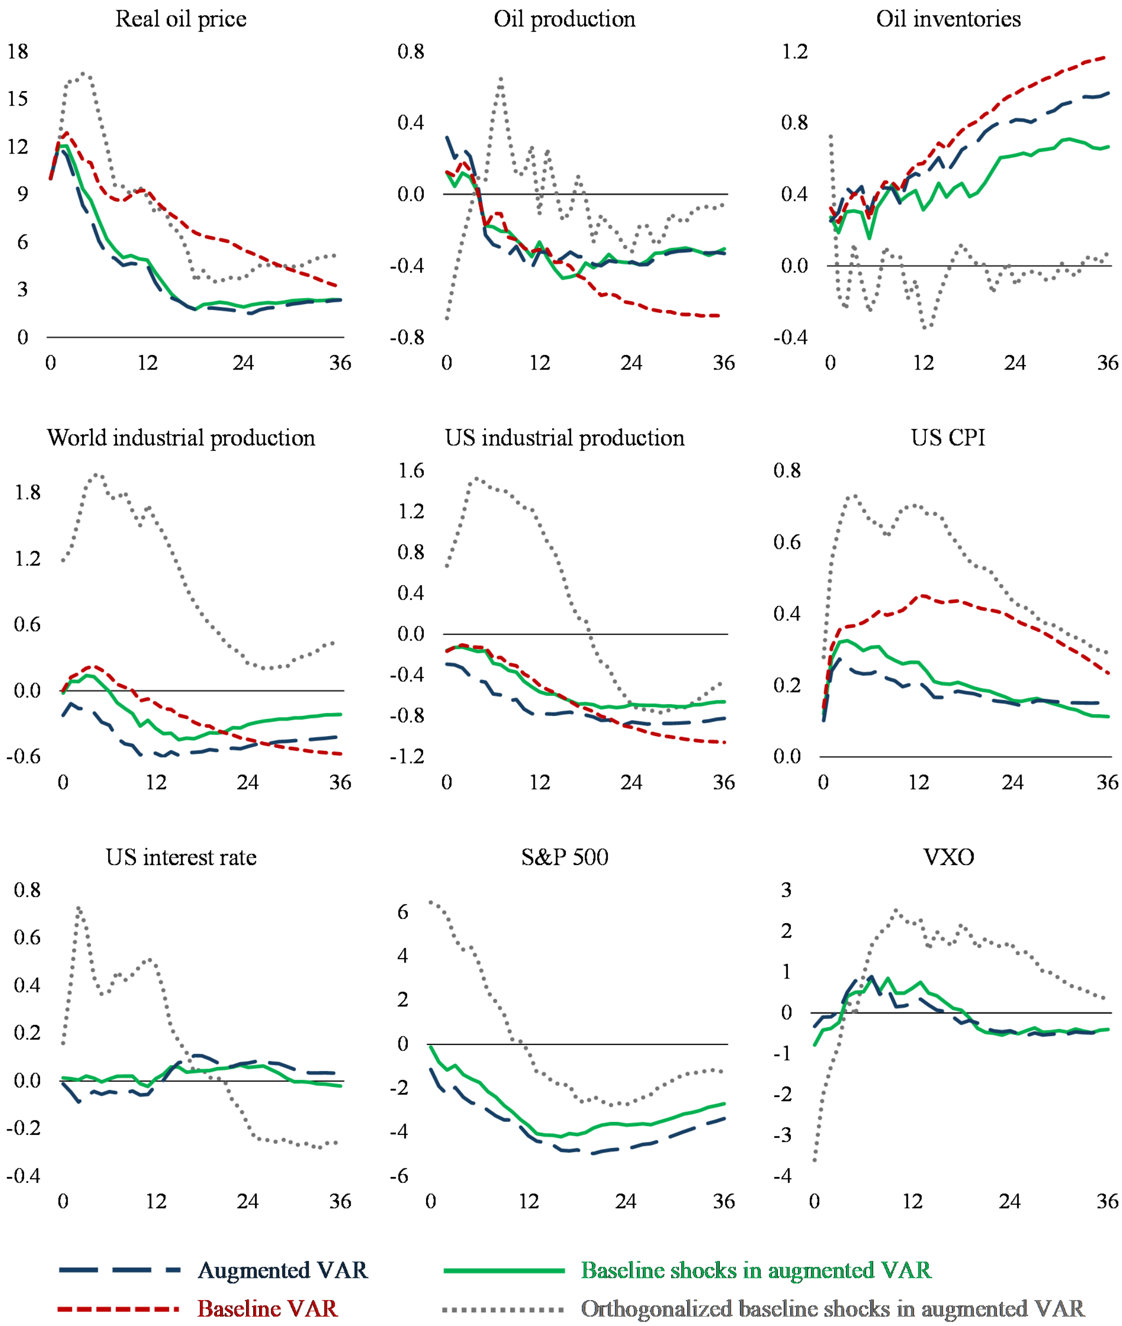

Note: Impulse responses to oil supply news shock that raises oil prices by 10% on impact for: (i) the baseline VAR, (ii) the augmented VAR with financial variables, (iii) the baseline shocks embedded in the augmented VAR, and (iv) the baseline shocks orthogonalized to the augmented shocks, embedded in the augmented VAR.

A closer inspection of the results reveals that the distortions caused by omitted financial variables cannot be attributed solely to their ability to forecast the instrument, oil price innovations, and structural shocks, as indicated by the Granger-causality tests. These tests suggest that financial markets anticipated some oil supply news shocks in prior months, distorting the estimated effects. However, financial variables are also crucial at the moment of impact—specifically for accurately retrieving the shocks from the VAR residuals, which is central to the invertibility condition discussed in Section 2.

This hypothesis is supported by two exercises reported in Figure 4. First, the green (solid) lines represent impulse responses to oil supply news shocks obtained from an "intermediate" VAR model. This model incorporates the estimated structural shocks series from the baseline VAR as the first variable in a VAR with the financial variables, identified via a recursive decomposition. By including the financial variables, this intermediate VAR accounts for the information contained in *past* financial variables and their influence on the VAR dynamics. However, because the shocks are ordered first, they remain contaminated by omitted *contemporaneous* information from financial variables.

As shown in Figure 4, the "intermediate" impulse responses deviate from the original baseline VAR results and align more closely with those of the augmented VAR. This suggests that the omitted lagged financial variables contribute to the distortions in the baseline VAR. However, a notable discrepancy remains between these intermediate responses and those of the augmented VAR (which does account for the contemporaneous information provided by the financial variables). For instance, for world industrial production, the immediate effect is 0.00 (compared to -0.22 in the augmented VAR), the impact after one year is -0.39 (instead of -0.60), while there is still a short-term output puzzle. Similarly, the impact on US industrial production is more subdued, while the inflationary effects are stronger. Another notable difference is the immediate response of the S&P 500: -0.12 in the intermediate VAR versus -1.14 in the augmented VAR. The weaker output effects, stronger inflationary effects and the negligible decline in the S&P 500 suggest that the intermediate shocks still capture endogenous oil price responses to changes in aggregate demand. This underscores the importance of contemporaneous financial information for achieving proper identification.

The second exercise presented in Figure 4 further highlights the importance of contemporaneous financial information to achieve identification. In this exercise, the baseline shocks are first regressed on the shocks obtained from the augmented VAR. The resulting residuals represent the oil supply news shocks captured in the baseline VAR but *not* in the augmented VAR. These residual shocks are then embedded as the first variable in a VAR with financial variables, identified using a recursive decomposition (i.e., they account for information contained in past financial variables). The impulse responses to these shocks (gray dotted lines) reveal that these oil price innovations are clearly driven by aggregate demand shocks. Specifically, they lead to strong increases in economic activity, inflation, and equity prices, alongside decreases in macroeconomic uncertainty and a monetary policy tightening.

Finally, following Caldara and Herbst (2018), who examined the role of bond premiums in identifying monetary policy shocks, we evaluate the presence of distortions by disregarding contemporaneous financial information by the contemporaneous elasticities of real oil prices (conditional on the identification of the structural oil supply news shocks) to changes in the variables included in the augmented VAR. Specifically, the endogenous responses of the real price of oil to unit changes in the S&P 500, the US interest rate, and the VXO are 0.70 (0.00), 2.95 (0.07), and 0.18 (0.23), respectively, with p-values in parentheses. The highly significant coefficient for the S&P 500 suggests that, through the lens of the augmented VAR, the baseline model identifies an oil supply news shock that is contaminated by the contemporaneous endogenous response of oil prices to equity prices.

In summary, the baseline VAR results are distorted by (expected) aggregate demand shocks due to the omission of contemporaneous, as well as lagged financial variables. These distortions can be effectively resolved by incorporating financial variables into the econometrician's information set, leading to more accurate and robust estimates of oil supply news effects.

4.5 Stability of the Results across Sample Periods

Another notable improvement is the greater stability across sample periods. Figure 5 shows the impulse responses when the VAR with financial variables is estimated over the periods 1974–2017 and 1982–2017, respectively. We consistently observe a sharp worldwide output contraction, and less persistent inflationary effects. The puzzling reversal in the monetary policy response observed in Figure 1 is also no longer present. Hence, by incorporating the financial variables, the instrumental variable becomes less susceptible to critiques related to sample period instability and output puzzles. For example, Degasperi (2021) interprets the output puzzle as evidence that the instrument also captures revisions in expectations about oil demand; but this puzzle disappears in the augmented VAR model.

The results also prove robust under various other checks. For example, the effects are similar when we exclude the first six years of the IV or estimate the augmented SVAR using a sample starting in 1988M4.[8] Furthermore, when additional variables are included in the VAR, such as the external bond premium or the principal components of McCracken and Ng (2016), the impulse responses remain remarkably consistent, indicating stability of the results. Finally, as shown in the appendix, the results remain stable when the sample period is extended to include data through 2024M8, covering the (post-)COVID era. It is worth noting that, while the inclusion of post-COVID data leads to a stronger impact of the shocks on consumer prices and there is a substantial contribution of oil supply news shocks to the increase in oil prices after COVID, the effects on the surge in inflation were relatively

[8] Kilian (2024) argues that the initial six years of daily OPEC surprises rely on unsuitable data, and shows that dropping these observations renders the VAR results sensitive to the estimation period.

modest. Given the extraordinary circumstances surrounding COVID, these results should obviously be interpreted with an appropriate degree of caution.

Figure 5: Stability of the SVAR-model augmented with financial variables

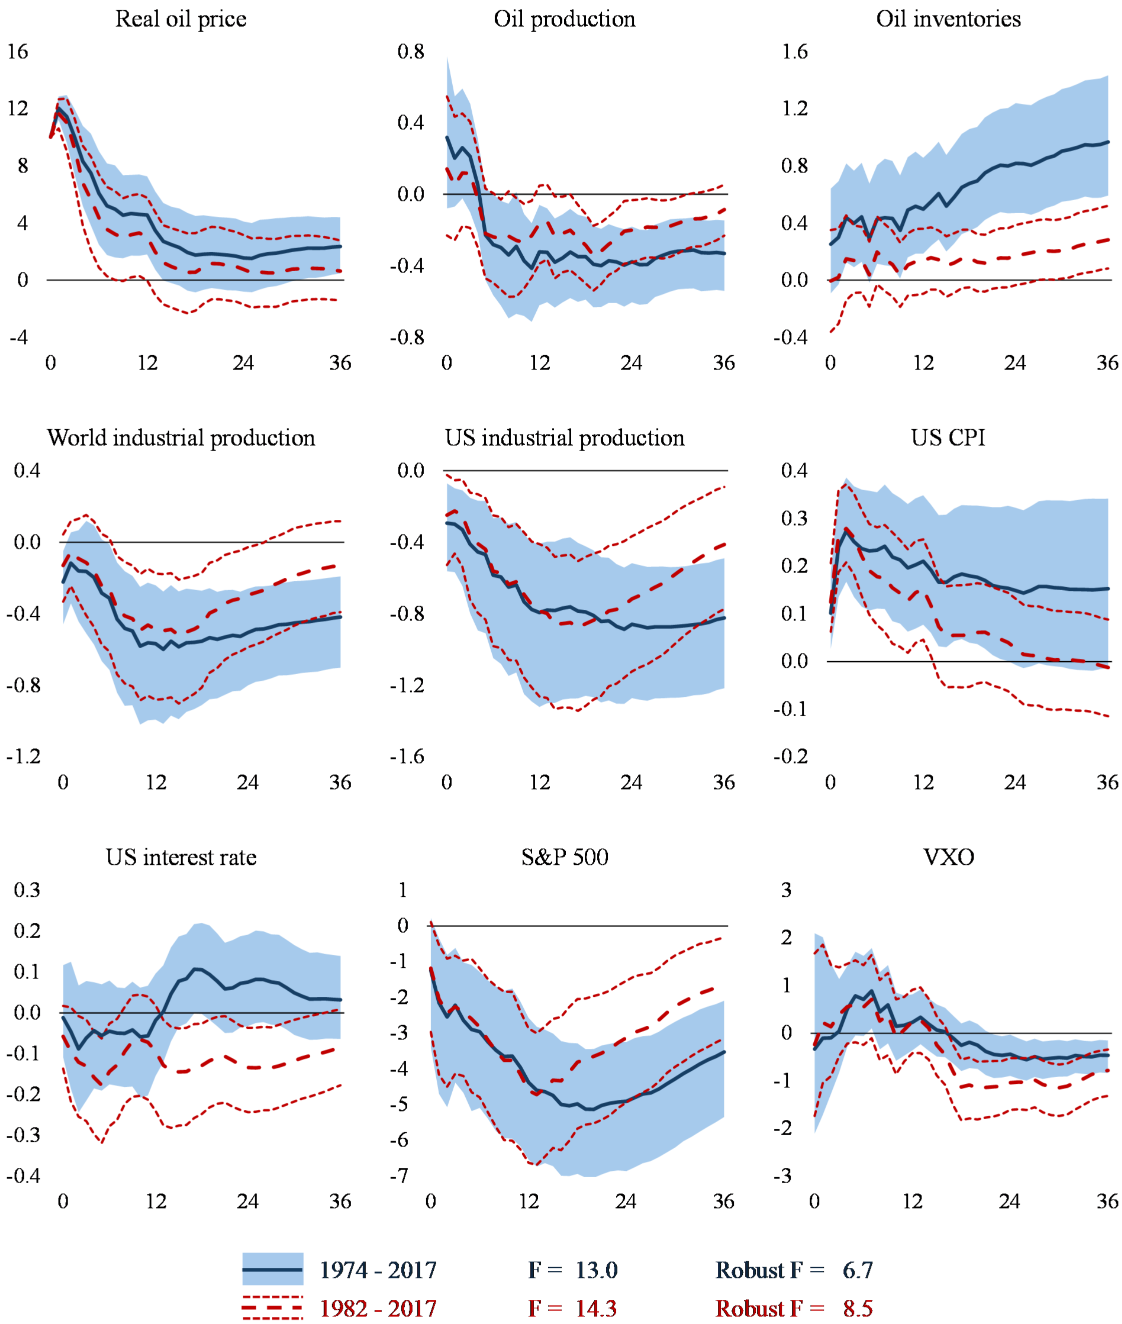

Note: Impulse responses to oil supply news shock that raises oil prices by 10% on impact, estimated over sample periods 1974-2017 and 1982-2017, respectively. 68% confidence intervals constructed using a moving block bootstrap.

5 Other Popular Oil-Market SVAR Models

Numerous papers have utilized SVAR models to estimate the dynamic effects of structural oil-market disturbances. Two prominent contributions are Kilian (2009) and Baumeister and Hamilton (2019). Kilian (2009) disentangles oil supply, aggregate demand and oil-specific demand shocks in a three-variable monthly VAR model that includes real oil prices, global oil production and an index of real economic activity. Baumeister and Hamilton (2019) estimate a four-variable VAR that also includes oil inventories, and identify oil supply, economic activity, oil consumption demand, and oil inventory demand shocks, respectively. Various studies employ alternative identification methods, but are based on the same information set (e.g., Peersman and Van Robays, 2009; Kilian and Murphy, 2014). Notably, the number of variables and structural shocks in most of these VARs are equal, which implies that the authors assume full invertibility; that is, all shocks can be obtained from linear combinations of the VAR residuals. Our findings suggest that these models may also suffer from informational deficiencies. In this section, we explore this for Baumeister and Hamilton (2019) and Kilian (2009).

Table 4 displays Granger-causality tests for the structural shocks estimated in both studies.[9] As discussed in section 3, external information is not supposed to Granger-cause the shocks. However, the table shows that the null hypothesis, which states that lagged external variables do not Granger-cause the structural shocks, is rejected for several variables not included in the VAR. Specifically, the common factors identified in McCracken and Ng (2016) jointly Granger-cause all the shocks in both studies. For all shocks in Kilian (2009), factor F7 is again the primary driver of this result. F7 also significantly predict the economic activity and oil consumption demand shocks identified in Baumeister and Hamilton (2019). Several financial indicators, such as the Global Financial Cycle and financial uncertainty support this observation. However, in contrast to the oil supply news shocks, other common factors also Granger-cause the structural shocks, particularly the aggregate demand shocks in Baumeister and Hamilton (2019) and Kilian's (2009) oil supply shocks. Notably, the OECD Global Composite Leading Indicator predicts oil consumption demand shocks, while past interest rates predict oil inventory demand shocks in Baumeister and Hamilton (2019), and the USD nominal effective exchange rate predicts aggregate demand shocks in Kilian (2009).

Overall, these findings imply that the shocks in these studies are not "structural", and influenced by endogenous responses to other (non-invertible) shocks. Again, it appears that the VARs lack information on market expectations about future demand and supply conditions, potentially causing a bias in the results due to omitted variables. It remains unclear whether these distortions are substantial.

[9] For consistency with the above analysis, all tests are based on the sample period 1974-2017. For Baumeister and Hamilton (2019), we have downloaded the shocks from Christiane Baumeister's website, while we have re-estimated the Kilian (2009) VAR model for this sample period.

Table 4: P-values Granger causality tests - Other oil-market VAR models

	Baumeister and Hamilton (2019)								Kilian (2009)					
	Oil supply		Econ. activity		Oil cons. demand		Oil inv. demand		Oil supply		Aggr. demand		Oil-spec demand	
	6*L*	12*L*	6*L*	12*L*	6*L*	12*L*	6*L*	12*L*	6*L*	12*L*	6*L*	12*L*	6*L*	12*L*
Common factors (all)	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	0.16	**0.00**	**0.01**	**0.00**
F1	0.43	0.57	**0.01**	**0.02**	0.10	0.36	0.11	0.33	0.15	0.08	0.67	0.42	0.90	0.44
F2	0.14	0.41	**0.00**	**0.01**	0.08	**0.01**	0.07	0.13	0.09	0.37	0.44	0.44	0.16	0.36
F3	**0.04**	0.15	**0.04**	**0.08**	0.41	0.28	0.08	**0.01**	**0.02**	**0.00**	0.68	0.81	0.16	0.43
F4	0.14	0.15	0.29	0.26	0.14	0.53	**0.01**	**0.02**	0.39	0.19	0.52	0.17	0.85	0.91
F5	0.15	**0.01**	**0.01**	**0.01**	0.45	0.22	0.09	**0.04**	**0.02**	**0.00**	0.67	0.49	0.52	0.33
F6	0.15	0.11	**0.01**	**0.03**	0.48	**0.04**	0.20	0.19	0.07	0.16	0.70	0.50	0.79	0.80
F7	0.06	0.08	**0.01**	**0.00**	**0.02**	**0.02**	0.20	0.35	**0.01**	**0.00**	**0.01**	0.07	**0.02**	**0.00**
F8	**0.01**	**0.02**	0.25	0.42	0.39	0.65	0.81	0.55	0.57	0.16	0.24	0.09	0.45	0.83
Global Financial Cycle	**0.01**	**0.03**	**0.01**	**0.00**	0.06	**0.01**	0.07	0.21	0.28	0.25	**0.00**	**0.00**	**0.03**	**0.04**
OECD Global CLI	0.47	0.15	**0.00**	**0.00**	**0.01**	**0.05**	0.56	0.19	0.19	0.32	0.06	0.06	0.06	0.17
S&P 500	**0.03**	0.06	**0.03**	0.09	0.29	0.09	**0.01**	0.10	**0.02**	**0.02**	0.16	0.49	0.37	0.10
MSCI World	0.35	0.13	**0.02**	**0.01**	0.22	0.40	**0.04**	0.18	0.44	0.53	**0.04**	0.09	0.13	0.40
VXO	0.60	0.28	**0.01**	**0.01**	0.48	0.29	0.11	0.17	**0.05**	0.21	0.20	0.09	0.20	**0.05**
Financial uncertainty	0.90	0.51	0.07	0.09	0.25	0.14	0.85	0.34	0.51	0.21	0.54	0.78	0.30	0.67
Macro uncertainty	0.92	0.49	**0.00**	**0.01**	**0.03**	**0.00**	0.34	**0.02**	**0.05**	**0.00**	**0.04**	**0.01**	0.18	0.19
Excess bond premium	**0.05**	0.08	**0.02**	**0.00**	0.06	0.20	0.57	0.65	0.42	0.34	0.33	0.41	**0.01**	**0.02**
BAA-AAA spread	**0.02**	0.30	**0.00**	**0.00**	0.20	0.45	0.09	0.06	0.11	0.29	0.57	0.57	0.44	0.87
US 1-year interest rate	0.22	**0.00**	0.18	0.09	0.06	0.16	**0.00**	**0.00**	**0.03**	**0.01**	0.50	**0.03**	0.45	0.84
USD nominal effective exchange rate	0.17	0.46	0.35	0.32	0.44	0.66	0.18	0.07	0.65	0.65	**0.00**	**0.00**	0.98	1.00

Note: P-values of (robust) F-statistic for the null hypothesis that the lagged variables do not Granger-cause the estimated structural shocks. Number of lags used in the test are 6 (6L) and 12 (12L). For SP500 and MSCI, we consider log differences. The common factors are obtained from McCracken and Ng (2016). Financial and macro uncertainty are collected from Ludvigson et al. (2021). The Global Financial Cycle are obtained from Miranda-Agrippino and Rey (2020), the S&P 500 from Robert Schiller's webpage, and the Excess bond premium from Gilchrist and Zakrajsek (2012). Numbers are in bold when p<0.05.

Unlike oil supply news shocks, the predictability of the shocks is not confined to financial variables, which requires other data to resolve the problem. Moreover, including additional variables in the information set necessitates assumptions for shock identification in these VARs. Addressing these challenges is beyond the scope of this paper.

6 Conclusions

Several studies use changes in oil future prices within a narrow window around OPEC quota announcements as an instrumental variable to estimate the macroeconomic effects of oil supply news shocks with SVAR-IV methods, which has greatly improved our understanding of the pass-through to the macroeconomy. However, in this paper we have shown that the reduced-form oil price innovations, the oil supply news shocks, and the external instrument in these studies are Granger-caused by financial variables. The predictability of the innovations and the structural shocks indicate informational deficiencies in the VAR model, and that the structural shocks are non-invertible. The predictability of the external instrument implies that even estimation methods not requiring invertibility, such as SVARs identified with an internal instrument or local projections using the VAR information set as control variables, are distorted.

To address this problem and quantify the distortions, we incorporated the omitted financial variables in the VAR specification. The results are markedly different: there is an immediate and more pronounced decline in economic activity, the impact on consumer prices is much lower and less persistent, while the shocks contribute far less to oil price variation. Additionally, the results are more stable over time, and puzzling responses disappear. At the same time, we find that oil supply news accounts for a substantial share of S&P 500 volatility. In summary, while the informational deficiencies and contamination of the instrument are quantitatively important, they can be relatively easily resolved by incorporating financial indicators into the econometrician's information set. The stability across sample periods and the disappearance of the puzzles also makes the instrumental variable less susceptible to critiques raised in the literature.

Finally, we demonstrated that the structural shocks identified in other prominent global oil-market SVAR models, particularly those of Baumeister and Hamilton (2019) and Kilian (2009), are also predictable by several variables not included in the VAR information set. The shocks in these models are, however, not solely predictable by financial variables. Addressing this issue and measuring the relevance of the omitted variables bias is left for future research.

References

Aastveit, Knut Are, Hilde C. Bjørnland, and Leif Anders Thorsrud (2015) "What Drives Oil Prices? Emerging Versus Developed Economies," *Journal of Applied Econometrics*, Vol. 30, pp. 1013–1028.

Bai, J. and Serena Ng (2002) "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, Vol. 70(1), pp. 191–221.

Baumeister, Christiane and James D. Hamilton (2019) "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks," *American Economic Review*, Vol. 109(5), pp. 1873–1910.

Beaudry, Paul and Franck Portier (2006) "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, Vol. 96(4), pp. 1293–1307.

Bruns, Martin (2021) "Proxy Vector Autoregressions in a Data-rich Environment," *Journal of Economic Dynamics and Control*, Vol. 123, p. 104046.

Bruns, Martin and Helmut Lütkepohl (2023) "Have the effects of shocks to oil price expectations changed?: Evidence from heteroskedastic proxy vector autoregressions," *Economics Letters*, Vol. 233, p. 111416.

Caldara, Dario and E. Herbst (2018) "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVAR," *American Economic Journal: Macroeconomics*, Vol. 11(1), pp. 157–192.

Caravello, Tomás and Pedro Martinez-Bruera (2024) "Disentangling sign and size non-linearities."

Castelnuovo, Efrem, Lorenzo Mori, and Gert Peersman (2024) "Commodity Price Shocks and Global Cycles: Monetary Policy Matters," Working Papers of Faculty of Economics and Business Administration, Ghent University, Belgium 24/1087.

Degasperi, Riccardo (2021) "Identification of Expectational Shocks in the Oil Market using OPEC Announcements," *Available at: https://www.riccardo-degasperi.com/research*.

Degasperi, Riccardo, Seokki Simon Hong, and Giovanni Ricco (2024) "The Global Transmission of U.S. Monetary Policy," Unpulbished Manuscript.

Forni, Mario, Alessandro Franconi, Luca Gambetti, and Luca Sala (2023) "Asymmetric Transmission of Oil Supply News."

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

Gagliardone, Luca and Mark Gertler (2023) "Oil Prices, Monetary Policy and Inflation Surges," NBER Working Paper 31263.

Gertler, Mark and Peter Karadi (2015) "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, Vol. 7(1), pp. 44–76.

Gilchrist, Simon and Egon Zakrajšek (2012) "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, Vol. 102(4), pp. 1692–1720.

Juvenal, Luciana and Ivan Petrella (2015) "Speculation in the Oil Market," *Journal of Applied Econometrics*, Vol. 30, pp. 621–649.

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

Kilian, Lutz (2009) "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, Vol. 99(3), pp. 1053–1069.

——— (2024) "How to Construct Monthly VAR ProxiesBased on Daily Surprises in Futures Markets," *Journal of Economic Dynamics and Control*, Vol. 168.

Kilian, Lutz and Daniel P. Murphy (2014) "The Role of Inventories and Speculative Trading in The Global Market for Crude Oil," *Journal of Applied Econometrics*, Vol. 29, pp. 454–478.

Kilian, Lutz and Cheolbeom Park (2009) "The impact of oil price shocks on the U.S. stock market," *International Economic Review*, Vol. 50, pp. 1267–1287.

Ludvigson, S., S. Ma, and S. Ng (2021) "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" *American Economic Journal: Macroeconomics*, Vol. 13(4), pp. 369–410.

McCracken, Michael W. and Serena Ng (2016) "FRED-MD: A Monthly Database for Macroeconomic Research," *Journal of Business and Economic Statistics*, Vol. 34(4), pp. 574–589.

Mertens, Karel and Morten Ravn (2019) "The dynamic effects of personal and corporate income tax changes in the United States: Reply to Jentsch and Lunsford," *American Economic Review*, Vol. forthcoming.

Mertens, Karel and Morten O. Ravn (2013) "The Dynamic Effecs of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, Vol. 103(4), pp. 1212–1247.

Miranda-Agrippino, S. and H. Rey (2020) "US Monetary Policy and the Global Financial Cycle," *Review of Economic Studies*, Vol. 87(6), pp. 2754–2776.

Miranda-Agrippino, Silvia and Giovanni Ricco (2021) "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, Vol. 13(3), pp. 74–107.

——— (2023) "Identification with External Instruments in Structural VARs," *Journal of Monetary Economics*, Vol. 135, pp. 1–19.

Montiel Olea, J, J. Stock, and M. Watson (2021) "Inference in SVARs Identified with an External Instruments," *Journal of Econometrics*, Vol. 225 (1), pp. 74–87.

Nguyen, Thuy Lan, Wataru Miyamoto, and Dmitriy Sergeyev (2024) "How Oil Shocks Propagate: Evidence on the Monetary Policy Channel," Federal Reserve Bank of San Francisco Working Papers 2024-07.

Patzelt, Paula and Ricardo Reis (2024) "Estimating the rise in expected inflation from higher energy prices," Discussion Papers 2411.

Peersman, Gert and Ine Van Robays (2009) "Oil and the Euro area economy," *Economic Policy*, Vol. 24, pp. 603–651.

Plagborg-Møller, Mikkel and Christian K. Wolf (2021) "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, Vol. 89(2), pp. 955–980.

Ramey, Valerie A. (2016) "Macroeconomic Shocks and Their Propagation," Vol. in: J.B. Taylor and H. Uhlig (Eds.): Handbook of Macroeconomics, Volume 2, Elsevier Science, pp. 71–162.

Stock, James H. and Mark W. Watson (2012) "Disentangling the Channels of the 2007-2009 Recession," *Brookings Papers on Economic Activity*, Vol. Spring, pp. 81–135.

——— (2018) "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, Vol. 128, pp. 917–948.

Appendix

A.1 Relationship between the VXO and the Instrumental Variable

The predictability of changes in oil future prices around OPEC quota announcements is surprising, given that future prices are expected to follow a random walk. The rejection of the null that financial variables do not Granger-cause the instrumental variable might be coincidental, possibly driven by in-sample noise or outlier observations, and could vanish asymptotically. If this is the case, incorporating variables correlated with the instrument into the VAR model, such as the VXO (in our case), may be problematic. These variables could "absorb" part of the actual effects of oil supply news.

In this appendix, we examine the relationship between the VXO and the IV in more detail. Specifically, if the correlation is accidental, it should be unstable within the sample. We explore this by regressing the IV on lagged changes in the VXO using the following specification:

$$IV_t = constant + \sum_{i=1}^{n} \beta_i \triangle VXO_{t-i} + \varepsilon_t \tag{6}$$

The results of this regression are presented in Table A1. Given that higher lag orders are statistically insignificant across all results we report, we set $n = 4$. The first four lags of changes in the VXO all exhibit a negative impact on the IV. The sum of the coefficients, as reported in the lower part of the table, is -0.158 and highly significant ($p < 0.01$). When the sample is restricted to months with OPEC announcements (i.e., when $IV_t \neq 0$), as shown in the second column, the relationship becomes even stronger, with a sum of -0.613. The fact that all four lags are significant and have the same sign provides an initial indication that this predictability is not accidental.

This conclusion is further supported by the observation that the negative relationship holds for both increases and decreases in the VXO, as reported in the third column, with symmetric coefficient sums (-0.156 for increases and -0.163 for decreases). Specifically, increases in the VXO tend to be followed by announcements that lower oil futures prices, while decreases in the VXO predict announcements that raise oil prices. The right panel of Table A1 demonstrates that the predictability and significant negative relationship also holds for changes in nominal oil prices in months with OPEC announcements. In contrast, the relationship between the VXO and oil prices is much weaker in months without announcements. This suggest that the VXO can predict OPEC announcements, rather than average oil price changes.

Finally, Figure A1, which plots the sum of the coefficients over a 15-year moving window, reveals that the relationship has remained remarkably stable throughout the sample period. In conclusion, the stability and symmetry of the instrument's predictability based on lagged changes in the VXO indicate that this relationship is neither coincidental nor the result of outlier observations. Identifying the reasons behind this relationship, however, is beyond the scope of this paper.

Table A1: Predictability of the Instrumental Variable and Oil Price Changes

	IV			Δp_{oil}	
	Full sample	Months with announcement	Full sample	Months with announcement	Months without announcement
$\triangle VXO_{t-1}$	-0.033* *(0.019)*	-0.202*** *(0.071)*		-0.539 *(0.338)*	-0.217 *(0.150)*
$\triangle VXO_{t-2}$	-0.046* *(0.024)*	-0.088** *(0.045)*		-0.527*** *(0.161)*	-0.168 *(0.157)*
$\triangle VXO_{t-3}$	-0.039** *(0.015)*	-0.174*** *(0.056)*		-0.581*** *(0.226)*	-0.014 *(0.086)*
$\triangle VXO_{t-4}$	-0.041** *(0.018)*	-0.148** *(0.059)*		-0.461* *(0.252)*	-0.170* *(0.088)*
$\sum_i \Delta VXO_{t-i}$	-0.158*** *(0.041)*	-0.613*** *(0.121)*		-2.108*** *(0.612)*	-0.569* *(0.293)*
$\sum_i \Delta VXO^{+}_{t-i}$			-0.156*** *(0.044)*		
$\sum_i \Delta VXO^{-}_{t-i}$			-0.163*** *(0.060)*		
$\overline{R}^2$	0.03	0.13	0.03	0.17	0.02
$\#obs$	417	117	417	117	300

Note: The table reports the sum of the coefficients of lagged changes in the VXO (4 lags). All estimations also include a constant (not reported). Dependent variables are the IV and changes in nominal oil prices, respectively. "+" and "-" are positive and negative changes in the VXO, respectively. Robust standard errors between parentheses. Significance as * p<0.10, ** p<0.05, *** p<0.01.

Figure A1: Predictability of the IV - 15-year moving window estimation

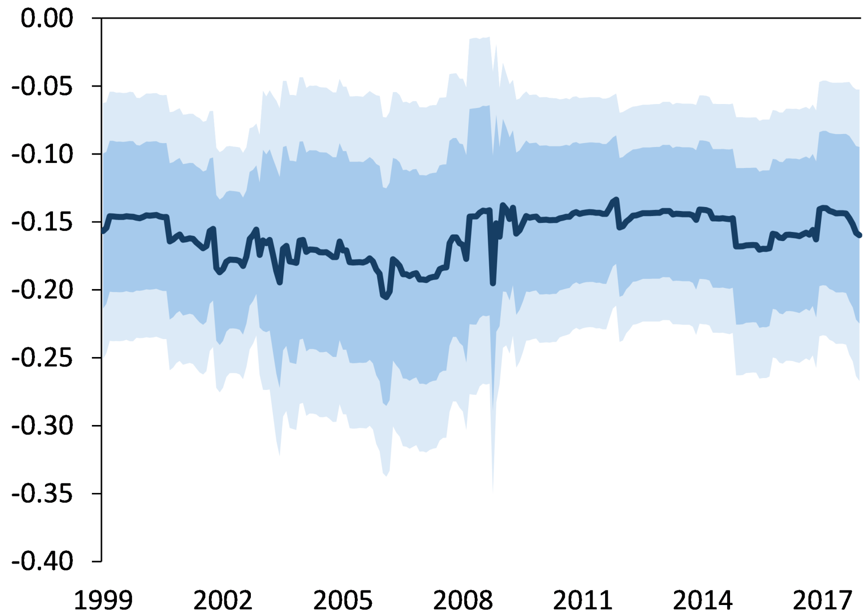

Note: The figure shows the sum of the coefficients of lagged changes in the VXO (4 lags), estimated over a moving window of 15 years (180 months). Dependent variable is the IV. The estimations also include a constant (not reported). 68% and 90% (robust) confidence intervals.

A.2 Robustness Checks

Figure A2, A3 and A4 present three robustness checks discussed in this paper. Figure A2 shows the instability in the original Känzig (2021) VAR-model (without the US interest rate) across sample periods. Figure A3 and A4 display the impulse responses of the baseline and augmented VAR models using alternative inference methods: the Bayesian approach from Miranda-Agrippino and Ricco (2021) and Miranda-Agrippino and Ricco (2023), and the weak-instrument robust inference method of Montiel Olea et al. (2021), respectively. Consistent with the discussion in Mertens and Ravn (2019), the confidence (credible) intervals generated by these alternative methods are somewhat narrower.

A.3 Historical Contribution of Oil Supply News Shocks

Figure A5 shows for the original VAR (top panel), as well as the augmented VAR model (middle panel), the cumulative historical contribution of oil supply news shocks to the real price of oil, together with the deviation of the actual real price of oil from its baseline evolution implied in the VAR. The bottom panel presents the cumulative historical contribution to the S&P 500 in the augmented VAR.

A.4 Results Covering the Post-COVID Era

Figure A6 demonstrates the stability of the results when the sample period is extended to 2024M8. To mitigate distortions caused by COVID-related outliers, five dummy variables are included in the VAR model for the period 2020M1-2020M5. Figure A7 presents the historical contribution of oil supply news shocks to the real price of oil, the S&P 500 and US CPI within this extended VAR model.

Figure A2: Instabilities in the original Känzig (2021) VAR-model

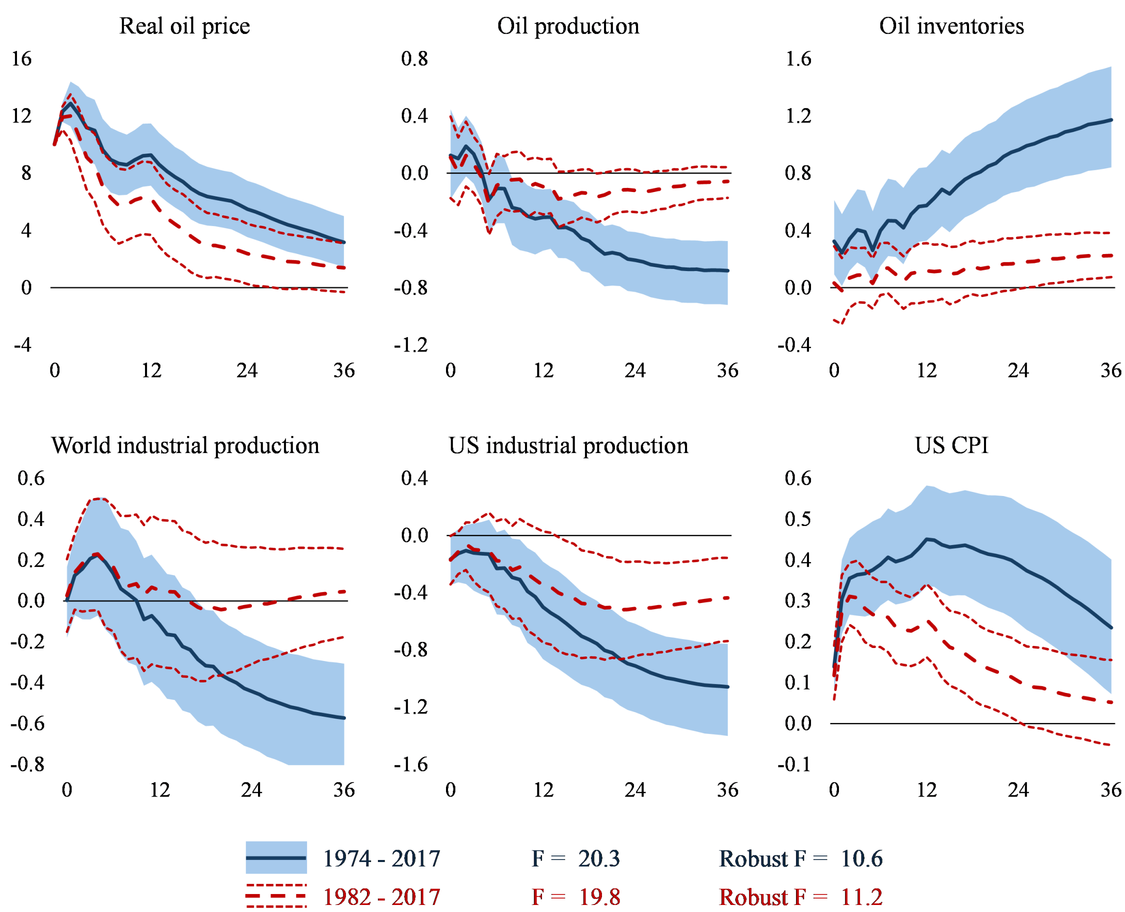

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact, for the sample periods 1974-2017 and 1982-2017, respectively. Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

Figure A3: Impact of oil supply news shocks - Bayesian SVAR

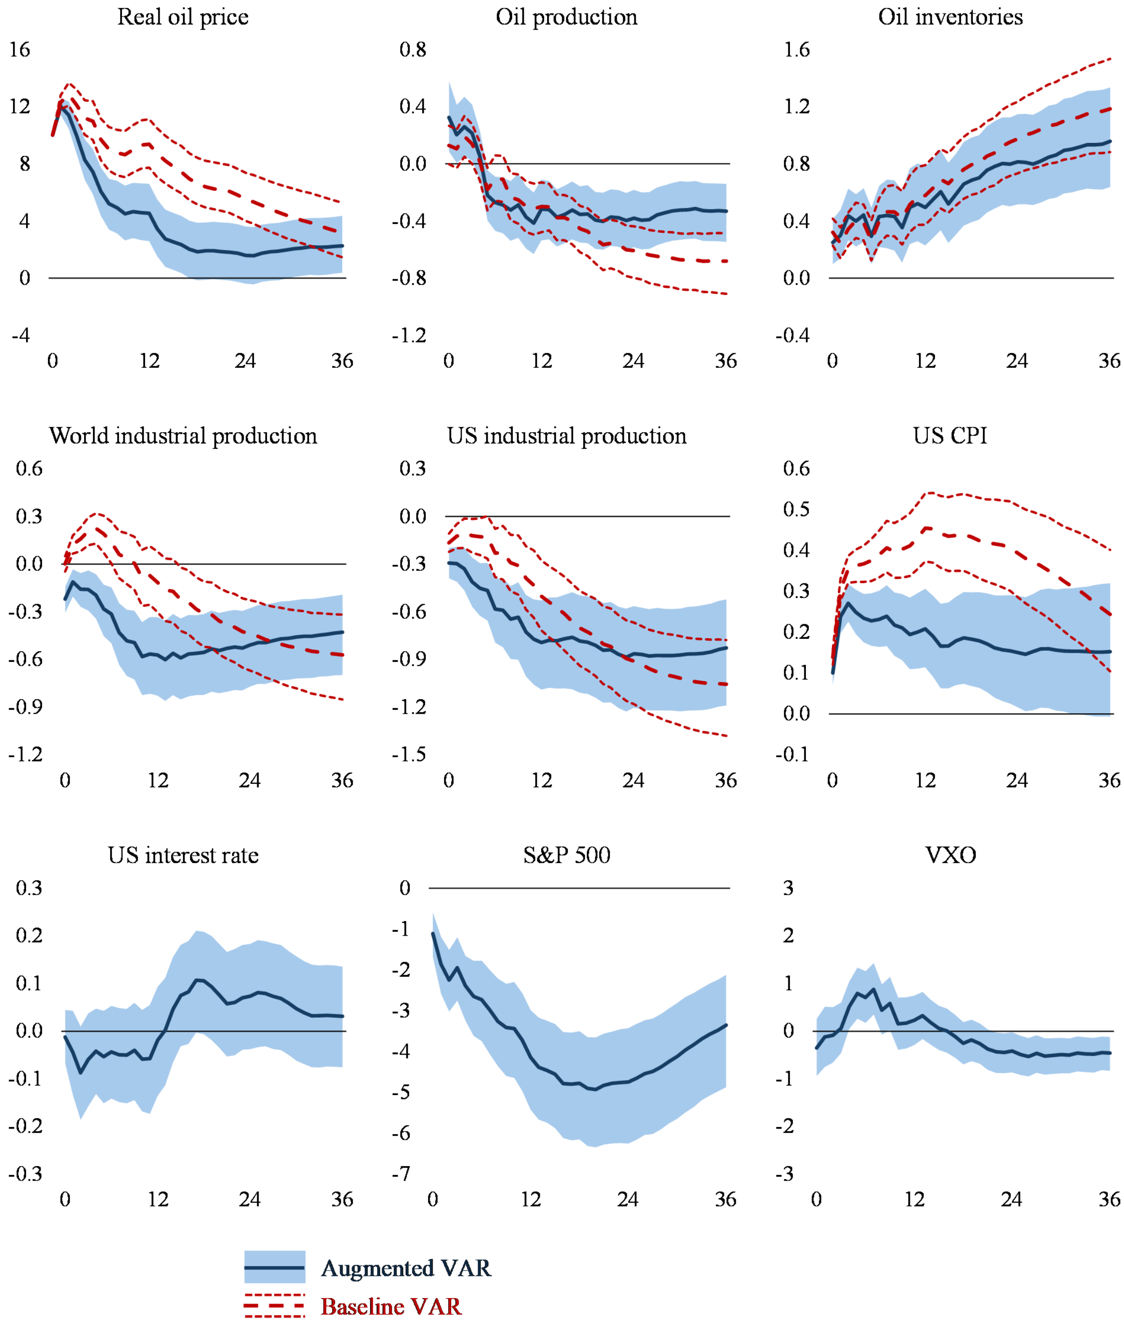

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Baseline VAR (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% credible intervals constructed using the Bayesian approach of Miranda-Agrippino and Ricco (2023), using a flat prior.

Figure A4: Impact of oil supply news shocks - Weak instrument robust inference

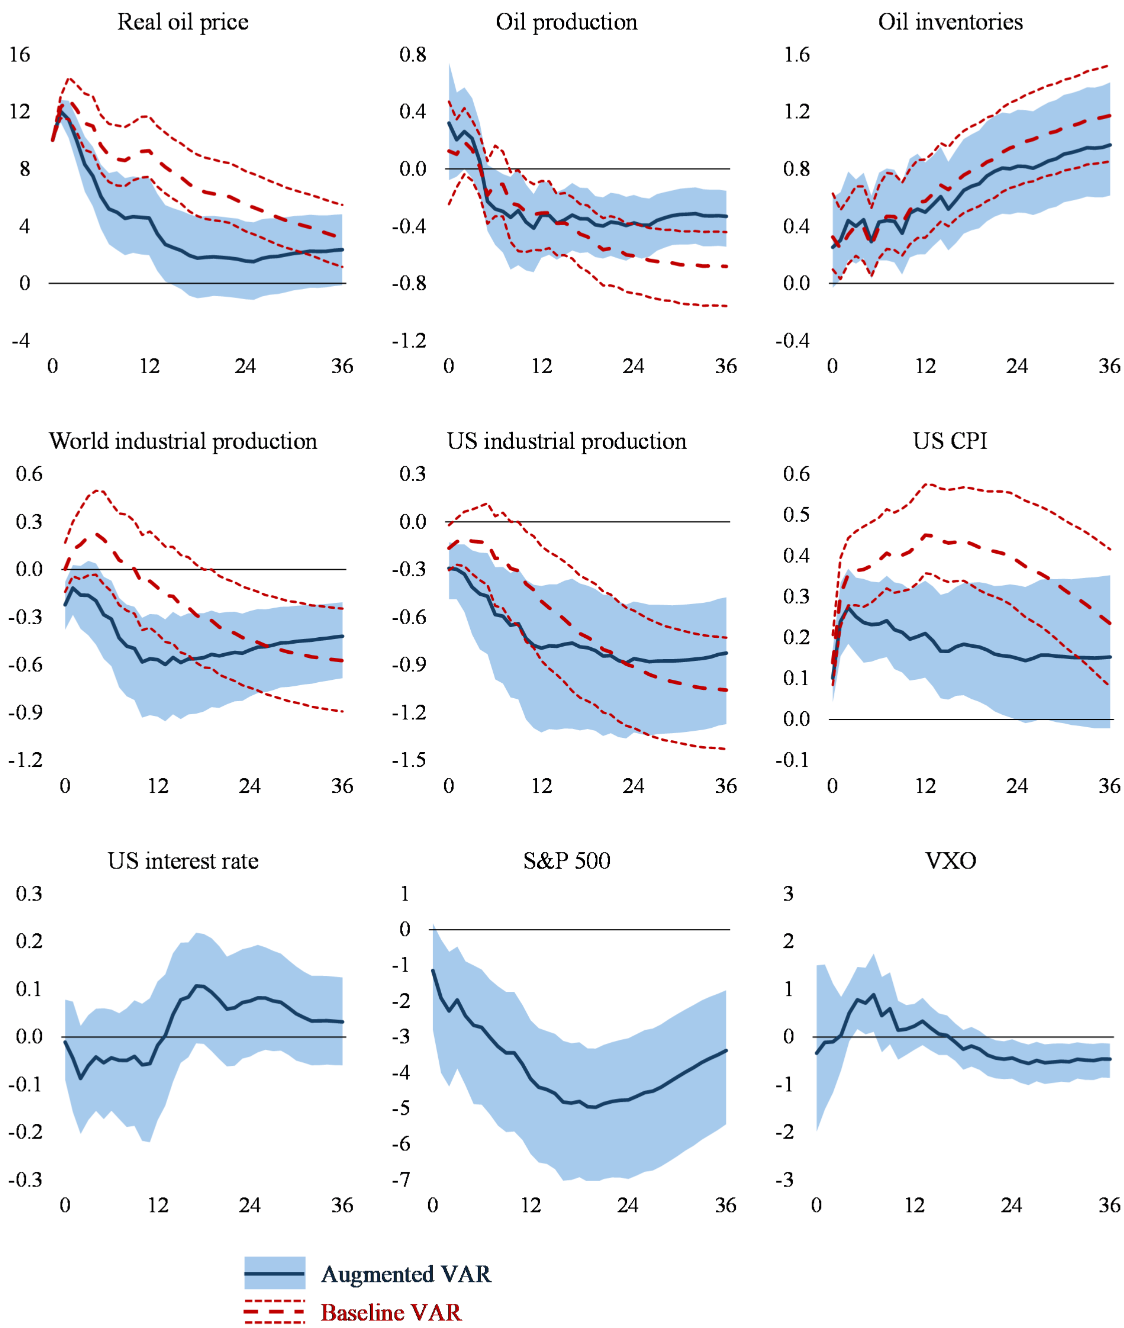

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Basline VAR (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% confidence intervals constructed using the weak instrument robust inference approach of Montiel Olea et al. (2021).

Figure A5: Historical contribution of oil supply news shocks

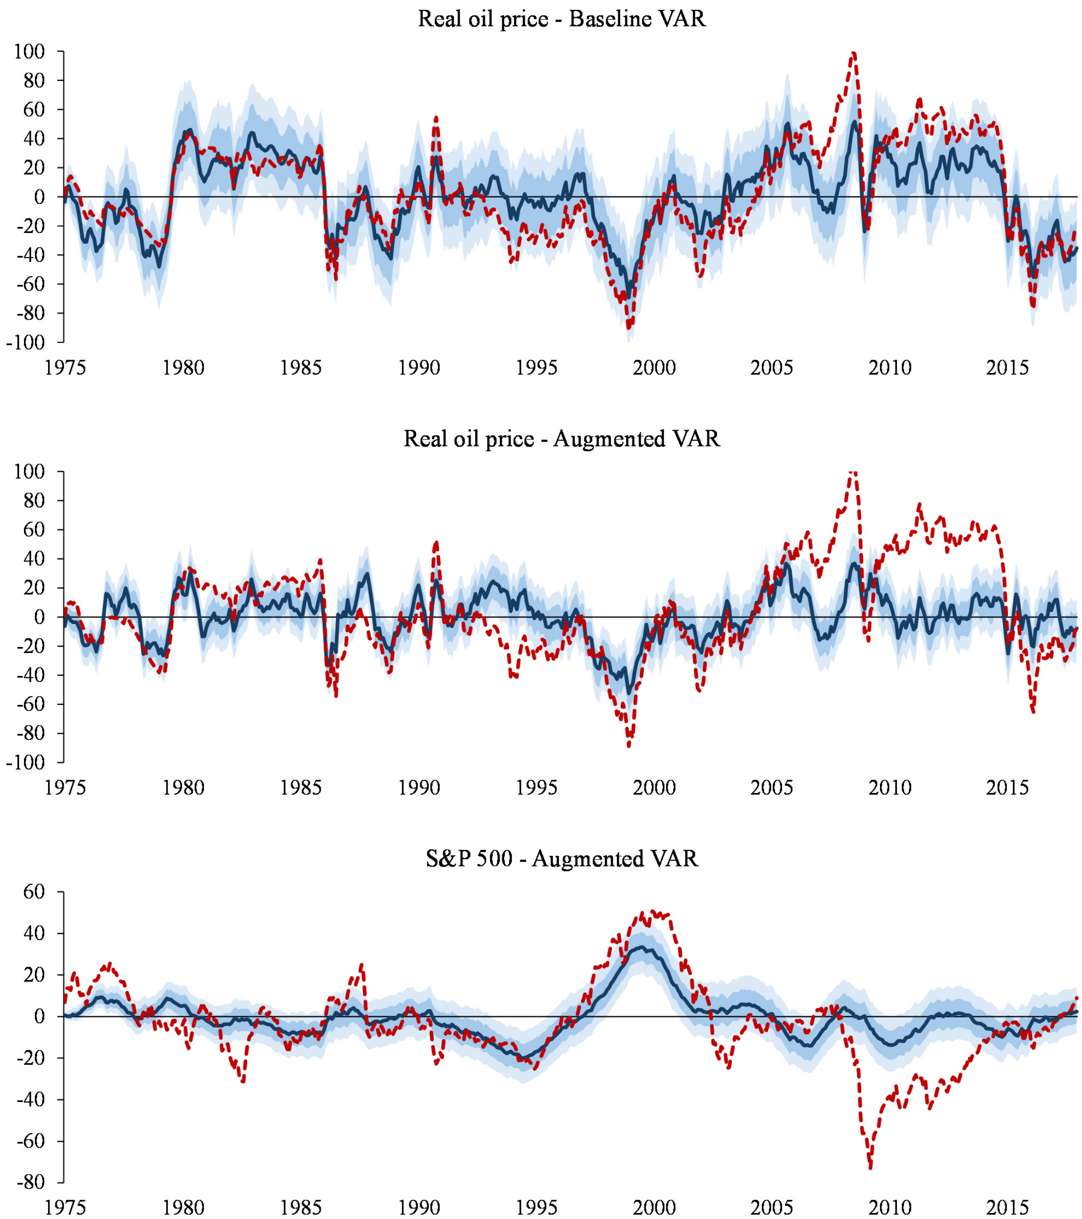

Note: Percentage points cumulative historical contribution of oil supply news shocks to the real price of oil and the S&P 500, together with the deviation of the actual real price of oil and S&P 500 from their baseline evolution implied in the VAR. The top panel is based on the baseline VAR, while the middle and bottom panel are based on the VAR augmented with financial variables. 68% and 90% confidence intervals constructed using a moving block bootstrap.

Figure A6: Stability of the augmented SVAR-model including the post-COVID era

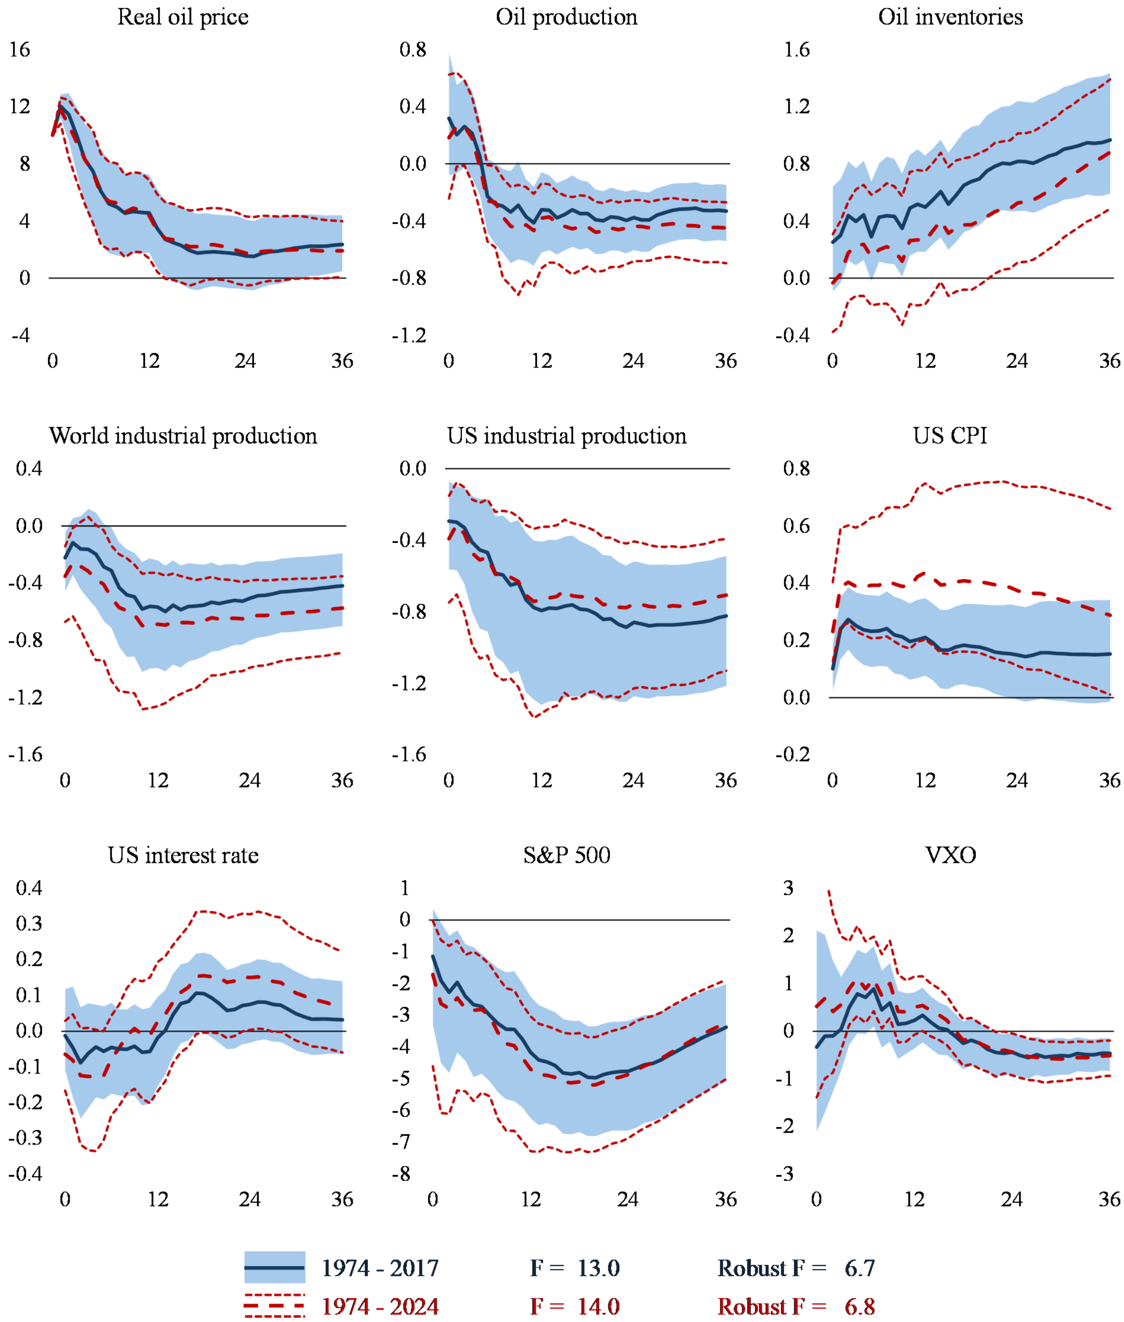

Note: Impulse responses to oil supply news shock that raises oil prices by 10% on impact, estimated over sample periods 1974-2017 and 1974-2024, respectively. 68% confidence intervals constructed using a moving block bootstrap.

Figure A7: Historical contribution of oil supply news shocks including the post-COVID era

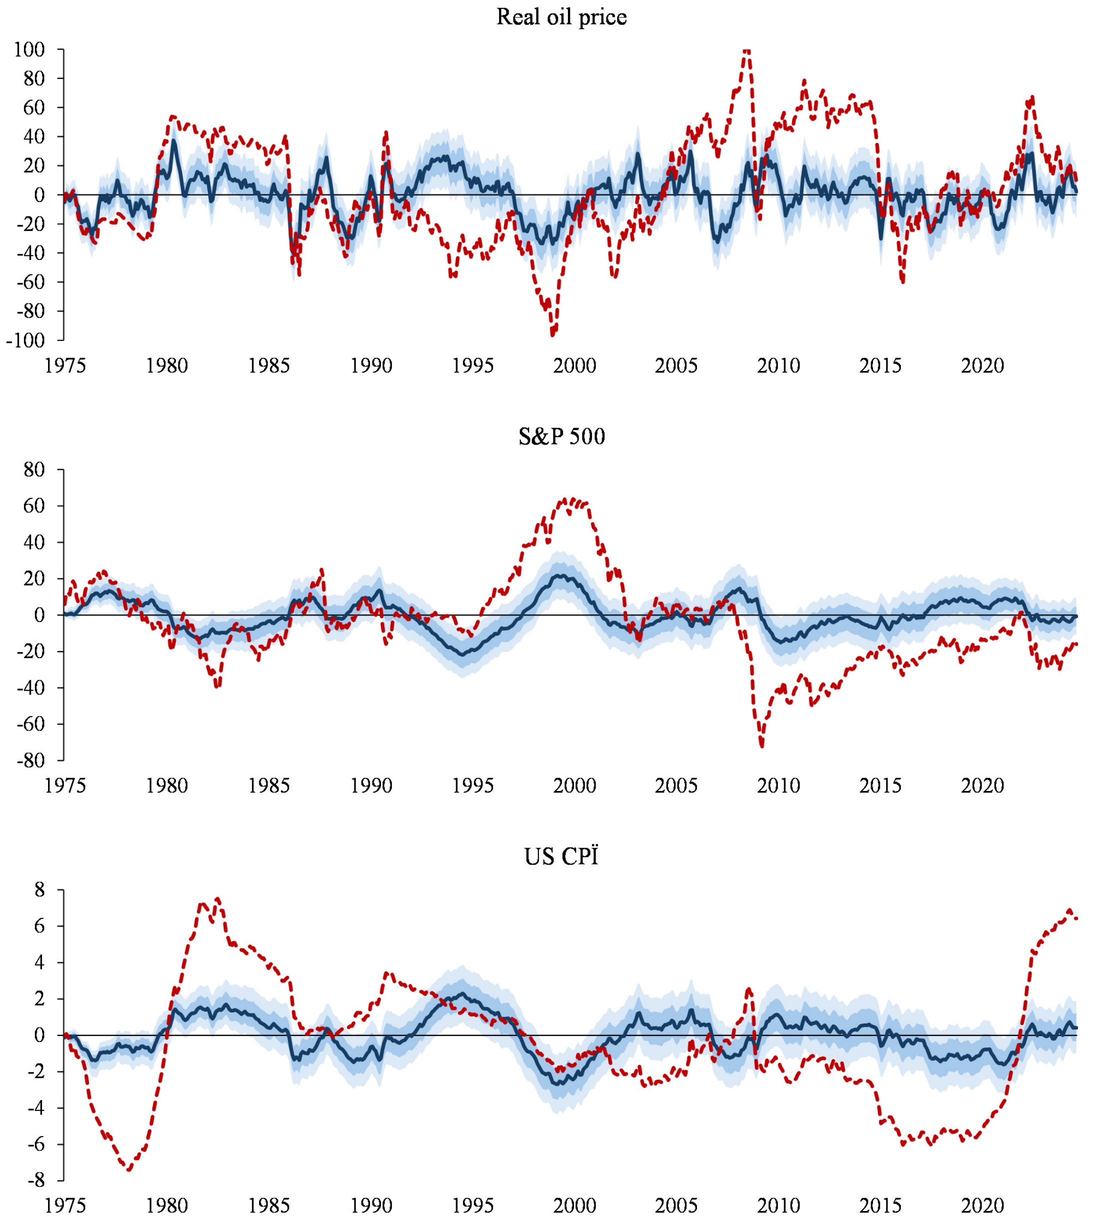

Note: Percentage points cumulative historical contribution of oil supply news shocks to the real price of oil, the S&P 500 and US CPI based on the augmented VAR model, together with the deviation of the variables from their baseline evolution implied in the VAR. 68% and 90% confidence intervals constructed using a moving block bootstrap.