BANCA D'ITALIA
EUROSISTEMA

Temi di discussione

(Working Papers)

Identification of expectational shocks in the oil market using OPEC announcements

by Riccardo Degasperi

February 2026

Number 1516

BANCA D'ITALIA
EUROSISTEMA

Temi di discussione

(Working Papers)

Identification of expectational shocks in the oil market using OPEC announcements

by Riccardo Degasperi

Number 1516 - February 2026

ISSN 2281-3950 (online)

Designed by the Printing and Publishing Division of Banca d'Italia

IDENTIFICATION OF EXPECTATIONAL SHOCKS
IN THE OIL MARKET USING OPEC ANNOUNCEMENTS

by Riccardo Degasperi*

Abstract

OPEC announcements reveal unanticipated information about future oil supply but may also lead imperfectly informed markets to revise their beliefs about demand conditions. As a result, surprises in oil futures prices around these announcements capture both a supply and a demand shock. Imposing an additional restriction on the sign of the comovement between oil futures surprises and stock price surprises results in clean instruments that separately identify these two components. A negative oil supply news shock has deep and long-lasting stagflationary effects, stronger than previously reported. This poses a challenge for monetary authorities and underscores the importance of accounting for information effects when identifying news shocks.

JEL Classification: C3, E3, Q4.

Keywords: oil supply news shocks, information frictions, information effects, OPEC announcements, high-frequency identification, external instruments, international transmission.

DOI: 10.32057/0.TD.2026.1516

---

* DG Economics, Statistics and Research; via Nazionale 91, 00184 Rome, Italy. Email: riccardo.degasperi@bancaditalia.it Web: www.riccardo-degasperi.com

1 Introduction[1]

How oil price fluctuations affect the global economy is an important question that has engaged economists and policymakers for at least four decades (Hamilton, 1983; Baumeister and Kilian, 2016). Quantifying these effects is challenging because oil prices and oil price expectations are endogenous to economic activity and respond to both supply and demand conditions. To analyse the effect of movements in oil prices, one has to identify the structural shocks that initiate them (Kilian, 2009). Oil supply news shocks are of particular interest since they have an immediate impact on current oil prices and can strongly affect global activity and financial markets. Oil prices are a forward-looking variable, which means that expectations of oil supply are as important as current oil supply in determining price movements (Kilian and Lee, 2014; Kilian and Murphy, 2014). In general, negative supply shocks tend to depress economic activity and increase prices, posing a challenge to monetary authorities seeking to stabilise both prices and output simultaneously (Darby, 1982; Bernanke et al., 1997). Unlike physical supply disruptions, oil supply news shocks affect prices through changes in market expectations about future oil production – often before any tangible change in production occurs. This paper improves on existing identification strategies to estimate the effects of oil supply news shocks on the global economy, showing that these effects are stronger and propagate more rapidly than previously reported (Känzig, 2021).

Announcements by the Organization of the Petroleum Exporting Countries (OPEC) regarding their output quota decisions can help address the identification problem as they generate variation in future oil supply conditions that is unexpected by market participants. OPEC's role as the key player in the oil market means that its decisions can

[1]I thank Giovanni Ricco and Ivan Petrella for their invaluable guidance. I also would like to thank Piergiorgio Alessandri, Christiane Baumeister, Dan Bernhardt, Apurav Yash Bhatiya, Hilde Bjørnland Christine Braun, Martin Bruns, Ambrogio Cesa-Bianchi, Natalie Chen, Anna Cieslak, Antonio Conti, Roger Farmer, James Fenske, Luca Gambetti, Andrea Gazzani, Georgios Georgiadis, Ana María Herrera Simon Hong, Matteo Iacoviello, Diego Känzig, Lutz Kilian, Robert Kollmann, Lorenzo Mori, Dennis Novy, Roberto Pancrazi, Gert Peersman, Claudia Rei, Federico Rossi, Raffaele Rossi, Moritz Schularick, Fabrizio Venditti, Giovanni Veronese, Marija Vukotić, and all the participants to the Warwick Macro & International Workshop, the 2022 T2M Conference, the 2022 RIEF Sciences Po, the 2022 BSE Summer Forum, the 2022 EEA-ESEM Conference, the 2023 IAAE Conference, the 2023 ICEEE Conference, the 2024 IWEEE Conference and seminar participants at the Federal Reserve Board, Bank of England, University of Oxford, University of Vienna, Bank of Italy, Université libre de Bruxelles, Collegio Carlo Alberto, Geneva Graduate Institute, University of Manchester, ECB IPA, and Corvinus University. Special thanks to Sergio Santoro for spotting an error in a previous draft. All errors are my own.

have a significant impact on the global price of oil and are closely scrutinised by financial markets (Lin and Tamvakis, 2010; Schmidbauer and Rösch, 2012; Brunetti et al., 2023).[2] Indeed, previous literature uses surprises in the price of oil futures computed on the last day of OPEC conferences, when the production quotas are publicly announced, to obtain a measure of exogenous changes in expected oil supply (Känzig, 2021). This measure can be used as an instrument to identify the effects of an oil supply news shock.

OPEC decisions on production quotas are themselves taken in response to changes in global economic conditions. However, under some assumptions, among which an assumption of perfect information, the surprises can be interpreted as revisions in oil prices that are a direct consequence of changes in markets' beliefs about oil supply. The conditions under which this is true are three. First, the event window within which the surprises are computed is narrow enough to minimize systematic contamination from other shocks. Second, OPEC and the markets share the same information and promptly price in all available information when making decisions. Third, risk premia do not change within the event window, so that any change in oil futures prices reflects a change in the expected future spot price.

I argue that the second condition – the assumption of perfect information – generally does not hold.[3] OPEC holds an informational advantage over the market regarding both current and future oil production.

Oil prices are often viewed as a timely barometer of global demand, however this interpretation only holds if supply is stable or fully observable (Hu and Xiong, 2013; Gazzani et al., 2024). When supply is fixed, fluctuations in price can be attributed entirely to changes in demand. In this sense, oil prices are a public signal of economic activity, with unobserved supply variation introducing noise into that signal. OPEC announcements, by explicitly or implicitly conveying information about supply, reduce this noise and might induce market participants to revise their expectations about underly-

[2]OPEC produces approximately 40% of the world's total crude oil and retains 80% of global crude oil reserves (OPEC, 2023). Despite recent efforts to speed up the shift to renewable sources, daily demand for crude oil worldwide has been increasing steadily from 85.3 million barrels per day (mmb/d) in 2006 to 100.1 mmb/d in 2022, and is expected to grow to 105.6 mmb/d by 2029 (IEA, 2021, 2024).

[3]The assumption on unchanged risk premia within the event window is also problematic. If the observed price change is not just a response to news about oil supply, but also reflects changes in investor attitudes purely toward risk and uncertainty, an oil supply news shock identified using the price surprises as proxy would be contaminated by an uncertainty shock, which potentially has different effects and a different transmission mechanism. I defer a discussion of this point to Section 4.2.

ing global demand conditions. Additionally they might reveal OPEC's own assessment of demand conditions. Therefore, the surprises in oil futures capture not only changes in expected oil supply, but also changes in expected demand, which invalidates the identification.[4]

Let me provide an example. Forecasters seeking to predict global aggregate demand by looking at the price of oil would do so conditioning their forecast to some estimate of oil supply. Assume that OPEC has been producing more than expected. Since an observed decline in the price of oil may be attributable either to a negative demand shock or to a positive supply shock, absent knowledge about supply developments, the forecasters may erroneously interpret the decline in oil prices as indicative of a weakening global demand. If many forecasters make this same mistake, that would put further downward pressure on the price of oil. The OPEC announcement, by revealing the true extent of oil production, induces the forecasters to revise what they previously had attributed to a decline in demand to an increase in supply. In other words, two things may happen to markets during the press release. First, they may revise their expectations about future oil supply: that is the oil supply news shock. Second, they may revise their estimate of aggregate demand based on oil prices, which they made conditional on the oil supply trajectory they had in mind before the announcement: that generates the information effect.[5,6]

I show this point formally within a model of information frictions in the oil market.[7] The takeaway is that any identification strategy based on surprises in oil futures alone would conflate the oil supply news shock with an aggregate demand shock, hence biasing the estimates of its dynamic effects. Indeed, the dynamic responses obtained with this identification strategy present various puzzles, as the underlying shocks have opposite

[4]In the paper, I refer interchangeably to this information effect – when it occurs in isolation – as 'demand shock' and 'information shock'. Since the information revealed by OPEC induces a revision in demand, the information shock is effectively a demand shock and indeed it propagates as such.

[5]In Section C.1 I provide evidence that, in a narrow event window around several OPEC announcements, professional forecasters revise their forecasts of GDP growth and inflation in the same direction as the surprise in the price of oil futures, consistently with the information channel described. This is true for at least 12 out of 48 OPEC announcements occurred between June 2011 and June 2023.

[6]It is important to notice that this is conceptually different from the information channel of monetary policy, where central banks, via the rate setting decision, implicitly signal to the markets their private view of the economy. In that case, markets receive a direct public signal of economic activity. In this case, the OPEC announcement reduces the noise in the signal about economic activity that is the global price of oil.

[7]I provide a simplified graphical intuition of the model's mechanism in Section A.

effects on macroeconomic aggregates.

Based on the model predictions, I propose an empirical method to disentangle the demand and supply components in the surprises, providing an exogenous measure of shifts in oil supply expectations that can be used to identify the shock of interest. This method exploits the high-frequency co-movement of oil futures prices and stock prices in a narrow window around OPEC announcements. The co-movement is informative because the demand shock (i.e. the information shock) moves both oil futures and stock prices in the same direction, whereas an oil supply news shock moves them in opposite directions. Imposing a restriction on the sign of this co-movement enables a clean identification of the shocks. By adding this restriction, I obtain two robust high-frequency instruments: one to identify the oil supply news shock and one for the information shock. In practice, I divide the surprises in oil futures into two groups based on their co-movement with stock prices. The instrument for oil supply news shocks consists only of those surprises that induce a negative co-movement with stock prices, while the instrument for the information shock uses those that induce a positive co-movement.

One might object that, since the argument hinges on the idea that markets cannot perfectly distinguish whether oil price movements are driven by demand or supply, it would be unwarranted to use high-frequency market reactions to identify these shocks. However, this is not a contradiction. While, in general, markets face uncertainty regarding the underlying drivers of oil price fluctuations, there are specific instances – such as OPEC announcements – where the source of the shock is effectively disclosed.

Importantly, for this result to hold true, OPEC does not need to have better knowledge of aggregate demand than the market. This assumption would be hard to justify, as market analysts and OPEC decision-makers are likely using the same sources to inform their forecasts of global oil and aggregate demand.[8] Conversely, it is reasonable to assume that OPEC knows more than the markets about oil production because, ulti-

[8]However, in Section 5.1, I provide some evidence that OPEC might indeed have superior knowledge of oil demand conditions.

mately, OPEC decides how much crude oil to produce.[9] For the markets, it is difficult to forecast OPEC production decisions for at least two reasons. First, production quotas are the outcome of closed-door negotiations among OPEC members. As with any cartel, OPEC decisions are not only a function of external demand-side developments, but also depend on the negotiating power of its members and their domestic goals. Second, OPEC members often do not respect the quotas, and the extent to which they plan to abide by them can partially be gauged by the comments and other forms of 'soft information' released during the conferences.

Figure 1 offers a preview of the main results. It shows the impulse responses of world industrial production to an oil supply news shock identified using the full set of daily surprises in the price of oil futures on OPEC announcement days as an instrument (left panel), and the oil supply news and demand (or information) shocks identified using the robust instruments developed in this paper (right panel). The responses are obtained from a six-variable VAR of the global oil market, similar to the one in Känzig (2021), estimated on the sample 1982:7–2023:5. In both cases, shocks are normalised to induce a 10% increase in the real oil price.

Economic theory suggests that an oil supply news shock that increases oil prices would have a contractionary effect on industrial production. However, when using the surprises as an instrument, world industrial production appears to increase for about 6 months after the shock hits and never contracts. This is a puzzling result and exemplifies the issues that may emerge when basing identification on the surprises in oil futures alone. On the other hand, the robust instruments based on the high-frequency

[9]To control oil supply, OPEC relies on an internal mechanism through which member countries self-report their production levels and compliance with agreed quotas. These figures, which are released with a delay in OPEC's Monthly Oil Market Reports (MOMRs), constitute part of OPEC's informational advantage over the average market participant. However, self-reported data can be unreliable, as member states may have political or economic incentives to misreport or delay accurate disclosure – especially during periods of quota negotiations or when facing sanctions. To monitor compliance with production targets, OPEC publishes both the self-reported figures (referred to as "direct communication") and estimates from secondary sources. Until recently, these secondary sources included the U.S. Energy Information Administration (EIA), the International Energy Agency (IEA), and private firms such as Rystad Energy. In 2022 OPEC stopped relying on IEA estimates and, as of early 2025, OPEC replaced the EIA and Rystad with private firms Kpler, OilX, and ESAI. These commercial sources are not exclusive to OPEC. Ultimately, OPEC's informational advantage stems from its role as the decision-maker. Because production adjustments are negotiated privately and not immediately disclosed, OPEC members necessarily hold superior knowledge about both current and future oil supply dynamics.

Figure 1: COMPARISON BETWEEN ALTERNATIVE IDENTIFICATION STRATEGIES

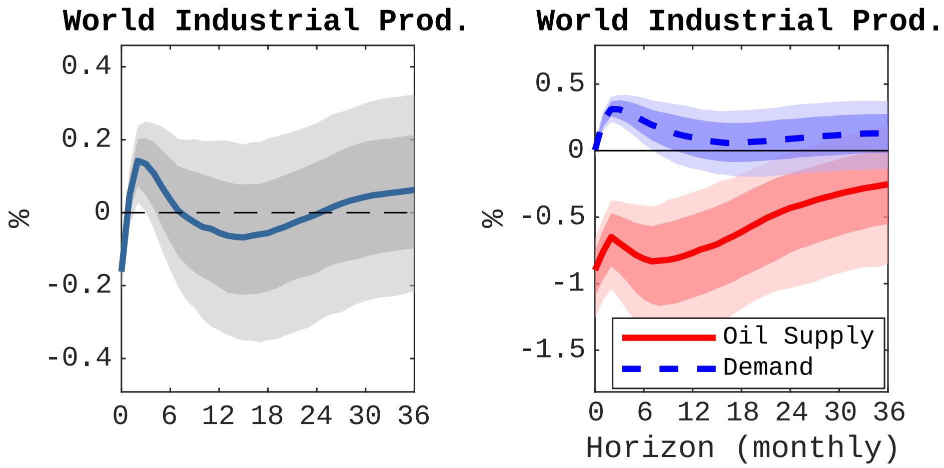

Note: Left panel – Impulse responses of world industrial production to an oil supply news shock identified using daily surprises in oil futures on OPEC conference days. Right panel – Solid red: responses to an oil supply news shock identified using the robust proxy proposed in this paper. Dashed blue: responses to an information shock identified using the residual variation in the surprises. All shocks are normalised to induce a 10% increase in the real price of WTI crude oil. The responses are obtained from the six-variable Bayesian VAR(12) described in Section 3.5. The responses for the variables that are not shown can be seen in Figures 4 and 5. Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1982:7–2023:5. Both proxies are available for the period 1983:7–2023:6. Results are similar when accounting for the pandemic period.

co-movement of oil futures and stock prices yield predictions that are consistent with theory (right panel). An oil supply news shock that raises the price of oil has an unambiguous contractionary effect on world industrial production. World industrial production contracts by almost 1% following an oil supply news shock normalised to increase real oil price by 10%. Conversely, the information shock produces an expansion in economic activity.

The main result of the paper is that information effects, which have been shown to be pervasive in many contexts, are also a prominent feature of the oil market. Failing to account for them can lead to biased estimates of the effects of the shock of interest. When accounting for these information frictions, the effects of oil supply news shocks are deeper and more immediate than previously reported. This represents a challenge for monetary and fiscal authorities worldwide. Lastly, as I will show in Section 4.4, I find novel evidence that the shock also propagates via the credit channel. Following an oil supply news shock that raises the price of oil, stock markets contract, financial uncertainty increases, and credit spreads widen. The credit channel acts as an amplification mechanism for the effects of the shock on the real economy (see also Gelain and Lorusso, 2025, for a similar result).

Related Literature. Recently, daily surprises in oil futures on OPEC conference days have been proposed as a measure of shocks to oil supply expectations (Känzig, 2021).[10] This is an important contribution because it shows that the expectational component of oil price movements can have macroeconomic effects as significant as actual shocks to oil supply.[11] I show that surprises in oil futures cannot be used "as is" to identify expectational shocks and propose a way to separate the supply and demand components in the surprises. In doing so, I draw a parallel with the literature that disentangles monetary policy and information shocks in the high-frequency surprises in interest rates that follow monetary policy announcements. These surprises have been used to identify monetary policy shocks but also capture an important informational component that acts as a confounding factor (Melosi, 2017; Nakamura and Steinsson, 2018). Two distinct methodologies can be used to disentangle the policy shock from the informational component. The first exploits the high-frequency co-movement between surprises in interest rates and asset prices (Jarociński and Karadi, 2020; Cieslak and Schrimpf, 2019; Cieslak and Pang, 2020). The second one separates the informational component from the policy shock by directly controlling for the central bank's information set (Miranda-Agrippino and Ricco, 2021; Romer and Romer, 2000). Here, I show how to adjust these methodologies to the study of oil markets and find that they are equally effective at disentangling the demand and supply components in the surprises. Importantly, my contribution does not simply translate results from the monetary policy literature to the oil literature but provides an important new insight into both by showing that in-

[10] Adopting a different methodology, Wu and Cavallo (2012) use OPEC announcements to obtain an exogenous measure of oil price shocks. They control for the predictable component of oil price revisions by regressing them on the spreads between oil spot and futures prices at different horizons on the day before the announcement. However, this does not solve the limitation pointed out in this paper.

[11] Baumeister (2023), using a measure of shocks to oil price expectations obtained by orthogonalising the difference between realized oil prices and oil price expectations – derived from futures prices using the model from Hamilton and Wu (2014) – with respect to the four fundamental oil market shocks in Baumeister and Hamilton (2019), shows that these shocks account for a non-trivial fraction of market-based oil price surprises.

formation effects may arise even when the agent originating the public signal does not have an informational advantage about the fundamental.

My paper fits into the large and well-established literature that studies the macroeconomic effects of oil price shocks.[12] One key takeaway from this literature is that, to determine the effects of fluctuations in oil prices, the underlying sources of the price movements must be identified (Kilian, 2009). A heated point of debate concerns the measurement of the price elasticity of oil supply, which determines the relative importance of oil supply shocks in driving fluctuations in oil prices compared to other shocks (Baumeister and Peersman, 2013; Kilian and Murphy, 2014; Bjørnland et al., 2018; Baumeister and Hamilton, 2019; Caldara et al., 2019). Some relevant contributions focus on specific determinants of oil price fluctuations: Juvenal and Petrella (2015) and Basak and Pavlova (2016) focus on the role of financial speculation; Anzuini et al. (2015), also exploiting OPEC announcements, study the macroeconomic effects of precautionary demand; Gambetti and Moretti (2017) analyse the role of noise shocks; Gazzani et al. (2024) propose an identification strategy for the sources of oil price fluctuations based on daily and real-time data. My focus is on oil price movements caused by revisions in market expectations, and I show that they can have deep and persistent effects on global economic activity.

Some very recent papers are particularly connected to my contribution. Kilian (2024) discusses issues related to the low liquidity of WTI futures contracts at long maturities during the 1980s and proposes a procedure to address temporal aggregation bias in the construction of proxies based on high-frequency surprises.[13] Mori and Peersman (2024) show that the 6-variable VAR used in Känzig (2021) suffers from informational insufficiency and that adding financial variables to the specification, as I do in Section 4.4, already helps alleviate the output puzzles I describe. Conversely, Forni et al. (2025), which provide evidence of asymmetric transmission of oil supply news shocks, find that Känzig (2021)'s 6-variable VAR is invertible. Bruns and Lütkepohl (2023) find evidence that the transmission of shocks to oil price expectations has changed around the 1990 Gulf War, which suggests that, on specific samples, the additional sign restrictions I employ in this study to disentangle the shocks might not be enough to correctly

[12] Kilian and Zhou (2023) offers a critical review of recent contributions to this literature.

[13] I discuss both points in Sections C.7 and C.8.

characterise the dynamics of the shocks. Herrera and Rangaraju (2025) find evidence of time-variation in the response of inflation expectations to oil news shocks. Alsalman et al. (2023), consistently with the results in this paper, show that the surprises of Känzig (2021), during Covid, are contaminated by lags of oil market variables. Moussa and Thomas (2023) use a Max-Share approach to identify oil supply news shocks, obtaining results qualitatively and quantitatively similar to those obtained here with a proxy-SVAR approach.

Finally, my paper relates to the event-study literature on the effect of OPEC announcements on oil prices (Draper, 1984; Demirer and Kutan, 2010; Lin and Tamvakis, 2010; Loutia et al., 2016) and on oil price volatility (Schmidbauer and Rösch, 2012).

The paper is structured as follows. In Section 2 I present a model of information frictions in the oil market and show that oil price revisions following an OPEC announcement capture changes in markets' expectations about both oil supply and aggregate demand. The model provides the rationale on which the identification strategy proposed in this paper is based. In Section 3 I describe the methodology adopted to disentangle the oil supply news and information shocks. Two different approaches are carried out and shown to deliver similar results. I also describe the data, the methodology to identify structural VARs by external instruments, and other details about the empirical exercises. In Section 4 I detail the main results of the paper. First, results obtained using the daily surprises in oil futures on OPEC announcement days to identify the oil supply news shock are sample-dependent and give rise to output puzzles. Second, the robust instrument developed to identify the oil supply news shock solves these puzzles. Third, an analysis of the transmission of the shocks to the world economy and to advanced and emerging economies shows that the shocks have powerful global effects. In Section 5 I provide robustness exercises. Finally, in Section 6 I draw some conclusions.

2 Information frictions in the oil market

The model presented in this section shows that revisions in the price of oil that follow OPEC announcements contain both supply and demand components and therefore cannot be used as such to identify the effects of an oil supply news shock on the economy. In the model, agents use the price of oil as a public signal of global economic conditions. Besides the price of oil, they only observe a noisy private signal of economic conditions and only know the distribution of oil supply shocks.[14] Oil supply shocks can be thought as the noise in the public signal. Since agents cannot be certain whether movements in the price of oil are driven by supply or demand shocks, they will attribute part of the decrease in the price of oil that follows a favourable supply shock to a deterioration of economic conditions, further depressing the oil price below what it would be under full information. Indeed, information frictions amplify the negative effect of the supply shock on prices.

An OPEC announcement, by revealing the oil supply shock, induces an upward revision in agents' expectations about economic conditions and, consequently, in the price of oil. This generates a positive correlation between surprises in the oil price and revisions of expectations about economic conditions, which renders the surprises endogenous. An upward price revision can be due both to the revelation of economic conditions and to a negative oil supply shock.[15]

The model also provides an additional restriction that allows me to disentangle the two underlying shocks. After an announcement, the revision in expectations about economic conditions moves in the same direction as the price revision following a demand shock, but the two revisions move in opposite directions following an oil supply shock. This suggests an identification strategy based on the sign of the co-movement between oil price revisions and revisions in expectations about economic conditions, proxied in the empirical exercise by a stock price index or a daily indicator of economic activity.[16] A positive co-movement identifies the demand shock, while a negative co-movement

[14] In the model, these are shocks to the actual supply of oil.

[15] Section A in the Appendix provides a simplified graphical illustration of the model's workings.

[16] To satisfy the identification assumption, surprises in the indicator of aggregate production must be measured over the same window as the surprises in the price of oil futures. For a detailed discussion, see Section 3.

identifies the oil supply shock.

2.1 Model setting

The argument builds on the model of information frictions in commodity markets of Sockin and Xiong (2015), which I briefly describe.[17,18] There is a continuum of islands that produce an island-specific good (Lucas, 1972). The domestic good can be consumed or traded with the good produced by another island. Production uses oil as input.[19] Following Angeletos and La'O (2013), the model has two periods. In period $t = 1$ the representative firms on each island make their production decisions and trade oil with OPEC to meet their production needs. At $t = 2$ all islands are randomly paired, trade their goods with each other, and consume.

OPEC is an independent entity that produces and supplies oil as a function of its price, but can deviate from this simple rule by a supply shifter ξ. OPEC's supply schedule is

$$\log X_S = k \log P_X + \xi, \qquad \xi \sim \mathcal{N}\left(\bar{\xi}, \tau_\xi^{-1}\right), \tag{1}$$

where P_X is the price of oil and k is a positive scalar.[20] Importantly, ξ is part of OPEC's information set, while the final-good producers only observe the parameters of its distribution.

The representative agent on island i wants to consume both foreign and domestic goods. Her utility function is given by

$$U\left(C_i, C_i^*\right) = \left(\frac{C_i}{1-\eta}\right)^{1-\eta} \left(\frac{C_i^*}{\eta}\right)^{\eta}, \tag{2}$$

where $\eta \in [0, 1]$, and C_i and C_i^* represent the consumption of domestic and foreign

[17]The focus of Sockin and Xiong (2015) is that, due to the informational role of commodity prices, supply shocks have an indeterminate effect on producers' demand, such that when the informational channel is strong enough the demand for commodities might *increase* with the price. In contrast, our focus is to show that announcements which improve the informativeness of commodity prices, even when the informational channel is not particularly strong, lead to revisions in price and demand of the same sign.

[18]I use the same notation. For more details and the complete derivations of the model, the reader is referred to the original paper.

[19]The number of islands is normalised to 1.

[20]OPEC's supply rule can be easily micro-founded so that OPEC maximizes its profit by assuming convex labour costs (see Sockin and Xiong, 2015).

goods respectively. From the perspective of the representative agent of island i, the goods produced by all other islands are perfect substitutes with each other, but they are complements with the domestically produced good. This gives agents a motive to trade at $t = 2$ and producers a reason to take into account the production decisions of their trading partner when making their own production plans in $t = 1$.

Representative firms on all islands produce according to a decreasing-returns-to-scale production function that uses oil as input:

$$Y_i = AX_i^{\phi}, \tag{3}$$

where $\phi \in (0, 1]$, A is the productivity that is common to all islands, and X_i is the oil input.[21] Productivity A is a random variable with a log-normal distribution,

$$\log A \sim \mathcal{N}\left(\bar{a}, \tau_A^{-1}\right). \tag{4}$$

Given the complementarity in consumption, the demand for the domestic good depends on the production of the island's trading partner. Since A determines the amount produced by both the domestic firm and its trading partner, it can be thought of as the strength of the overall economy. Ultimately, A determines demand. Importantly, producers observe A only after production has taken place at $t = 2$ and only know the parameters of its distribution. Producers use the price of oil, which is determined in equilibrium in $t = 1$, as a public signal of aggregate productivity A. Moreover, at $t = 1$, producers receive a private signal of A.

$$s_i = \log A + \varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}\left(0, \tau_s^{-1}\right), \tag{5}$$

They use both public and private signals to form expectations about aggregate productivity and to determine their production decisions. The random noise ε_i is orthogonal to $\log A$ and to the noise in the signal of the other producers.

I will now use the model to study the implications of a supply shifter $\xi \neq 0$ for the equilibrium in the oil market (in $t = 1$) and how this equilibrium is affected when an

[21] Adding an idiosyncratic component to differentiate productivity across islands does not change the results but complicates the analysis.

OPEC announcement reveals ξ to the producers.

2.2 Equilibrium in the oil market and OPEC announcements

One advantage of the model in Sockin and Xiong (2015) is that there is a unique log-linear equilibrium in closed form. The equilibrium conditions for the oil market at $t = 1$ are given in Eq. (6) and (8). Oil price is a log-linear function of productivity A and the oil supply shifter ξ:

$$\log P_X = h_A \log A + h_\xi \xi + h_0, \tag{6}$$

where $h_A > 0$ and $h_\xi < 0$. Eq. (6) illustrates that, from the perspective of the producers, the oil price aggregates all dispersed information about $\log A$, partially revealing it, and the supply shifter ξ works as a noise. In equilibrium, producer i's demand for oil is given by

$$\log X_i = l_s s_i + l_P \log P_X + l_0, \tag{7}$$

where $l_s > 0$ and l_P is indeterminate.[22] The aggregate demand for oil can be derived integrating the demand from producers over the noise in their private signal:

$$\log X_S = l_P h_\xi \xi + (l_s + l_P h_A) \log A + l_0 + l_P h_0 + \frac{1}{2} l_s^2 \tau_s^{-1}. \tag{8}$$

The expressions for h_A, h_ξ, h_0, l_s, l_P, and l_0 are given in the original paper.

Now we can focus on what happens when an OPEC announcement reveals information about oil supply. Assume that the announcement reveals the oil supply shifter ξ. It is clear from Eq. (6) that this allows the producers to disentangle the contributions of ξ and $\log A$ to the price of oil. In other words, the model becomes one of perfect

[22]As described so far, the model shows that a higher commodity price increases costs for producers, but also signals a stronger economy, which pushes firms to increase production. The parameter l_P captures this tension. Sockin and Xiong (2015) discusses the conditions under which the informational effect dominates the cost effect, leading to a positive price elasticity of producers' demand for the commodity. Importantly, the results of this section do not require $l_P > 0$. In other words, results go through even if the informational effect is not particularly strong. See Kilian and Zhou (2023) for a critique of the conclusions that Sockin and Xiong (2015) draw from their model.

information.[23] The equilibrium conditions under perfect information are given by

$$\log P'_X = \frac{1}{1+k(1-\phi)}\log A - \frac{1-\phi}{1+k(1-\phi)}\xi + \frac{1}{1+k(1-\phi)}\log\phi, \tag{9}$$

$$\log X'_S = \frac{k}{1+k\,(1-\phi)}\log A + \frac{1}{1+k\,(1-\phi)}\xi + \frac{k}{1+k\,(1-\phi)}\log\phi, \tag{10}$$

where $\log P'_X$ is the price of oil, $\log X'_S$ is the aggregate demand for oil, and the prime distinguishes the equilibrium conditions under perfect information from those under information frictions.

Case 1: OPEC only reveals ξ

Let us first consider the case in which OPEC reveals the oil supply shifter ξ, without announcing any new unexpected changes in supply. In other words, OPEC reveals the *current* oil supply to the producers. Without loss of generality, assume for the rest of the section that $\xi > 0$, which means an unexpected increase in oil supply. In this scenario, the price and demand surprises that follow the announcement are obtained by subtracting Eq. (6) from Eq. (9), and Eq. (8) from Eq. (10):

$$\log P'_X - \log P_X = \tilde{h}_A \log A + \tilde{h}_\xi \xi + \tilde{h}_0, \tag{11}$$

$$\log X'_S - \log X_S = k\tilde{h}_A \log A + k\tilde{h}_\xi \xi + k\tilde{h}_0, \tag{12}$$

According to Proposition 4 in Sockin and Xiong (2015), under information frictions, $h_A > 0$ and $h_\xi < 0$ are both lower than their corresponding values under perfect information. Therefore

$$\tilde{h}_A = \frac{1}{1+k(1-\phi)} - h_A > 0,$$

$$\tilde{h}_\xi = -\frac{1-\phi}{1+k(1-\phi)} - h_\xi > 0.$$

Equations (11) and (12) imply that an OPEC announcement that reveals the current supply shifter ξ induces *positive* revisions both in the price of oil and the quantity

[23] The same reasoning and mechanisms would apply if we were to assume that OPEC sends a public signal about ξ, instead of revealing ξ itself. Although more realistic, that assumption would complicate the derivations without adding much to the intuition.

demanded. The reason for this is that under imperfect information (before the OPEC announcement), markets attribute part of the decrease in prices due to ξ to a drop in $\log A$, which depresses prices below the full-information level.

Case 2: OPEC reveals ξ and announces an additional shift ξ'

Let us now assume that OPEC also surprises the market with a new oil supply shock ξ', so that the total oil supply shifter under full-information becomes $\xi + \xi'$. Again, without loss of generality, assume $\xi' > 0$. The price and demand surprises become

$$\begin{aligned}\log P'_X - \log P_X &= \tilde{h}_A \log A + \tilde{h}_\xi \xi - \frac{1-\phi}{1+k(1-\phi)}\xi' + \tilde{h}_0, \\ \log X'_S - \log X_S &= k\tilde{h}_A \log A + k\tilde{h}_\xi \xi + \frac{1}{1+k(1-\phi)}\xi' + k\tilde{h}_0,\end{aligned}$$

As expected, the positive oil supply shock ξ' has a negative effect on the price of oil and a positive effect on the aggregate demand for oil. This is – within the limitations of the model – the oil supply news shock that Känzig (2021) wants to isolate.[24] Instead, as in the previous scenario, the revelation of ξ has a positive effect on both price and aggregate demand. This shows that revisions in the price of oil following the announcement can be driven both by an oil supply shock and an information shock. Therefore, surprises in the price of oil alone cannot be used as an exogenous measure of oil supply news shocks.

This result provides the rationale behind the identification strategy proposed in this paper. A positive co-movement of surprises in the oil price and surprises in the stock price – which I use as a proxy for production decisions – identifies the information shock, while a negative co-movement identifies the oil supply news shock.

[24] By 'limitations' I mean that, in the model, an oil supply shock affects supply contemporaneously. This is different from an oil news shock, which affects prices on impact but not necessarily quantities. In order to obtain the results described here in a more general model in which ξ represents a shift in the future trajectory of oil production, one would need to add differently informed producers. Assuming a positive shift in the trajectory ($\xi > 0$), more informed producers know that oil supply will be higher, which puts downward pressure on the price. Less informed producers misinterpret the price drop as a negative demand shock, which further depresses the price. From there, the same logic determining the results in the model presented here applies.

3 Methodology and data

I estimate the impulse responses of a wide set of global macroeconomic aggregates to the oil supply news shock and to the information shock. I consider three main specifications: the six-variable VAR of Känzig (2021), a medium-scale 16-variable global VAR, and a set of 30 VARs for advanced and emerging economies that are aggregated to obtain the dynamic responses of the median country. All models include 12 lags of the endogenous variables. The models are estimated using Bayesian techniques that efficiently deal with the high dimensionality of the systems. The priors imposed are standard Normal-Inverse-Wishart. The shocks of interest are identified using external instruments (Stock and Watson, 2012; Mertens and Ravn, 2013; Stock and Watson, 2018).[25] Obtaining an exogenous instrument for the identification of oil supply news shocks is the main methodological contribution of this paper.

3.1 Construction of the robust instruments

Two distinct methodologies can be adopted to disentangle the demand and the supply components in the surprises in the price of oil futures. Both methodologies have been used in the context of monetary policy to separate the policy and the information components in the high-frequency surprises in interest rates computed around monetary policy announcements. The first one exploits the high-frequency co-movement between oil futures and asset prices (Jarociński and Karadi, 2020; Cieslak and Schrimpf, 2019; Cieslak and Pang, 2020). The second one separates the informational component of the surprises in oil futures by directly controlling for the information set of OPEC (Miranda-Agrippino and Ricco, 2021; Romer and Romer, 2000). I adapt both methodologies to the study of the oil market and show that they are equally effective in disentangling the demand and the supply components in the surprises. However, the second methodology is presented as a robustness check because the OPEC's Monthly Oil Market Reports – used to approximate the information set of OPEC – are only available for a relatively short span, limiting the sample size.

[25] In Section C.4 I show that results are robust to using an internal instrument approach.

3.1.1 Construction of the daily surprises in oil futures

The daily surprises in the price of oil futures are provided by Känzig (2021).[26] They are obtained as follows. First, compute the daily change in the price of West Texas Intermediate (WTI) crude futures contracts with maturities from the front month to the 12th month on the closing days of OPEC conferences. These contracts have been traded on the New York Mercantile Exchange (NYMEX) since March 1983. These maturities are the most liquid (Alquist and Kilian, 2010). Second, estimate the first principal component in the term structure of the daily surprises. Between July 1983 and June 2023 there were 151 OPEC announcement days. The surprises, that are now at OPEC meeting frequency, are aggregated at monthly frequency to match the frequency of the data used in the VAR analysis. In months with more than one OPEC announcements, the monthly surprise is the sum of the daily surprises. In months with no OPEC announcement, the monthly surprise is set to zero.

The conference dates can be obtained from the OPEC website. OPEC conferences bring together delegations from each member country and are held at least twice a year.[27] Their duration varies between one day and one week. The conference dates are well publicised. Two press meetings are held during the conferences. One at the beginning and one at the end. In the final press meeting, any decision to adjust production quotas is formally announced via a communiqué, followed by a Q&A session. Importantly, before announcing the production decision, the communiqué provides a review of the oil market outlook. The final press release usually takes place after the end of the meetings, so that markets effectively observe quota decisions on the last day of the conference.

[26]The surprises updated to June 2023 can be downloaded from Diego Känzig website at `https://github.com/dkaenzig/oilsupplynews`. In a previous draft I had extended these surprises from 2017 onwards myself. There were a few minor differences relative to Känzig's version.

[27]OPEC was founded in 1960 by Iran, Iraq, Kuwait, Saudi Arabia, and Venezuela. Currently, it counts 12 members. In addition to the founding members, the following countries are part of OPEC accession date in parentheses): Algeria (1969), Congo (2018), Equatorial Guinea (2017), Gabon (1975–1995; 2016–present), Libya (1962), Nigeria (1971), and the United Arab Emirates (1967). Other countries have been part of OPEC over the years: Qatar (1961–2019), Indonesia (1962–2009; 2016:1–2016:11), Ecuador (1973–1992; 2007–2020), and Angola (2007–2024).

3.1.2 Disentangling demand and oil supply components

The first methodology relies on the high-frequency co-movement of oil futures prices and asset prices (Jarociński and Karadi, 2020; Cieslak and Schrimpf, 2019; Cieslak and Pang, 2020). As the model in Section 2 shows, a positive co-movement of revisions in oil futures and economic conditions – which I proxy with stock prices – identifies the information shock, while a negative co-movement identifies the oil supply news shock. This suggests a division in two of the time series of daily surprises in the price of oil futures. For each announcement day, the revisions in oil futures that induce a surprise in the stock price index of the opposite sign are stored in the instrument for the oil news shock. The revisions that induce a surprise in the stock price index of the same sign are stored in the instrument for the information shock. Of the 151 announcement days that occur between 19/07/1983 and 04/06/2023, 74 are classified as events revealing the oil news shock, and the remaining 77 are classified as events revealing the information shock. The surprises are then aggregated at monthly frequency as follows. For both instruments, in months with more than one OPEC announcement that induces the co-movement of the correct sign, the monthly surprise is the sum of the daily surprises. In months with no OPEC announcement that induces the co-movement of the correct sign, the monthly surprise is set to zero.

The stock price index used is the S&P 500, which is an index of the 500 largest publicly-traded companies in the U.S. weighted by market capitalisation. However, the results depend neither on the use of a specific stock price index nor on the country where the index is based. This is consistent with the evidence on the worldwide co-movement in risky asset prices and with the global nature of the shocks that affect the oil market (Miranda-Agrippino and Rey, 2020, 2021). In Section 5.3, the main results of the paper are replicated using alternative stock price indices. Namely, the Datastream (DS) World stock price index, the DS Airlines index, the TOPIX (an index of all firms in the first section of the Tokyo Stock Exchange), and the KOSPI 200 (an index of the 200 largest companies traded on the Korean Stock Exchange). I do not use a stock price index specific for the oil sector because the price of energy stocks tend to comove with the oil price regardless of the type of shock, and that would invalidate the identification strategy. This issue could potentially affect also the S&P 500, but the share of energy-related companies listed in that index is quite limited, around 3%.

The two robust instruments obtained with this methodology are shown in the upper panel of Figure 2. The instrument for the oil news shock is represented by the blue bars. The instrument for the information shock is represented by the red bars. The lower panel displays the evolution of the nominal WTI crude price over the same period. Three takeaways are worth mentioning. First, surprises that co-move positively with stock prices are not bigger or more frequent during recessions. This is an indication that the surprises are not simply capturing the downward co-movement of oil prices and stock prices that is common during recessions. Second, larger surprises do not necessarily correspond to larger swings in the spot price. This suggests that the identification of the shocks is not driven by a few peculiar events. Third, surprises are evenly distributed over time, indicating that information effects are not limited to a specific period.

One might worry that revisions in the price of oil futures on OPEC announcement days might be similar in magnitude to the price revisions on any other day. This would be concerning because the surprises might capture background noise rather than the consequences of the announcement, therefore invalidating the identification. Figure 3 shows that the variance of the movements in the price of oil futures is higher on announcement days compared to non-announcement days (see also Pescatori and Nazer, 2022). The left panel compares announcement and non-announcement days characterised by a negative co-movement of oil futures and stock prices. The right panel displays the same comparison for positive co-movements. Each subplot displays the distribution of surprises on announcement days (solid blue) and on non-announcement days (dashed red) for a specific maturity of the futures contracts. In all cases, the density on announcement days has fatter tails than the density on non-announcement days.[28] The subplots also report the p-values of Brown–Forsythe tests for the equality of group variances performed for each maturity. In all cases except one, the null hypothesis of equal variances is rejected at 10%.[29] Non-announcement days are randomly selected in equal number to the announcement days. Repeating the tests with different draws of non-announcement days does not alter the results.

[28]This is immediately seen by comparing the height of the dashed-red density at the mode with the height of the solid-blue density at the mode.

[29]This is true also for the front month contract, which is not shown to fit the figure.

Figure 2: Robust instruments

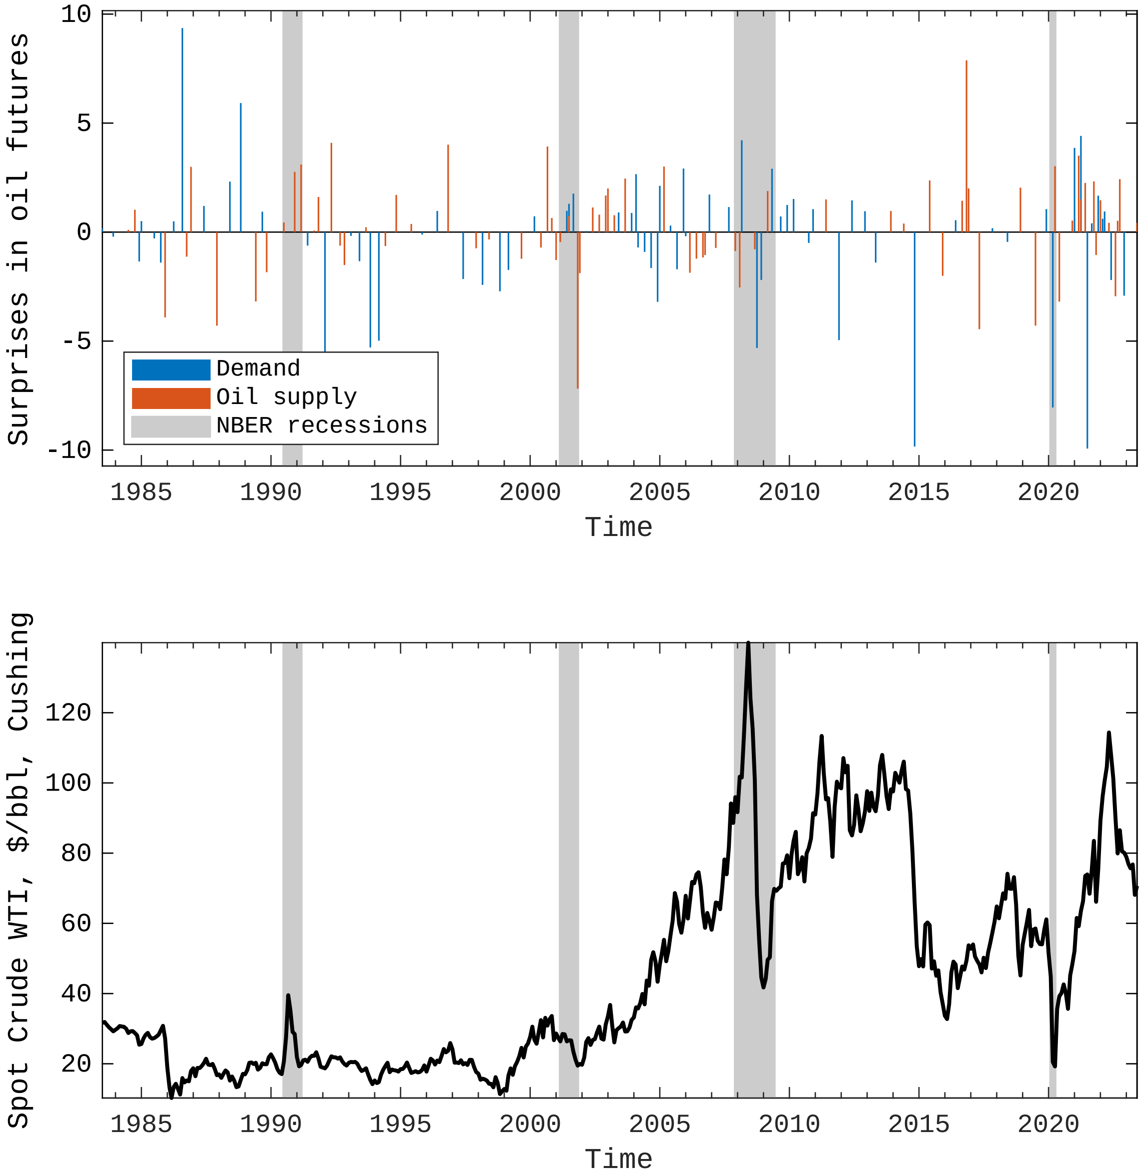

Note: Upper panel: the robust instruments for oil supply news shocks (red) and information shocks (blue) obtained by separating the surprises in oil futures on the basis of their co-movement with the daily changes in stock prices. Lower panel: nominal spot crude WTI price in dollars per barrel, for delivery in Cushing (OK). Grey areas represent NBER recessions.

3.2 Identification by external instruments

Identification of the structural shocks of interest is based on the Proxy SVAR/IV-SVAR approach (Stock and Watson, 2012; Mertens and Ravn, 2013; Stock and Watson, 2018).[30]

[30] Unless otherwise stated, the shocks are identified using the instruments sequentially rather than jointly.

Figure 3: Surprises on announcement and non-announcement days

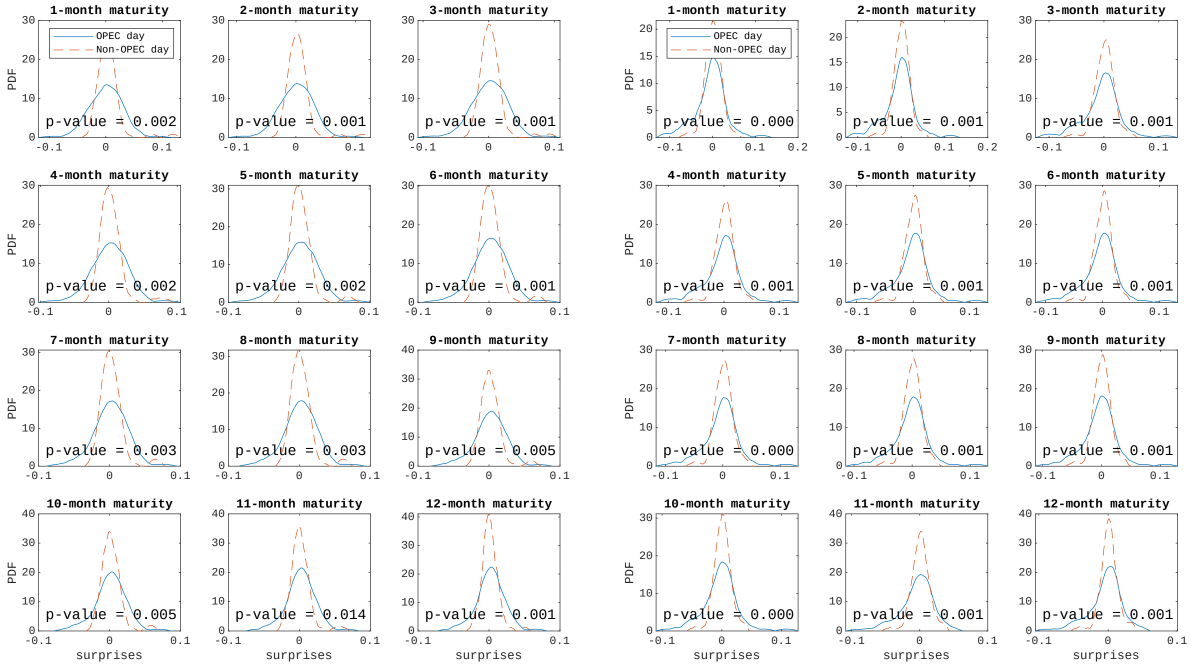

(a) Negative co-movement (b) Positive co-movement

Note: Comparison between daily surprises in oil future prices on OPEC announcement days (solid blue) compared to non-announcement days (dashed red). If the dashed-red bell is 'higher' than the solid-blue one, then surprises on announcement days are larger in magnitude than on any other day. Left panel: days characterised by a negative co-movement of oil futures and stock prices; right panel: days characterised by a positive co-movement. Reported p-values are for Brown–Forsythe tests for the equality of group variances performed for each maturity. The front month contract is not shown. Announcement and non-announcement groups contain the same number of observations.

Consider the following reduced-form VAR(p) model,

$$\boldsymbol{Y}_t = \boldsymbol{c} + \sum_{\ell=1}^{p} \boldsymbol{A}_\ell \boldsymbol{Y}_{t-\ell} + \boldsymbol{e}_t, \tag{13}$$

where $\boldsymbol{Y}_t$ is a $n \times 1$ vector of endogenous variables, $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_p$ are $n \times n$ matrices collecting the autoregressive coefficients, $\boldsymbol{c}$ is a $n \times 1$ vector of intercepts, p is the lag order, and $\boldsymbol{e}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$ is a $n \times 1$ vector of mean-zero innovations with covariance matrix $\boldsymbol{\Sigma}$.

Assume that the innovations are linear combinations of the structural shocks such that the following condition holds,

$$\boldsymbol{e}_t = \boldsymbol{B}\boldsymbol{u}_t, \tag{14}$$

where $\boldsymbol{u}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})$ is a $n \times 1$ vector of structural shocks and $\boldsymbol{B}$ is a $n \times n$ matrix of impacts. The identification problem arises because the covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{B}\boldsymbol{\Omega}\boldsymbol{B}'$ only provides $n(n+1)/2$ restrictions to identify the n^2 free parameters in $\boldsymbol{B}$.

For ease of exposition, assume that we are interested in identifying the impact of only one shock which, without loss of generality, is ordered first in the vector $\boldsymbol{u}_t$.[31] Partition Eq. (14) as follows:

$$\begin{bmatrix} e_t^1 \\ \boldsymbol{e}_t^{2:n} \end{bmatrix} = \begin{bmatrix} \boldsymbol{b}_1 & \boldsymbol{B}_{2:n} \end{bmatrix} \begin{bmatrix} u_t^1 \\ \boldsymbol{u}_t^{2:n} \end{bmatrix},$$

where the notation $2:n$ indicates columns 2 to n of the underlying matrix, or element 2 to n of the underlying vector. The vector $\boldsymbol{b}_1$, which is the object that we want to identify, represents the $n \times 1$ impact vector for the shock of interest u_t^1. The Proxy SVAR/IV-SVAR identification strategy exploits external instruments to isolate exogenous variation in the innovations of the VAR that is due to the structural shock of interest. A valid instrument needs to satisfy the two conditions of relevance and exogeneity. These conditions can be summarized as follows:

$$\mathbb{E}\left(z_t u_t^1\right) = \phi \neq 0, \tag{15}$$

$$\mathbb{E}\left(z_t \boldsymbol{u}_t^{2:n}\right) = \mathbf{0}. \tag{16}$$

If the conditions are satisfied, then the impact vector $\boldsymbol{b}_1$ is identified up to sign and scale, as shown in the following expression:

$$\mathbb{E}\left(z_t \boldsymbol{e}_t\right) = \boldsymbol{B}\mathbb{E}\left(z_t \boldsymbol{u}_t\right) = \begin{bmatrix} \boldsymbol{b}_1 & \boldsymbol{B}_{2:n} \end{bmatrix} \begin{bmatrix} \mathbb{E}\left(z_t u_t^1\right) \\ \mathbb{E}\left(z_t \boldsymbol{u}_t^{2:n}\right) \end{bmatrix} = \boldsymbol{b}_1 \phi \tag{17}$$

In practice, the impact vector $\boldsymbol{b}_1$ is then normalised so that a unitary impulse in

[31] This will be the case for the rest of the paper. However, the methodology allows to identify all n shocks, conditional on having at least n instruments that satisfy the conditions of relevance and exogeneity. Point identification in this case also requires the imposition of additional identifying restrictions (see for instance Giacomini et al., 2021).

u_t^1 induces a unitary response in one of the endogenous variables. Given the following partitioning,

$$\mathbb{E}\left(z_t \boldsymbol{e}_t\right) = \begin{bmatrix} \mathbb{E}\left(z_t e_t^1\right) \\ \mathbb{E}\left(z_t \boldsymbol{e}_t^{2:n}\right) \end{bmatrix} = \begin{bmatrix} b_{1,1}\phi \\ \boldsymbol{b}_{2:n,1}\phi \end{bmatrix}$$

one can immediately obtain

$$\boldsymbol{b}_{2:n,1} b_{1,1}^{-1} = \mathbb{E}\left(z_t \boldsymbol{e}_t^{2:n}\right) \mathbb{E}\left(z_t e_t^1\right)^{-1}$$

which identifies the impact vector $\boldsymbol{b}_1$ up to a scale.

3.3 Bayesian Vector Autoregressions

The prior adopted in the empirical analysis is a standard Normal-Inverse-Wishart prior (Litterman, 1986; Kadiyala and Karlsson, 1997). It formalises the view that an independent random-walk model for each variable in the system is a reasonable centre for the beliefs about their time series behaviour (Sims and Zha, 1998). The prior is imposed by setting the following moments for the prior distribution of the coefficients:

$$\mathbb{E}\left[(\boldsymbol{A}_\ell)_{ij}|\boldsymbol{\Sigma}\right] = \begin{cases} \delta_i & j = i,\ \ell = 1 \\ 0 & \text{otherwise} \end{cases}, \qquad \mathbb{V}\left[(\boldsymbol{A}_\ell)_{ij}|\boldsymbol{\Sigma}\right] = \begin{cases} \frac{\lambda^2}{\ell^2} & \text{for } j = i, \forall \ell \\ \frac{\lambda^2}{\ell^2}\frac{\sigma_i^2}{\sigma_j^2} & \text{for } j \neq i, \forall \ell \end{cases}, \tag{18}$$

where $(\boldsymbol{A}_\ell)_{ij}$ denotes the coefficient on variable j in equation i at lag ℓ and δ_i is either 1 for variables that display a trending behaviour or 0 for variables that are bounded or appear as stationary.[32] The prior assumes the coefficients $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_p$ to be a priori independent and normally distributed. It also assumes that the most recent lags of a variable tend to be more informative than distant lags. This is represented by ℓ^2. The hyperparameters $\sigma_1^2, \ldots, \sigma_n^2$ are set using sample information and equal the variance of the residuals from a univariate autoregressive model of order 1 for each variable in the system. The term σ_i^2/σ_j^2 accounts for differences in the scales of variable j relative to variable i. The hyperparameter λ controls the overall tightness of the prior. The tightness is estimated using the optimal prior selection approach proposed by Giannone

[32] This reflects the idea that a white noise process, rather than a random walk, is a better reference for variables characterised by high mean reversion.

et al. (2015). The prior is cast by means of dummy observations (Bańbura et al., 2010).

3.4 Estimation of median-group responses

To estimate the dynamic response of the median advanced and emerging economies to the shocks I rely on the median-group estimator used in Degasperi et al. (2021). This estimator aggregates the country-specific responses to obtain the median response across countries. Importantly, it allows both the intercepts and the slope parameters to vary across countries. This accounts for the potentially high degree of dynamic heterogeneity across countries, especially in the case of emerging markets (see Ciccarelli and Canova, 2009, for a discussion).

The median responses are estimated in two steps. First, the country-specific models are estimated for all countries in the group. The models are VAR(12). The credibility regions for the impulse responses of each model are estimated using a standard Gibbs sampler. Second, for each country-specific variable in the VAR, I stack the structural impulse responses across countries and I compute the median across countries at each horizon. For global variables, I simply stack the structural impulse responses without computing the median. This delivers a set of median structural impulse responses for the underlying group of countries. These responses are then summarised by displaying the median response, the 68% and 90% credibility regions.

3.5 Data

All variables used in the empirical exercises are collected at monthly frequency. If series are available at a daily frequency, the end-of-month value is used.[33]

Six-variable VAR. To ensure comparability, the specification used for the six-variable VAR follows Känzig (2021). The variables included are: the real oil price, world oil production, world oil inventories, world industrial production, U.S. industrial production, and U.S. CPI. In figures displaying the impulse responses, the shocks are always normalised to induce a 10% increase in the real oil price. The real oil price is constructed by deflating the end-of-month WTI spot crude oil price by the U.S. consumer

[33] Using monthly averages instead does not change the results.

price index (CPI). The measure of world oil inventories is taken from Kilian and Murphy (2014). World industrial production measures the production of OECD countries plus six major emerging markets (Brazil, China, India, Indonesia, Russia, and South Africa) and is taken from Baumeister and Hamilton (2019).

Global VAR. The global system contains 16 variables. Three variables relate to the global oil market: the real oil price, world oil production, and world oil inventories. Four variables capture global economic conditions: world industrial production, CPI for advanced economies (AE) excluding the U.S., CPI for emerging economies (EM), and a stock price index for OECD economies (excluding North America). The system also includes the U.S. nominal effective exchange rate (narrow) and a commodity price indices that excludes energy-related commodities (provided by the World Bank). Finally, I include seven variables for the U.S. economy: the Federal Funds rate, the 1-year and 10-year constant maturity treasury rates, the CPI, the S&P 500, the excess bond premium by Gilchrist and Zakrajšek (2012), and the CBOE VIX index. The excess bond premium is a measure of corporate credit spreads and captures risk appetite in the corporate bond market. The VIX is a measure of volatility in the S&P 500 and captures uncertainty in the financial markets.[34]

Advanced and emerging economies. The results on the transmission of shocks to the global economy are complemented with an additional exercise that separately focusses on advanced and emerging economies, following Degasperi et al. (2021). VARs for 15 advanced and 15 emerging economies are estimated and aggregated to obtain the responses of the median advanced and emerging economies. Each system includes 12 variables (11 in the case of emerging economies, as core CPI in not included). The first 6 variables represent the domestic economy: industrial production, CPI, core CPI, stock price index, nominal bilateral exchange rate, and policy rate. The remaining 6 capture the global economy and oil market: real oil price, world oil production, and world oil inventories, world industrial production, VXO, and U.S. 1-year constant maturity treasury rate. The countries included in the analysis and the relative sample size are reported in Table B.1.

[34]The VIX, before 1990:1, is replaced by the VXO. The VXO, prior to 1986:1, is reconstructed following Bloom, 2009.

Inflation expectations. To study the transmission of the shocks to inflation expectations, I augment the 6-variable VAR model with financial variables (the S&P 500 and the VIX) and with a set of inflation expectations measures. The measures used are the St. Louis Fed's 5-Year–5-Year forward inflation expectation rate, computed from 10-year and 5-year nominal and inflation adjusted Treasury securities, the University of Michigan 1-year inflation expectation of U.S. households, and the Cleveland Fed's 1-year and 5-year inflation expectations, which are based on a model that uses as input Treasury yields, inflation data, inflation swaps, and survey-based measures of inflation expectations.

Data for robustness exercises. The high-frequency co-movement between oil futures and the stock price index can be used to disentangle the shocks because stock prices are a high-frequency proxy for economic activity. However, they are not the only high-frequency measure of economic activity available. In a robustness check, the shocks are separated based on the co-movement of oil futures and the daily measure of U.S. real business conditions proposed by Aruoba et al. (2009). In another exercise, it is shown that results are robust to using alternative measures of world industrial production. The measure used throughout the paper is the OECD-plus-six index by Baumeister and Hamilton (2019). The alternative measures are the Dallas Fed world (excluding U.S.) industrial production and the OECD industrial production from the OECD Main Economic Indicators.

4 Results

In this section I present the results. First, using the full set of daily surprises in the price of oil futures as an instrument to identify the oil news shock gives rise to puzzles. Second, these puzzles disappear when the two underlying shocks – the oil news and the information shocks – are identified using the robust instruments based on the high-frequency co-movement of oil futures and stock prices. An oil news shock that raises oil prices has an unambiguous contractionary effect on world industrial production. Conversely, the information shock, similarly normalised, transmits as a demand shock and drives an expansion in economic activity. Third, the robust instruments can be used to identify the effects of the shocks in larger models that better capture features of the global economy and address potential issues of informational insufficiency in smaller VARs. Results are presented for a model of the global economy, and for the median advanced and emerging economies. The shocks have significant effects on global stock markets, exchange rates, credit conditions, and uncertainty measures. Moreover, they induce an endogenous response from monetary policy authorities worldwide.

4.1 Price revisions as instrument give rise to puzzles

The results obtained using the full set of daily surprises as instrument for the oil news shock are not robust to changes to the length of the sample. In particular, shortening the estimation sample gives rise to output puzzles in which an oil news shock that increases the price of oil appears to *increase* industrial production. This is in line with the predictions of the model, that the surprises confound oil supply shocks with demand shocks. As seen in Section 2, these two shocks move economic activity in opposite directions, therefore giving rise to puzzles when used as exogenous measures of changes in oil supply expectations. An oil news shock that increases the price of oil, by immediately increasing marginal costs, is supposed to reduce economic activity.

Figure 4 documents the lack of robustness to changes in the sample size and the presence of puzzles. The left panel replicates the main result in Känzig (2021) using the extended series of surprises in oil futures.[35] The shock is normalised to induce a 10% increase in the real oil price. Results are virtually identical to the original despite the minor differences in methodology. World oil production, and world and U.S. industrial production contract with a lag, while world oil inventories and U.S. CPI expand.

However, repeating the estimation on a shorter sample substantially alters the results. The right panel shows the impulse responses to the same shock identified on the samples 1982:7–2019:12 (in red) and 1990:1–2019:12 (in blue). World oil production does not contract. World industrial production expands significantly for at least 6 months

[35]The series of surprises used in Känzig (2021) spans the period 1983:4 to 2017:12. Here they are extended to 2023:6. Moreover, the VAR sample used in that paper starts in 1974:1, whereas here it starts in 1975:1. Another difference is that the confidence bands for the impulse responses in Känzig (2021) are obtained by bootstrap. Here, I use Bayesian methods that in general deliver smoother and tighter credibility regions. Minor differences might be present also in the series used. Specifically, I do not seasonally adjust the world oil inventories series.

Figure 4: Sample dependence under non-robust identification

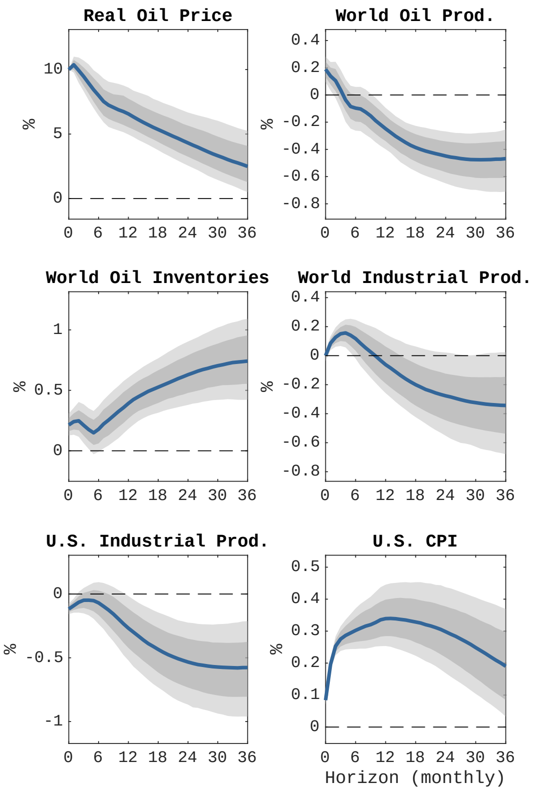

(a) Replica of Känzig (2021)

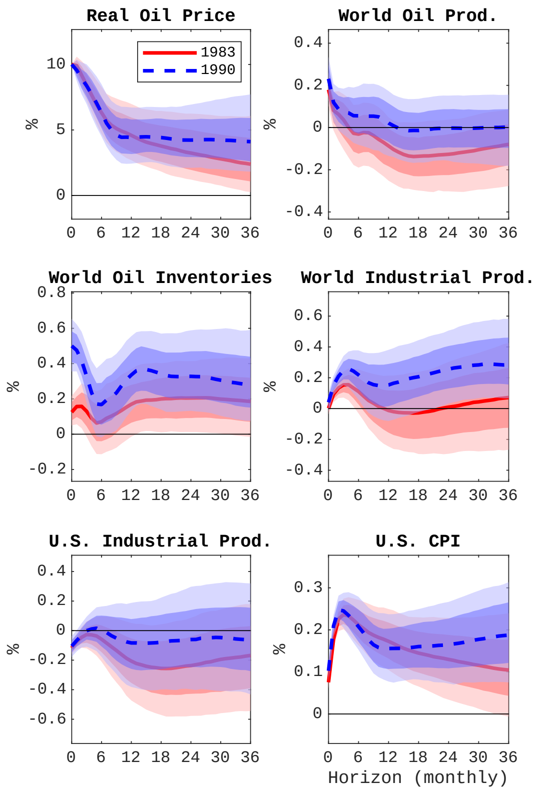

(b) Shorter samples

Note: Impulse responses to an oil supply news shock normalised to induce a 10% increase in real oil price. The shock is identified using the proxy provided in Känzig (2021), which spans the period 1983:7–2023:6. Left panel: BVAR(12) estimated on the sample 1975:1–2019:12. Right panel: BVAR(12) estimated on the samples 1982:7–2019:12 (solid red) and 1990:1–2019:12 (dashed blue). Shaded areas represent 68% and 90% posterior coverage bands.

after impact. The contraction in U.S. production and in world oil production attenuates and completely disappears on the sample starting in 1990. World oil inventories and U.S. CPI show an expansion, but the dynamics differ substantially from those obtained on the longer sample. Overall, the responses obtained on the shorter samples do not conform to the description of an oil supply shock and represent a puzzle.

4.2 Puzzles disappear using robust instruments

Results obtained using the robust instruments based on the high-frequency co-movement of oil futures and stock prices show no sign of the puzzles just discussed. A negative oil supply news shock has an unambiguous contractionary effect on world industrial

Figure 5: Responses under robust identification

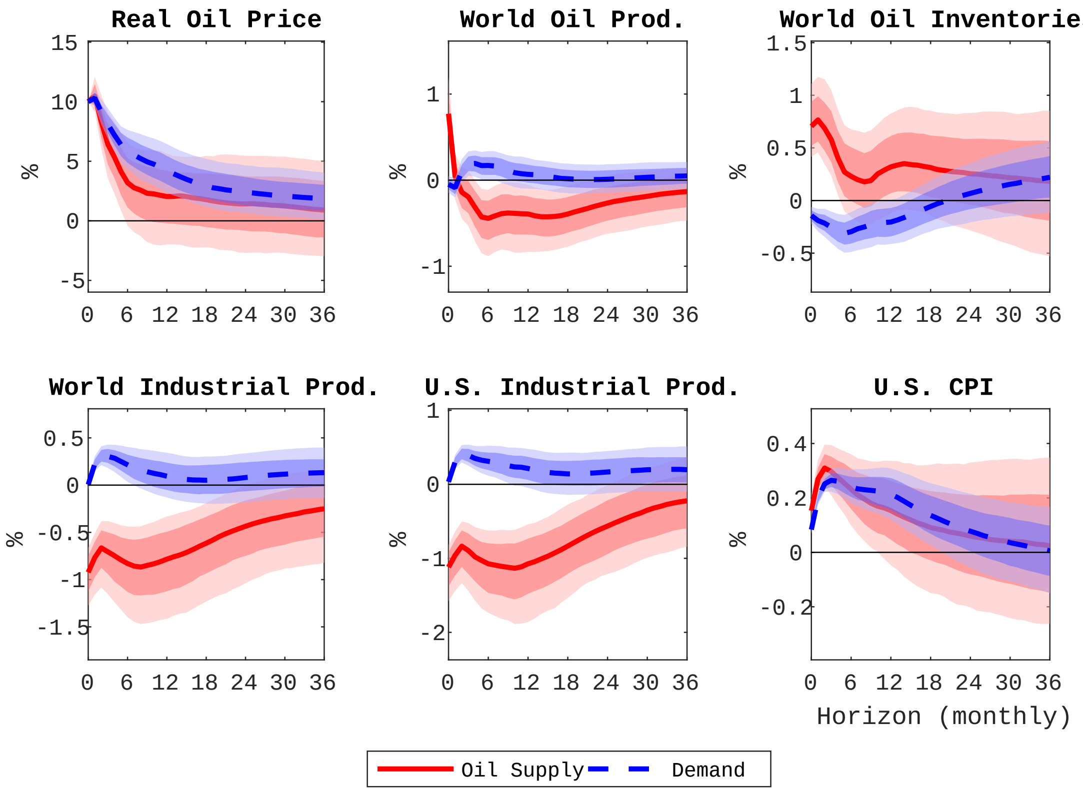

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1982:7–2023:5. Both proxies span the period 1983:7–2023:6.

production. Conversely, the information shock causes an expansion in economic activity. Additionally, responses to both shocks are now consistent across samples. Figure 5 only shows the results obtained on the sample 1982:7–2023:5, but results are robust to changing the sample size, as can be seen in Figure C.2.

An oil supply news shock normalised to increase the oil price by 10% causes roughly a 1% contraction in world and in U.S. industrial production on impact. This is true for all three estimation samples considered. The contractionary effects of the shock are persistent, lasting up to 36 months in the baseline setting. World oil production is also contracting with a lag, which is consistent with the lag in the implementation by OPEC of the announced quota cuts. The trough is at around -0.5%, 6 months after the shock.

U.S. CPI expands on impact and continues to increase, reaching a peak at +0.3% around 3 months from impact. World oil inventories also expand. In the baseline setting, the peak is reached at +0.7%, roughly 3 months from impact. However, the dynamics for this variable are slightly different across samples. This could be due to oil inventories being less responsive to shocks when closer to full capacity.

Importantly, the contemporaneous contraction in production and increase in inflation represents a challenge to U.S. monetary policy authorities. An increase in the policy rate would rein in inflation, but would put even more downward pressure on economic activity.

The effects of the information shock – which propagates like a demand shock – are one order of magnitude smaller than the effects of the supply shock. However, this explains why the responses identified with the surprises that mix the two shocks are unstable and puzzles emerge. Miranda-Agrippino and Ricco (2023) show that even small contamination of the proxy may lead to considerable bias in the estimated impulse responses. The size of the bias depends on how strongly the proxy correlates with the non-target shock and by how much variance of the variables of interest the non-target shock explains relative to the target one. Puzzles can emerge even in the presence of small contamination from pervasive shocks, such as demand shocks.

The information shock, normalised to increase real oil price by 10%, causes a 0.3% expansion in both world and U.S. industrial production that lasts around 6 months before reverting to trend. World oil inventories are also contracting by 0.3%, consistently with increased demand for oil. U.S. CPI and world oil production are increasing persistently, with a peak response of roughly +0.3% 3 months from the shock. Both shocks have a persistent effect on real oil price, with the variable not reverting to trend for at least 6 months for the oil supply news shock, and for at least 36 months for the information shock.[36]

Recently, Mori and Peersman (2024) show that the 6-variable VAR used in Känzig (2021) – and for comparability also here – is informationally insufficient (Forni and

[36] In the main text, the two shocks are identified one at the time, using the instruments sequentially. Figure C.3 in the Appendix reports the impulse responses to the two shocks when they are jointly identified using the two proxies and an additional recursive restriction (Mertens and Ravn, 2013). Results are essentially identical.

Gambetti, 2014). To solve the issue one needs to augment the set of endogenous variables with financial variables. It also shows that the puzzles I document here are attenuated when financial variables are included in the VAR. In Section C.5, in the Appendix, I show that solving the problem of informational insufficiency is not enough to obtain unbiased dynamic responses. The problems deriving from the failure of the proxy exogeneity assumption remain.[37]

In Section C.5, using the extended VAR model proposed by Mori and Peersman (2024) – which augments the 6-variable specification with the 1-year Treasury rate, the S&P 500, and the VIX – I study how much oil supply news and information shocks contribute to the variation in the real oil price between 1976 and 2023 (Figure C.9). I find that the contributions of the two shocks are similar in magnitude, although the estimate for the information shock is less precise. After 2000, the information shock gains prominence, while the contribution of the oil supply news shock fades. The historical decomposition captures well the developments in the oil market during the first wave of Covid-19. In February and March 2020, the information shock contributed significantly, indicating that negative demand shocks were weighing on the real oil price, whereas in April and May 2020, the oil supply news shock was putting upward pressure on prices, suggesting a contraction in supply. To counter collapsing demand and falling oil prices during the pandemic, in April 2020 OPEC+ agreed to cut production by 9.7 mbbl/day, and these cuts were implemented in May 2020.

I also provide a forecast error variance decomposition (Figure C.10). Both shocks account for a non-negligible share of the forecast error variance of the endogenous variables. The oil supply news shock accounts for about 15% of the variance in the real oil price, declining to 5% at horizon 36. The explained share of world oil production increases steadily, reaching 25% over the forecast horizon. This shock also accounts for a substantial share of the variation in world and U.S. industrial production, as well as U.S. CPI. Interestingly, the information shock explains the bulk of the forecast error variance in the real oil price. It also accounts for a sizeable portion of the variation in

[37]Nonetheless, informational insufficiency of the 6-variables VAR appears to be an issue. While the dynamic responses obtained with the robust instruments from a VAR that includes financial variables are essentially identical when estimated on the pre-Covid sample and on the full sample, estimates from the 6-variables VAR on the pre-Covid sample are problematic and substantially different from the results presented in Figure 5.

industrial production and CPI.[38]

As mentioned in the Introduction, another critical assumption for the validity of the identification strategy is that risk premia need not move within the event window. If they did, the price surprises would capture not only the oil supply news shock and the information effect that I isolate in this paper, but also an additional uncertainty shock. I do not attempt here to assess how relevant this concern is in practice or to isolate this additional shock (see Baumeister and Kilian, 2014; Baumeister, 2023, for a reference on how to approach the task). However, it is worth noting that such a shock would likely induce a negative comovement with the stock market. In this case, the methodology I use would misattribute part of the variation in the price surprises which is due to the uncertainty shock to the supply news shock. Therefore, even if risk premia moved within the window, that would not offer an alternative explanation for the output puzzles. Rather, it would attenuate the estimated impact of supply shocks.

4.3 Strength of the instruments

Although the relevance of the robust instruments is only marginally lower than that obtained using all surprises, the estimates discussed in this section point to a potential problem of instrumental weakness. Table 1 reports the F statistics for the regressions of the reduced-form VAR innovations corresponding to the real oil price equation on the instruments. The regression model is estimated for three instruments: (1) the whole set of surprises in oil futures (i.e. the instrument proposed by Känzig, 2021), (2) the robust instrument for the oil supply news shock, and (3) the instrument for the information shock. Results are obtained using residuals from the six-variables VAR estimated on two different samples: 1975:1–2023:5 and 1982:7–2023:5.

Stock et al. (2002) recommends a threshold for the F statistic of 10 or above to rule out weak instrument problems. The proxy for the oil supply news shock (column 2) falls marginally below this threshold (the robust F is 9.3 on the baseline sample and 10 on the extended one). The proxy for the information shock (column 3), although the standard F is quite high (28.7 on the baseline sample and 31.5 on the extended one),

[38]For both the historical and the forecast error variance decomposition, the shocks are identified sequentially. Results remain unchanged when the shocks are jointly identified by imposing an additional recursive restriction.

Table 1: First-stage F statistics

	1975:1–2023:5			1982:7–2023:5		
	(1)	(2)	(3)	(1)	(2)	(3)
Coefficient	1.535***	1.030*	1.878**	1.479***	1.059*	1.768***
N	479	479	479	479	479	479
F	34.53	5.438	31.54	32.99	5.947	28.71
F Robust	11.31	9.992	7.194	11.23	9.342	6.894

Note: F statistics for the regression of the reduced-form VAR innovations corresponding to the real oil price equation on the instrument and a constant. Results are reported for the six-variable VAR(12) estimated on two different samples (starting in 1975:1 and 1982:7 respectively) and for three different instrumental variables. (1) is the non-robust proxy of Känzig (2021), extended to 2023:6, (2) is the robust proxy for oil supply news shocks, (3) is the proxy for information shocks. The samples include the initial observations. F Robust allows for heteroscedasticity. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

is weak when considering the robust F statistic (6.9 on the baseline sample and 7.2 on the extended one). The relatively high F statistic obtained when using all surprises as instrument (column 1) is consistent with the fact that the surprises are conflating two different shocks that move the price of oil in the same direction. In this sense, contamination from other structural shocks can inflate the F statistics.

Weak instruments are problematic because they compromise the large-sample validity of standard inference (Montiel Olea et al., 2020). However, they do not invalidate posterior inference in a fully Bayesian setting (Caldara and Herbst, 2019; Arias et al., 2021, 2025). To check whether instrumental weakness might represent a problem in this context, I jointly identify the shocks and estimate impulse responses that fully incorporate the uncertainty about the correlation between the proxies and the shocks using the Bayesian algorithm of Arias et al. (2021). Figure 11 shows that, although credibility regions are generally wider, all results are verified. This suggests that the problem of instrumental weakness is not particularly severe.

Moreover, in Section C.4, I identify the shocks using an internal instrument Proxy-SVAR, which does not rely on the relevance assumption. Credibility regions are noticeably wider – especially for the oil supply news shock – but the results are confirmed in both the full and the pre-pandemic samples (Figures C.4 and C.5). The wider credibility regions may be attributed to noise in the instrument: noise that the external instrument approach effectively removes (see, for instance, Bruns and Lütkepohl, 2025).

4.4 Transmission to global aggregates

The study of the transmission of the two shocks can be extended to larger systems that incorporate important macro-financial indicators omitted in the six-variable VAR. Figure 6 shows the impulse responses of global macro aggregates to both shocks obtained from the global VAR. A negative oil supply news shock has strong and long-lasting contractionary effects on the global economy. A positive information shock has a relatively short-lived expansive effect that is smaller in magnitude.

Consistent with the results of the six-variable VAR, world industrial production and world oil production contract by about 1%, while oil inventories expand, following an oil supply news shock normalised to increase real oil price by 10%. Adding to the results from the smaller VAR, the shock contracts OECD and U.S. stock prices by roughly 5%, widens credit spreads in the U.S. (as measured by the excess bond premium), and causes an increase financial market uncertainty (as measured by the VIX). Importantly, the shock causes a 2% contraction in the price of non-energy commodities, indicating that the new instrument is unlikely to be contaminated by other demand shocks.

The shock has a similar, positive effect on U.S. CPI and the CPI indices of advanced economies. Contrarily, it depresses the CPI index of emerging markets, indicating that for this group of countries the wealth effect might dominate in the propagation to prices. The dollar appreciates. This might reflect the transition of the U.S. to net oil exporter around 2020. Indeed, the only visible difference between the responses shown here and those obtained on the pre-Covid sample is that, in the latter, the dollar depreciates, although the response is significant only at 68% (Figure C.18). The Fed appears to react by easing the monetary stance, although the response of interest rates is not significant. This might be due to the tradeoff between price and output stabilization.

On the other hand, following an information shock, normalized to increase the real oil price by 10%, real activity, CPI, world oil production, and equity prices expand sharply, while credit spreads narrow and uncertainty subsides. U.S. monetary policy endogenously responds to this demand shock with a monetary easing. The monetary stimulus normally transmits along the yield curve to higher maturities, indicating that

Figure 6: Transmission to the Global Economy

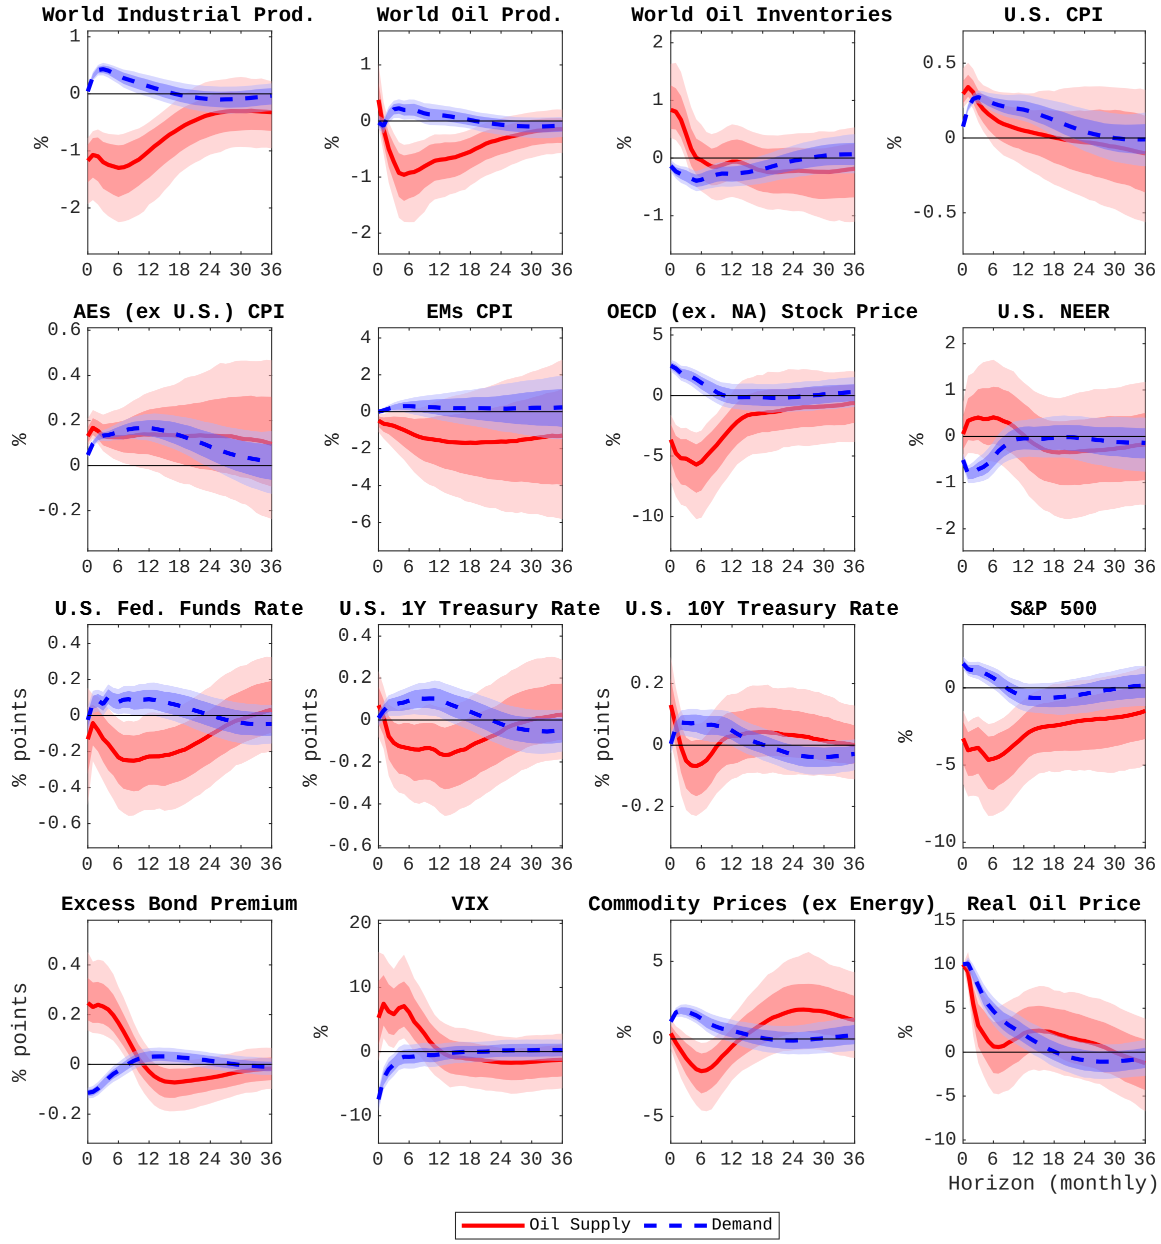

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1980:3–2023:5. Both proxies span the period 1983:7–2023:6.

the shock does not affect U.S. risk premia. The shock also increases the price of other commodities, consistently with an upward revision in demand expectations.

4.5 Transmission to advanced economies

This section is a complement to the result on global aggregates. Here I analyse the transmission of the two shocks to a set of country-specific VARs. By doing so I can study the effects of the shocks at a disaggregated level and compare the responses of specific groups of countries. In this section I focus on advanced economies and results for emerging markets are presented below.

Figure 7 shows the responses of the median advanced economy to the two shocks. Specifically, these are the median group responses computed by aggregating the individual country responses from VAR models that include domestic and global variables. The six variables shown model the median economy while the remaining six variables are global controls (Figure C.20). Table B.1 reports the sample size used for each country. Importantly, the sample for each country in the analysis excludes the Covid-19 pandemic. Results are in general consistent with those presented in the previous section. A negative oil supply news shock has a long-lasting contractionary effect on industrial production and stock prices in the median advanced economy, and a positive effect on prices. Although only marginally significantly, core CPI is also increasing, indicating that the inflationary effect of the shock spills over to non-energy sectors. The domestic exchange rate does not appear to move significantly vis-à-vis the U.S. dollar. There is also a short-lived tightening of the domestic monetary stance visible in the response of the policy rate, which may amplify the transmission to activity and stock prices.

A positive information shock causes an expansion in the median advanced economy. Industrial production and stock prices rise, the exchange rate appreciates relative to the dollar, and the policy rate unambiguously increases.

4.6 Transmission to emerging markets

Figure 8 shows the median-group responses of emerging markets to the two shocks. Results are stronger in magnitude when compared to those of advanced economies, with some striking differences in sign. A negative oil supply news shock has deep and persistent recessionary effects on the median emerging market. Following a shock that induces a 10% increase in the real oil price, industrial production contracts by 1%, stock

Figure 7: Transmission to advanced economies

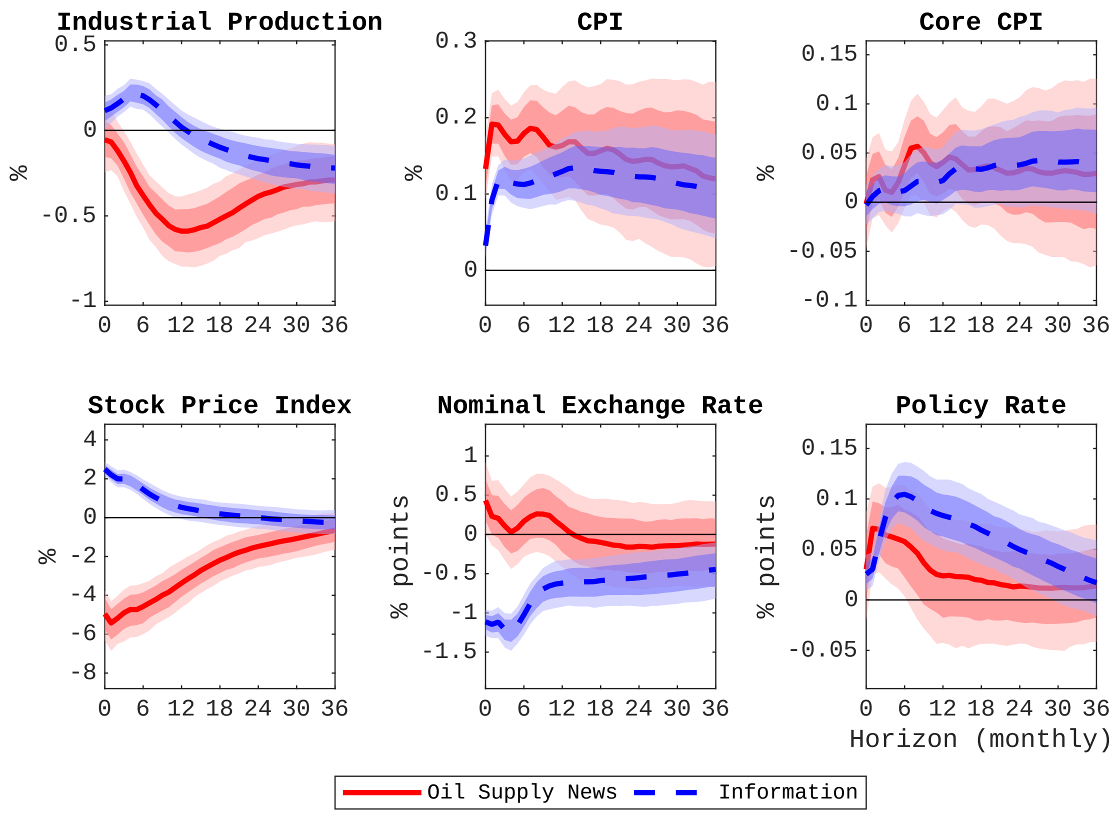

Note: Selected impulse responses for the median advanced economy. Solid red: responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1. Both proxies span the period 1983:7–2023:6. Full IRFs reported in Figure C.20.

prices drop by almost 20%, and the currency depreciates by 3% vis-à-vis the U.S. dollar. Prices contract, with a lag that matches the response of monetary policy, which is to loosen the stance. However, the response of the policy rate masks a high degree of underlying heterogeneity across countries. Heterogeneity in the response of the policy variable is to be expected, as these economies feature substantially different institutional and regulatory settings, and might decide to respond differently to the trade-off between higher inflation and larger contraction in production.

Responses to a positive information shock are expansionary and similar in sign, magnitude, and dynamics to the responses of the median advanced economy, with the exception of the policy rate, whose response is small and not significant.

Figure 8: Transmission to emerging markets

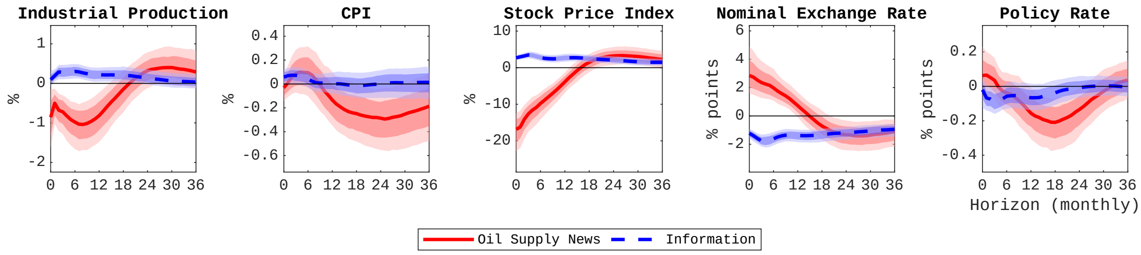

Note: Selected impulse responses for the median emerging market. Solid red: responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1. Both proxies span the period 1983:7–2023:6. Full IRFs reported in Figure C.21.

The responses show a remarkable degree of homogeneity across the 15 emerging markets in the sample. Following a negative oil supply news shock, all domestic currencies depreciate against the dollar, with the exception of South Africa (Figure C.22). The currencies of Hungary, India, and Russia do not respond significantly to the shock. Stock prices, with the exception of Colombia and South Africa, unambiguously contract for all EMs (Figure C.23). Table B.1 reports the sample size used for each country. As for the advanced economies, the sample excludes the Covid-19 pandemic.

4.7 The shocks affect long-run inflation expectations

A relevant question for policymakers is whether and by how much oil shocks can disanchor inflation expectations. Figure 9 reports the impulse responses of various measures of inflation expectations to the two shocks.[39] The 5-Year, 5-Year forward inflation expectation rate produced by the St. Louis Fed – a commonly used measure of average expected inflation over the five-year period that begins five years from today – responds significantly to the oil supply news shock (red line). Although short-lived, this effect indicates that the shock has the potential to move long-run inflation expectations away from the Fed's target. Also the 5-year inflation expectations measure produced by the

[39] The response of the 5-Year, 5-Year forward inflation expectation rate is obtained by estimating a VAR that augments the 6-variable VAR with the VIX, the S&P 500, and the expectation measure itself. The estimation sample spans the period 2003:1–2023:5. The responses of the other three expectation measures are estimated jointly in the 6-variable plus VIX and S&P 500 VAR. The sample is 1982:7–2023:5.

Figure 9: Transmission to inflation expectations

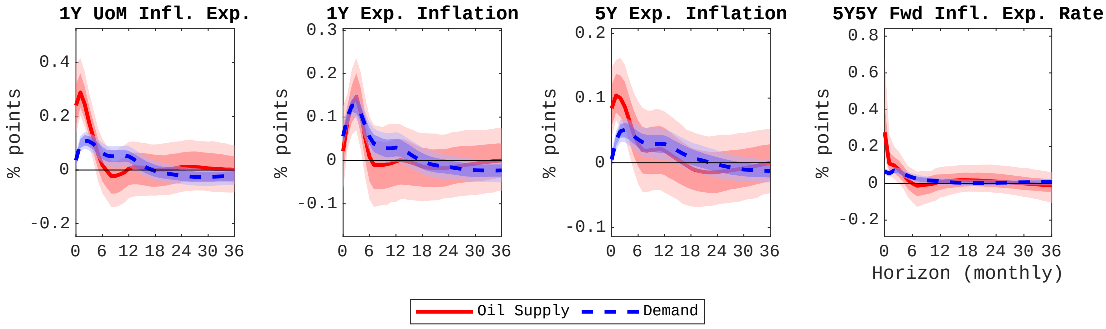

Note: Impulse responses of different measures of inflation expectations. Solid red: responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Both proxies span the period 1983:7–2023:6.

Cleveland Fed responds significantly. Moreover, inflation expectations move also at shorter horizons. Both the 1-year household inflation expectations measure produced by the University of Michigan and the measure produced by the Cleveland Fed point to a positive and significant response.

5 Robustness

This section provides additional support for the baseline results. First, I show that qualitatively similar results can be obtained by using an alternative methodology to separate supply and demand components in the surprises in oil futures that directly controls for the information set of OPEC. Second, I show that jointly identifying the shocks using the Bayesian algorithm of Arias et al. (2021), which fully incorporates in the credibility regions of the IRFs the uncertainty about the correlation between instruments and shocks of interest, delivers similar results to the baseline. Third, I show that the supply and demand components in the surprises can be separated using the high-frequency co-movement between surprises in oil futures and in a set of alternative stock price indices. Finally, I briefly describe several additional robustness checks, which are presented in the Appendix.

5.1 Directly controlling for information effects

The second methodology to disentangle the oil supply news shock from the information shock consists in directly controlling for the potential information asymmetry between OPEC and financial markets (Romer and Romer, 2000; Miranda-Agrippino and Ricco, 2021). If part of the variation in the surprises is due to markets learning OPEC's private assessment of oil supply or demand conditions, then orthogonalising the surprises with respect to a good measure of the OPEC's information set would separate the information component (the fitted values) from the oil supply news component (the residuals).

I obtain a measure of OPEC's information set from the OPEC Monthly Oil Market Reports (MOMR). These publications are available from January 2001 and provide forecasts for a variety of oil-related variables. The five concepts that I can consistently obtain for the 2001–2021 period are: (i) the nowcast of world and U.S. GDP and relative revisions; (ii) the 1-year backcast, nowcast, and revisions of oil demand for the world and OECD economies; (iii) the 1-year backcast, nowcast, and revisions of non-OPEC global oil supply; (iv) the nowcast, month-on-month change, and revisions of OPEC oil supply as reported by secondary sources; and finally (v) the 1-year backcast, nowcast, and revisions of the global oil demand-supply balance.[40]

The methodology consists of two steps. In the first step I regress the surprises in oil futures on the information in the MOMRs at OPEC conference frequency. The residuals of this regression correlate with the oil news shock, while the fitted part captures the information shock. The model is the following:

$$Surprise_d = \alpha + \underbrace{\sum_{j=-1}^{0} \theta_j F_d x_{y+j}}_{\text{MOMR forecasts}} + \underbrace{\sum_{j=-1}^{0} \vartheta_j \left[F_d x_{y+j} - F_{d-1} x_{y+j}\right]}_{\text{MOMR revisions}} + IV_d^{oil}, \qquad (19)$$

where the subscript d indexes the day of the announcement, $F_d x_y$ represents the MOMR

[40]The MOMRs are by no means a perfect measure of OPEC's private information at the time of the announcement. There are two main limitations. First, the MOMRs are released at specific dates that do not necessarily correspond to OPEC conference days. Consequently, they might not capture all the information available to OPEC at the time of the announcement. However, in most cases the time gap that separates the release of the MOMR from the announcement is limited to a few days. Second, the MOMRs are publicly available from 2001 and there is no clear reason why markets should not have incorporated this information into their pricing decisions before OPEC conferences. However, MOMRs were not available publicly prior to 2001 and became easily accessible via the OPEC website only recently.

forecast for variable x at yearly horizon y on the day of announcement d, and $F_d x_y - F_{d-1} x_y$ is the revision in the forecast for variable x at yearly horizon y from the MOMR release associated with the day of the last OPEC meeting. Forecasts and OPEC conference days are aligned such that the latest edition of the MOMR is always associated with the upcoming OPEC meeting.

One important prediction of models of information frictions is that agents only gradually adjust their beliefs to new information (Coibion and Gorodnichenko, 2015; Bordalo et al., 2020). Because of this, revisions of expectations might be autocorrelated and might contain information on both current and past structural shocks. This suggests an additional step to obtain a clean measure of current structural shocks. The residual and fitted components from the previous step are aggregated at monthly frequency and consequently regressed on their own lags. The model is the following:

$$\overline{Z}_m = \phi_0 + \sum_{j=1}^{12} \phi_j \overline{Z}_{m-j} + Z_m, \tag{20}$$

where $\overline{Z}_m$ is either the residual or the fitted part of Eq. (19). The residuals Z_m from this regression are either the (alternative) instrument for the oil supply new shock or the (alternative) instrument for the information shock, both at monthly frequencies, according to the dependent variables selected.[41]

I find that the MOMRs contain information that helps predict the surprises in the price of oil futures. The results of the regression in Eq. (19) are reported in Table 2. The first column reports the results for the regression that uses all 24 MOMR forecast and revision series as covariates. The null hypothesis of joint non-significance of the coefficients is rejected at the 1% confidence level. The other columns focus on the specific concepts: forecasts of GDP, forecasts of global oil demand, forecasts of non-OPEC supply, forecasts of OPEC supply, and forecasts of global demand-supply balance. In all these cases, except for column (5), the null of joint non-significance is rejected at 5% confidence level.

Although the significance of the individual coefficients is not particularly important in this context, as the objective is to maximise the fit, a few interesting results emerge.

[41] These regressions use only observations in months when there was at least one OPEC meeting.

Table 2: Information content of the Monthly Oil Market Reports

	(1)	(2)	(3)	(4)	(5)	(6)
R^2	0.380	0.070	0.132	0.148	0.021	0.035
F	91.603	7.506	5.769	4.475	1.299	8.494
p-value	0.000	0.000	0.000	0.003	0.279	0.000
N	76	76	76	76	76	76

Note: Measures of fit for the projection of daily surprises in oil futures on the OPEC Monthly Oil Market Reports forecasts and revisions. (1) projection on all MOMR forecasts and revisions; (2) only forecasts of GDP; (3) only forecasts of global oil demand; (4) only forecasts of non-OPEC supply; (5) only forecasts of OPEC supply from secondary sources; (6) only forecasts of global demand-supply balance. Table B.3 details the full set of results.

Table 3: Autoregressive component

	Residual	Fitted
R^2	0.155	0.355
F	1.456	9.621
p-value	0.170	0.000
N	68	68

Note: Measures of fit for the projection of the residual and fitted components of Equation (20) on their own lags. The lag order used is 12. Table B.4 details the full set of results.

The full set of estimates, reported in Table B.3, shows that the concepts that correlate the most with the surprises are OPEC's forecasts and revisions for world and OECD oil demand. Most coefficients on these concepts are significantly different from zero. Also the coefficients on the forecasts and revisions for world GDP are significant. Surprisingly, among the concepts relating to OPEC and non-OPEC's oil supply, only the coefficient on the nowcast revision of OPEC's supply is significant, but it is very close to zero.

The results also show that the information component of the surprises (i.e. the fitted values of Eq. 19) is autocorrelated. In other words, the variation in past surprises that is explained by the information set of OPEC helps predict current surprises. Table 3 reports the results for the regression in Eq. (20). In the autoregression for the fitted values, the null hypothesis of joint non-significance of the autoregressive coefficients is rejected at the 1% level.

The results obtained using this alternative methodology to disentangle the shocks are consistent with those obtained using the baseline identification strategy. Figure 10

Figure 10: Alternative information-robust identification

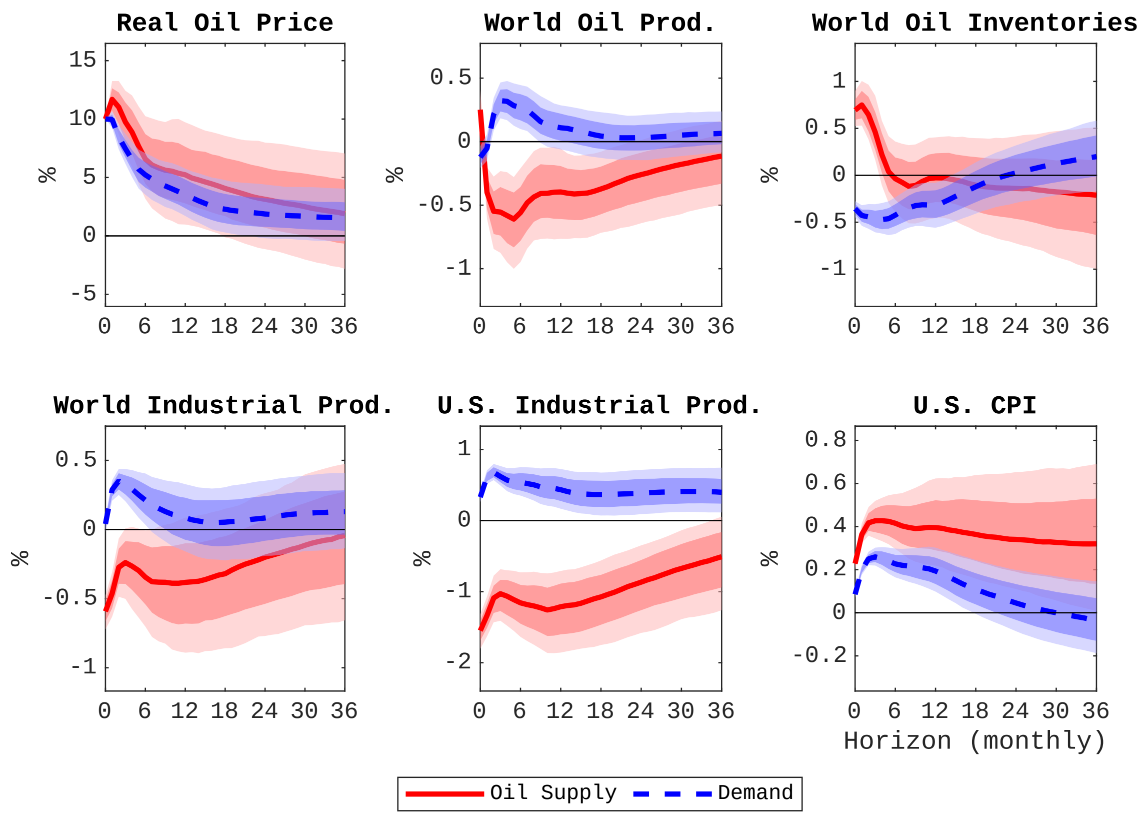

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the alternative robust proxies presented in Section 5.1. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1982:7–2023:5. IV sample: 2002:1–2021:3.

displays the impulse responses. A negative oil supply news shock contracts world and U.S. industrial production, contracts world oil production, and expands U.S. CPI. World oil production decreases, with a lag, while world oil inventories increase. This supports the baseline results and shows that both methodologies can be employed to separate the shocks in the surprises in oil futures. In Section C.5 I identify the shocks using these alternative proxies in the 6-variable VAR augmented with financial variables to avoid confounding the issues deriving from informational insufficiency with those deriving from the failure of the proxy exogeneity assumption. Figure C.8 shows that the results are confirmed.

Figure 11: Joint identification – fully Bayesian approach

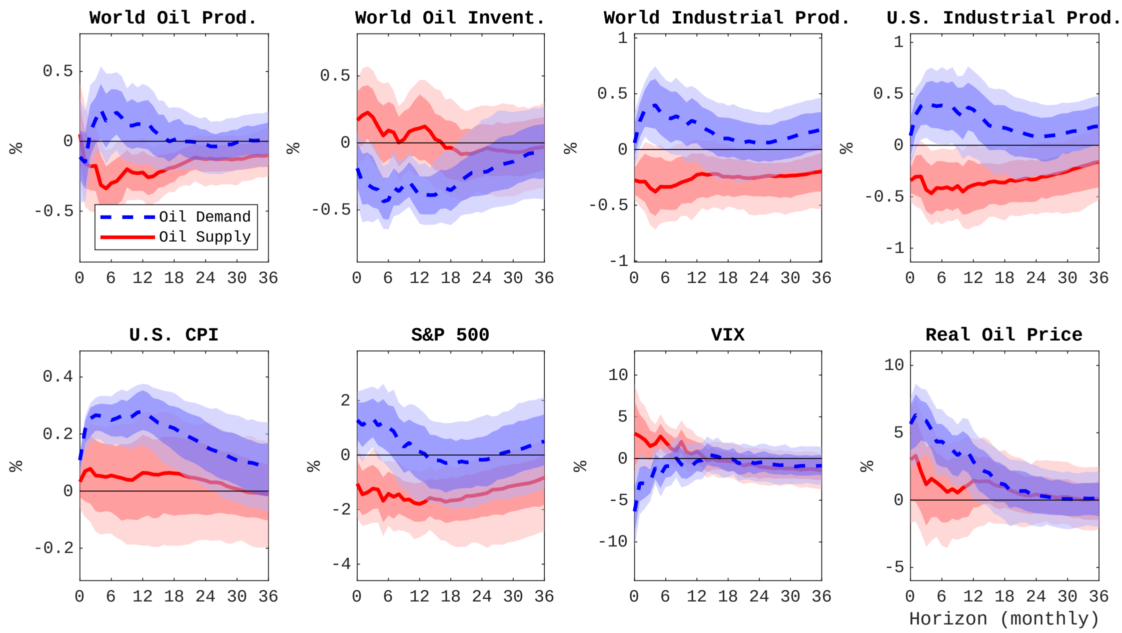

Note: Solid red: impulse responses to a one-standard-deviation oil news shock. Dashed blue: responses to a one-standard-deviation information shock. The shocks are jointly set-identified using the robust proxies and an additional set of sign restrictions on the cross-correlations of shocks and proxies, following Arias et al. (2021). BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1983:7–2023:5.

5.2 Joint identification à la Arias et al. (2021)

To provide further robustness to the identification, especially to ease potential concerns related to instrumental weakness, I jointly identify the two shocks by estimating the structural form of the 6-variable VAR augmented with the the S&P 500 and the VIX using the fully Bayesian algorithm of Arias et al. (2021). The sample covers the period from July 1983 to May 2023. The additional sign restrictions used to set-identify the shocks require that each proxy be positively correlated with its respective shock, and that this correlation be higher than the correlation between the same proxy and the other shock. The prior belief is also specified so that at least 5% of the variance in each proxy is attributed to the relevant shock. The impulse responses to both shocks are in line with the baseline results (Figure 11).

5.3 Identification based on alternative stock prices

One might worry that the results might depend on the specific stock price index used to separate the two shocks. As a test that results are robust to the use of alternative stock price indices, the exercise in Section 4.4 is repeated by identifying the shocks using robust proxies constructed based on the co-movement on OPEC conference days between daily surprises in oil futures and changes in different stock price indices. The stock price indices used are the S&P 500 (i.e. the baseline) the DS World stock price index, the DS Airlines index, the TOPIX, and the KOSPI. A quick look at Figure 12 shows that using alternative stock price indices does not alter the results.

Figure 12: Alternative stock price indices

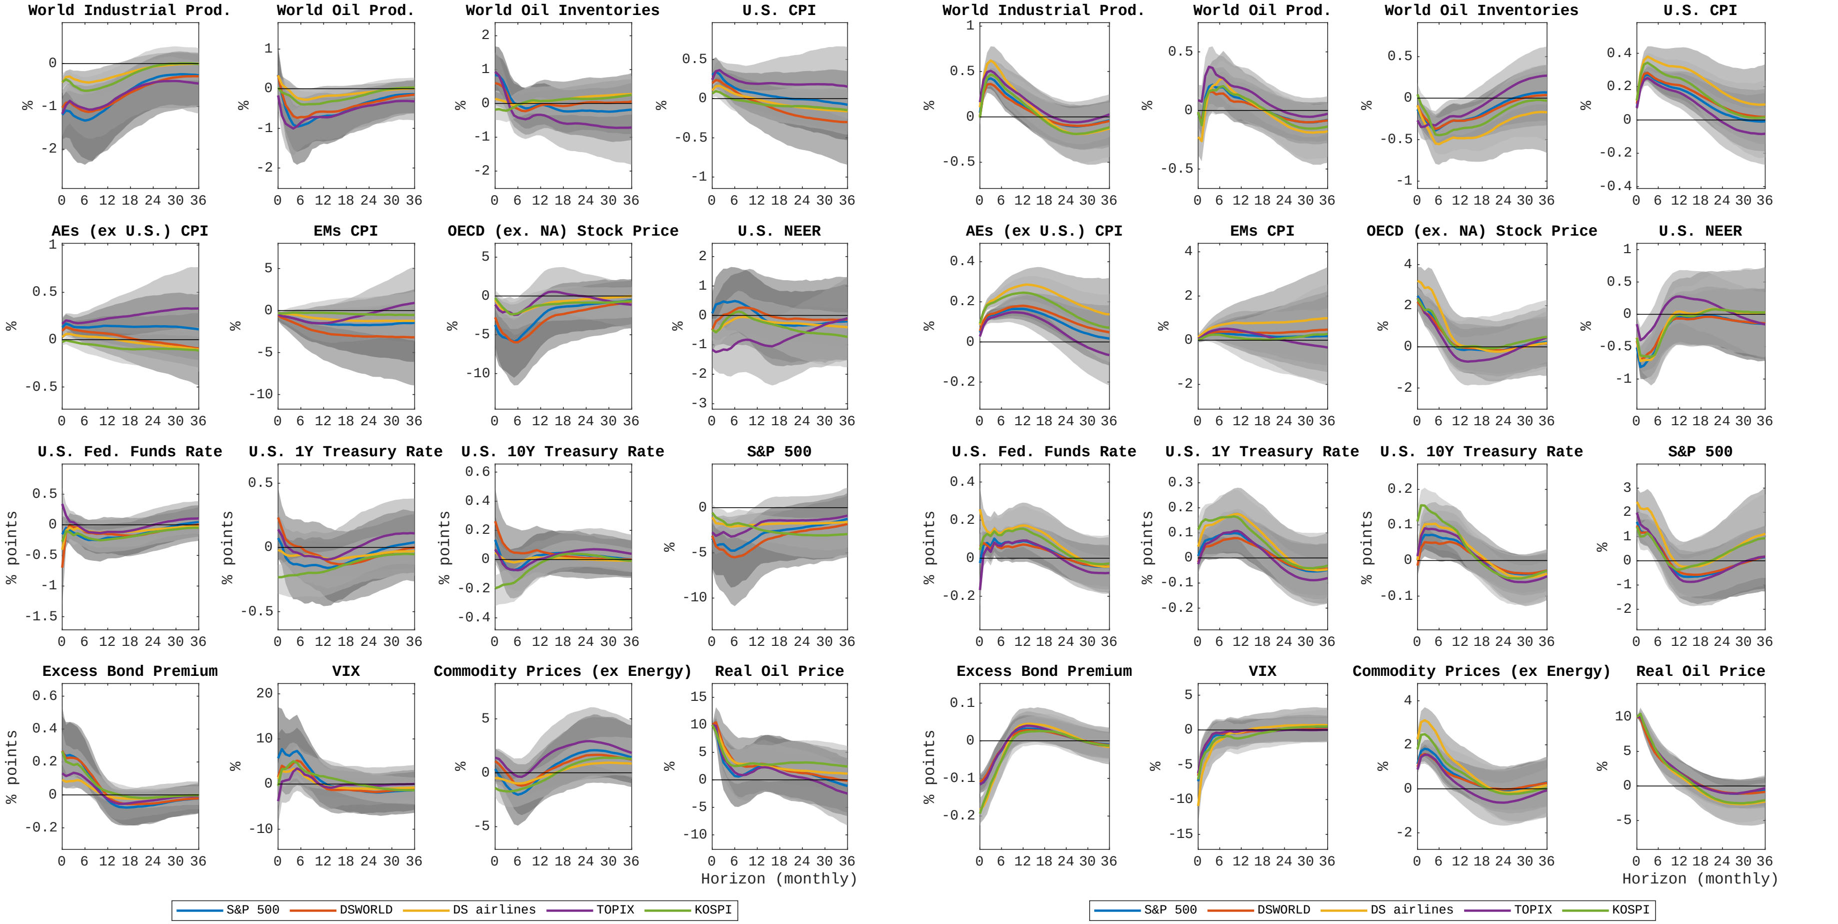

(a) Oil supply news shock

(b) Information shock

Note: Left panel: impulse responses to an oil supply news shock. Right panel: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using robust proxies, constructed based on the comovement on OPEC conference days between daily surprises in oil futures and changes in different stock price indices. **Blue**: S&P 500; **Red**: DS World; **Yellow**: DS Airlines; **Purple**: TOPIX; **Green**: KOSPI. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1980:3–2023:5. All proxies span the period 1983:7–2023:6.

5.4 Additional robustness checks

In Section C, in the Appendix, I present some additional robustness checks. The results are summarised as follows. Results obtained with the robust instruments do not suffer from sample-dependence (Figure C.2). Jointly identifying the shocks, which requires an additional restriction – in this case a recursive restriction whereby the oil supply news shock is ordered first – does not alter the results (Figure C.3). Results are robust to using an internal instrument approach, both on the full and pre-pandemic samples (Figures C.4 and C.5). Extending the 6-variables VAR with financial variables, as suggested in Mori and Peersman (2024), does not guarantee a clean identification of the shocks as the issues stemming from the endogeneity of the proxy remain (Figures C.6, C.7, and C.8). Constructing the instruments using only maturities between 1 and 4 months – that are reasonably liquid even in the early sample and are available from March 1983 – does not alter the results (Figure C.11). Similarly, constructing the proxies using only data from April 1989 onwards – as recommended in Kilian (2024) – delivers similar results, although the proxies obtained are weak (Figure C.12). Moreover, Figure C.13 shows that results are robust to reweighting the surprises as suggested in Kilian (2024). Using the co-movement of oil futures and the daily measure of U.S. real business conditions proposed by Aruoba et al. (2009) to identify the two shocks delivers similar results (Figure C.14). Removing from the sample OPEC conferences that happened during market holidays does not alter the results (Figure C.15). Removing the most influential observations from the first-stage regression also does not alter the results (Figure C.16). Results are robust to using alternative measures of world industrial production (Figure C.17). Finally, results are robust to cutting the Covid-19 period from the sample (Figure C.18) or to dealing with the Covid discontinuity using the pandemic priors of Cascaldi-Garcia (2022) (Figure C.19).

6 Conclusion

I study the macroeconomic effects of oil supply news shocks by exploiting institutional features of OPEC. My paper contributes to the literature in three distinct ways. First, I identify an issue with the predominant identification strategy in the literature. Surprises in the price of oil futures computed on a daily window around OPEC conference announcements about future production quotas capture revisions in market expectations about both oil supply and global demand. Therefore, they cannot be considered an exogenous measure of shifts in oil supply expectations and any identification based on them, by conflating shocks to oil supply and demand expectations, will return biased estimates. Indeed, dynamic responses obtained by identifying the oil supply news shock using the surprises in oil futures alone present output puzzles. An oil supply news shock that increases the price of oil, depending on the sample, appears to have an *expansionary* effect on the global economy. This contrasts with the well-established theoretical result that a negative oil supply shock has a recessionary impact on macroeconomic conditions, via an increase in marginal costs. I show – aided by a model of information frictions in the oil market – that OPEC announcements, by revealing oil supply conditions, induce imperfectly informed markets that use the price of oil as a public signal of economic conditions to revise their expectations about aggregate demand. In other words, there is an information effect of OPEC announcements – similar to the information effect of monetary policy – that leads markets to revise up (down) their forecasts of demand when OPEC reveals that oil supply conditions were looser (tighter) than expected. Importantly, for this effect to occur, it is not necessary for OPEC to be more informed than the markets about demand conditions.

Second, I provide a solution to this identification problem by exploiting the high-frequency co-movement of oil futures and stock prices in a narrow window around OPEC announcements, in line with the theoretical predictions of the model. This co-movement is informative because the oil supply news shock moves oil futures and stock prices in opposite directions, while the information shock moves both in the same direction. This additional restriction on the sign of the co-movement allows me to obtain two robust high-frequency instruments: one to identify the oil supply news shock and one for the information shock. Impulse responses for the six-variables VAR of Känzig (2021) identified with the robust instruments do not show any trace of the puzzles that are present when using as an instrument the full set of surprises.

Third, having obtained exogenous instruments, I show that oil supply news shocks have powerful effects – larger than previously documented – on a large set of macroeconomic aggregates. A negative shock has deep and long-lasting contractionary effects on the global economy. World industrial production and U.S. industrial production contract by roughly 1% following an oil supply news shock normalised to increase real oil price by 10%. CPI increases by 0.3%. World oil production also contracts by 0.5%, while oil inventories expand by 0.5%. The shock also contracts OECD (excluding North America) and U.S. stock prices by 5%, widens credit spreads in the U.S. (as measured by the excess bond premium), and causes an increase in uncertainty on financial markets (as measured by the VIX). The stagflationary nature of the shock puts central banks around the world in the difficult position of having to choose between price and output stabilisation. The response of the policy rate of the median advanced economy suggests that central banks face a challenging trade-off between even higher inflation or an even deeper recession.

Conversely, there is no such trade-off in the case of a positive information shock, which has a positive effect on both industrial production and prices, and transmits as a demand shock. In this instance the direction of monetary policy is unambiguous. A positive information shock has a relatively short-lived and mild expansionary effect. Following such a shock, real activity, CPI, world oil production, and equity prices expand sharply, while credit spreads narrow and uncertainty subsides. An information shock normalised to increase the real oil price by 10% causes a 0.3% expansion in both world and U.S. industrial production. U.S. CPI increases persistently, with a peak response of roughly +0.2% 3 months after the shock. World oil production increases while world oil inventories get depleted.

References

Alquist, Ron and Lutz Kilian, "What do we learn from the price of crude oil futures?," *Journal of Applied Econometrics*, jun 2010, *25* (4), 539–573.

Alsalman, Zeina, Ana María Herrera, and Sandeep Kumar Rangaraju, "Oil news shocks and the US stock market," *Energy Economics*, 2023, *126*, 106891.

Angeletos, George-Marios and Jennifer La'O, "Sentiments," *Econometrica*, 2013, *81* (2), 739–779.

Anzuini, Alessio, Patrizio Pagano, and Massimiliano Pisani, "Macroeconomic Effects of Precautionary Demand for Oil," *Journal of Applied Econometrics*, sep 2015, *30*, 968–986.

Arias, Jonas E., Juan F. Rubio-Ramírez, and Daniel F. Waggoner, "Inference in Bayesian Proxy-SVARs," *Journal of Econometrics*, 2021.

Arias, Jonas E, Juan F Rubio-Ramírez, and Daniel F Waggoner, "Uniform priors for impulse responses," *Econometrica*, 2025, *93* (2), 695–718.

Aruoba, S. Borağan, Francis X. Diebold, and Chiara Scotti, "Real-time measurement of business conditions," *Journal of Business and Economic Statistics*, 2009, *27* (4), 417–427.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin, "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, jan 2010, *25* (1), 71–92.

Basak, Suleyman and Anna Pavlova, "A Model of Financialization of Commodities," *Journal of Finance*, 2016, *71* (4), 1511–1556.

Baumeister, Christiane, "Measuring market expectations," *Handbook of economic expectations*, 2023, pp. 413–441.

_ **and Gert Peersman**, "The role of time-varying price elasticities in accounting for volatility changes in the crude oil market," *Journal of applied econometrics*, 2013, *28* (7), 1087–1109.

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

_ **and Lutz Kilian**, "A general approach to recovering market expectations from futures prices with an application to crude oil," 2014.

_ **and** _ , "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 S., Mark Gertler, and Mark Watson, "Systematic monetary policy and the effects of oil price shocks," *Brookings Papers on Economic Activity*, 1997, *1997* (1), 91–157.

Bjørnland, Hilde C., Vegard H. Larsen, and Junior Maih, "Oil and macroeconomic (in)stability," *American Economic Journal: Macroeconomics*, 2018, *10* (4), 128–151.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, "Overreaction in macroeconomic expectations," *American Economic Review*, 2020, *110* (9), 2748–2782.

Brunetti, Celso, Marc Joëts, and Valérie Mignon, "Reasons Behind Words: OPEC Narratives and the Oil Market," 2023.

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

— **and —** , “Comparing external and internal instruments for vector autoregressions,” *Journal of Economic Dynamics and Control*, 2025, p. 105131.

Caldara, Dario and Edward Herbst, “Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs,” *American Economic Journal: Macroeconomics*, 2019, *11* (1), 157–192.

— , Michele Cavallo, and Matteo Iacoviello, “Oil price elasticities and oil price fluctuations,” *Journal of Monetary Economics*, 2019, *103*, 1–20.

Cascaldi-Garcia, Danilo, “Pandemic priors,” *International Finance Discussion Paper*, 2022, (1352).

Ciccarelli, Matteo and Fabio Canova, “Estimating Multicountry VAR Models,” *International Economic Review*, 2009, *50* (3), 929–959.

Cieslak, Anna and Andreas Schrimpf, “Non-monetary news in central bank communication,” *Journal of International Economics*, 2019, *118*, 293–315.

— and Hao Pang, “Common Shocks in Stocks and Bonds,” *SSRN Electronic Journal*, 2020, (June 2019).

Coibion, Olivier and Yuriy Gorodnichenko, “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 2015, *105* (8), 2644–2678.

Darby, Michael R., “The Price of Oil and World Inflation and Recession,” *The American Economic Review*, 1982, *72* (4), 738–751.

Degasperi, Riccardo, Seokki Simon Hong, and Giovanni Ricco, “The Global Transmission of U.S. Monetary Policy,” 2021.

Demirer, Riza and Ali M. Kutan, “The behavior of crude oil spot and futures prices around OPEC and SPR announcements: An event study perspective,” *Energy Economics*, 2010, *32* (6), 1467–1476.

Draper, Dennis W., “The behavior of event-related returns on oil futures contracts,” *Journal of Futures Markets*, 1984, *4* (2), 125–132.

Forni, Mario, Alessandro Franconi, Luca Gambetti, and Luca Sala, “Asymmetric transmission of oil supply news,” *Quantitative Economics*, 2025, *16* (3), 947–979.

— and Luca Gambetti, “Sufficient information in structural VARs,” *Journal of Monetary Economics*, 2014, *66*, 124–136.

Gambetti, Luca and Laura Moretti, “News, Noise and Oil Price Swings,” 2017.

Gazzani, Andrea, Fabrizio Venditti, and Giovanni Veronese, “Oil price shocks in real time,” *Journal of Monetary Economics*, 2024, *144*, 103547.

Gelain, Paolo and Marco Lorusso, “Oil price fluctuations, US banks, and macroprudential policy,” *Journal of Monetary Economics*, 2025, *156*, 103848.

Giacomini, R, T Kitagawa, and M Read, “Robust Bayesian Inference in Proxy SVARs,” *Journal of Econometrics*, 2021.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri, “Prior Selection for Vector Autoregressions,” *Review of Economics and Statistics*, may 2015, *97* (2), 436–451.

Gilchrist, Simon and Egon Zakrajšek, “Credit spreads and business cycle fluctuations,” *American Economic Review*, 2012, *102* (4), 1692–1720.

Hamilton, James D., "Oil and the macroeconomy since world war II," *Journal of Political Economy*, 1983, *91* (2), 228–248.

Hamilton, James D and Jing Cynthia Wu, "Risk premia in crude oil futures prices," *Journal of International Money and Finance*, 2014, *42*, 9–37.

Herrera, Ana María and Sandeep Kumar Rangaraju, "The Time-Varying Effects of Oil News Shocks," 2025.

Hu, Conghui and Wei Xiong, "The informational role of commodity futures prices," *NBER Working Paper*, 2013, (May), 1–42.

IEA, "Oil 2021," *International Energy Agency*, 2021.

_ , "Oil 2024 - Analysis and forecast to 2030," *International Energy Agency*, 2024.

Jarociński, Marek and Peter Karadi, "Deconstructing monetary policy surprises-The role of information shocks," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 1–43.

Juvenal, Luciana and Ivan Petrella, "Speculation in the Oil Market," *Journal of Applied Econometrics*, jun 2015, *30* (4), 621–649.

Kadiyala, K . Rao and Sune Karlsson, "Numerical Methods for Estimation and Inference in Bayesian VAR-Models," *Journal of Applied Econometrics*, 1997, *12* (2), 99–132.

Känzig, Diego R., "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, apr 2021, *111* (4), 1092–1125.

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

_ , "How to construct monthly VAR proxies based on daily surprises in futures markets," *Journal of Economic Dynamics and Control*, 2024, *168*, 104966.

_ **and Daniel P. Murphy**, "The Role Of Inventories And Speculative Trading In The Global Market For Crude Oil," *Journal of Applied Econometrics*, apr 2014, *29* (3), 454–478.

_ **and Thomas K. Lee**, "Quantifying the speculative component in the real price of oil: The role of global oil inventories," *Journal of International Money and Finance*, 2014, *42*, 71–87.

_ **and Xiaoqing Zhou**, "The Econometrics of Oil Market VAR Models," 04 2023.

Lee, Yeon Jik and Tatevik Sekhposyan, "The relevance of temporal aggregation for the propagation of macroeconomic shocks," 2024.

Lin, Sharon Xiaowen and Michael Tamvakis, "OPEC announcements and their effects on crude oil prices," *Energy Policy*, 2010, *38* (2), 1010–1016.

Litterman, Robert B., "Forecasting with bayesian vector autoregressions—five years of experience," *Journal of Business and Economic Statistics*, 1986, *4* (1), 25–38.

Loutia, Amine, Constantin Mellios, and Kostas Andriosopoulos, "Do OPEC announcements influence oil prices?," *Energy Policy*, 2016, *90*, 262–272.

Lucas, Robert E. Jr., "Expectations and the neutrality of money," *Journal of Economic Theory*, 1972, *4* (2), 103–124.

Melosi, Leonardo, "Signalling Effects of Monetary Policy," *The Review of Economic Studies*, 2017, *84*, 853–884.

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

Miranda-Agrippino, Silvia and Giovanni Ricco, "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 2021.

_ and _ , "Identification with external instruments in structural VARs," *Journal of Monetary Economics*, 2023, *135*, 1–19.

_ and Hélène Rey, "U.S. Monetary Policy and the Global Financial Cycle," *The Review of Economic Studies*, 2020, *87* (6), 2754–2776.

_ and _ , "The Global Financial Cycle," *CEPR Discussion Paper*, 2021.

Montiel Olea, José Luis, James H. Stock, and Mark W. Watson, "Inference in Structural Vector Autoregressions identified with an external instrument," *Journal of Econometrics*, 2020.

Mori, Lorenzo and Gert Peersman, "Estimating the Macroeconomic Effects of Oil Supply News," 2024.

Moussa, Zakaria and Arthur Thomas, "Identifying oil supply news shocks and their effects on the global oil market," 2023.

Nakamura, Emi and Jón Steinsson, "High-frequency identification of monetary non-neutrality: The information effect," *Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

OPEC, "OPEC Annual Statistical Bulletin 2023," *Organization of the Petroleum Exporting Countries*, 2023, *58.*

Pescatori, Andrea and Yousef Nazer, "OPEC and the Oil Market," *IMF Working Papers*, 9 2022, *2022* (183), 1.

Romer, Christina D. and David H. Romer, "Federal Reserve Information and the Behavior of Interest Rates," *The American Economic Review*, 2000, *90* (3), 429–457.

Schmidbauer, Harald and Angi Rösch, "OPEC news announcements: Effects on oil price expectation and volatility," *Energy Economics*, 2012, *34* (5), 1656–1663.

Sims, Christopher A. and Tao Zha, "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 1998, *39* (4), 949–968.

Sockin, Michael and Wei Xiong, "Informational Frictions and Commodity Markets," *Journal of Finance*, 2015, *70* (5), 2063–2098.

Stock, James H. and Mark W. Watson, "Disentangling the Channels of the 2007–09 Recession," *Brookings Papers on Economic Activity*, 2012, (Spring).

_ and _ , "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, may 2018, *128* (610), 917–948.

_ , Jonathan H. Wright, and Motohiro Yogo, "A survey of weak instruments and weak identification in generalized method of moments," *Journal of Business and Economic Statistics*, 2002, *20* (4), 518–529.

Wu, Tao and Michele Cavallo, "Measuring Oil-Price Shocks Using Market-Based Information," *IMF Working Papers*, 2012, *12* (19), 1.

Appendices

A A simplified graphical description of the model

Figure A.1: Price and quantity revisions following OPEC announcements

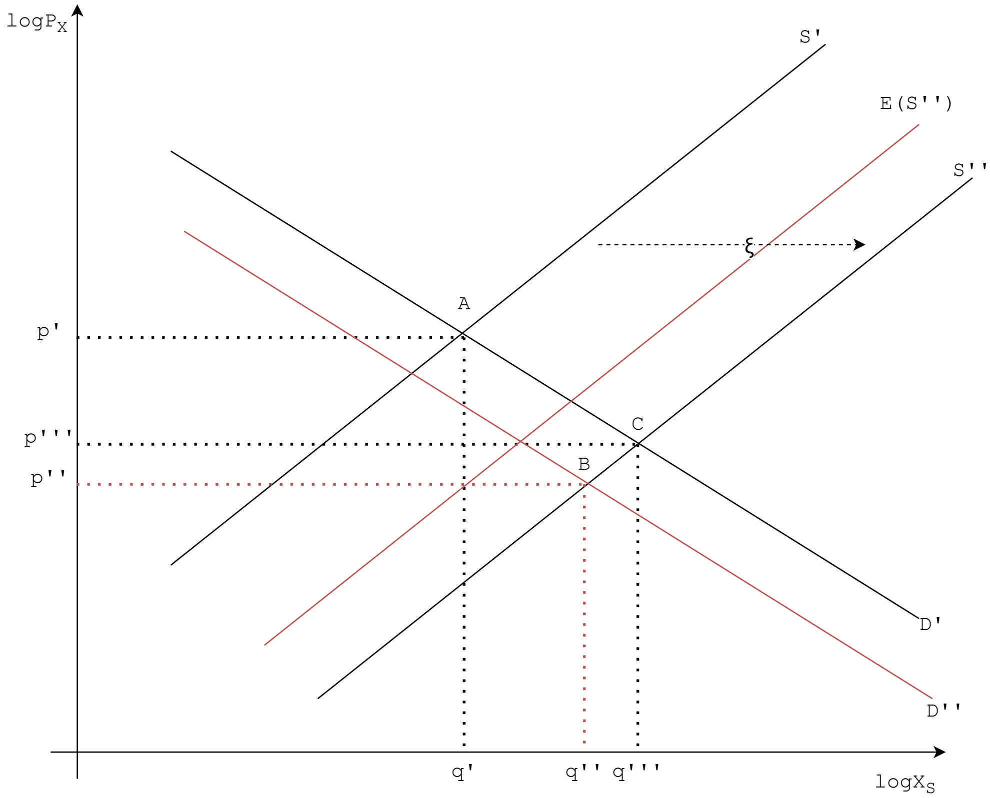

This subsection provides a graphical intuition for the model in Section 2. The equilibria under the full information benchmark and with information frictions are displayed in Figure A.1. Starting from an equilibrium defined by the intersection of the oil supply curve S' and oil demand curve D' (point A), the oil supply shifter ξ moves the oil supply curve from S' to S''. Agents, not knowing for certain whether movements in oil price are due to a favorable supply shock or to weaker economic conditions, attribute part of the decrease in oil price to a drop in demand. They expect oil supply to shift to $\mathbb{E}(S'')$, while the demand for oil shifts from D' to D''. The new equilibrium is reached where D'' intersects the actual supply curve S'' (point B). Observe that these price and quantity (p'' and q'', loosely corresponding to Eq. (6) and (8) in the main text) are both lower than the price and quantity that would emerge under the full information benchmark (p''' and q''', corresponding to Eq. (9) and (10) in the main text). When the OPEC announcement reveals ξ, the equilibrium moves from B to C, determining a positive revision in the price and a positive reassessment of economic conditions, which leads to a positive revision in the quantity demanded. This generates a positive correlation between the surprises in prices and economic conditions, which renders the surprises endogenous.

B Additional tables

Table B.1: Sample coverage for country exercises

Advanced Economies			Emerging Markets		
Australia	1986:11	2018:8	Brazil	1994:7	2018:8
Austria	1984:1	2018:8	Chile	1995:5	2018:6
Belgium	1984:1	2018:8	China	1994:5	2018:8
Canada	1984:1	2018:8	Colombia	1995:4	2018:8
Denmark	1984:1	2018:8	Czech Republic	1995:12	2018:8
Finland	1988:3	2018:8	Hungary	1991:6	2018:8
France	1984:1	2018:8	India	1990:1	2018:4
Germany	1984:1	2018:8	Malaysia	1995:11	2017:12
Italy	1984:1	2018:8	Mexico	1998:11	2018:2
Japan	1984:1	2018:8	Philippines	1996:1	2018:7
Netherlands	1985:6	2018:8	Poland	1994:3	2018:8
Norway	1984:1	2018:8	Russia	1998:1	2018:8
Spain	1987:3	2018:8	South Africa	1990:1	2018:8
Sweden	1984:1	2018:8	Thailand	1999:1	2018:7
UK	1984:1	2018:8	Turkey	1990:1	2018:8

Table B.2: Variables used, sources, and transformations

Variable	Description	Source	Codes	Logs	RW	(1)	(2)	(3)	(4)	(5)	(6)
Real Oil Price	Spot Crude Oil Price: WTI, $/bbl, Monthly, NSA. End-of-month from 1986:1. Deflated by U.S. CPI.	FRED	WTISPLC; DCOILWTICO; CPIAUCSL	•	•	•	•	•	•	•	•
World Oil Production	Crude Oil Production, World, Mbbl/day – Turnover by volume	Datastream	EIA1955	•	•	•	•	•	•	•	•
World Oil Inventories		Datastream, Kilian and Murphy (2014)	EIA1976; EIA1533; EIA1541	•	•	•	•	•	•	•	•
World Industrial Production	Industrial production of OECD + 6 Major Emerging Markets (Brazil, China, India, Indonesia, Russia and South Africa)	Baumeister and Hamilton (2019)		•	•	•	•	•	•	•	•
U.S. Industrial Production	Industrial Production: Total Index, 2012=100, Monthly, SA	FRED	INDPRO	•	•	•				•	•
U.S. CPI	CPI for All Urban Consumers: All Items in U.S. City Average, 1982-1984=100, Monthly, SA	FRED	CPIAUCSL	•	•	•				•	•
AEs (ex. U.S.) CPI	Headline CPI, nominal GDP weights, 2005 = 100	Dallas Fed, Global Economic Indicators		•	•		•				
EMs CPI	Headline CPI, nominal GDP weights, 2005 = 100	Dallas Fed, Global Economic Indicators		•	•		•				
OECD Stock Price	DEVD.MKTS.EX-NA-DS Market - PRICE INDEX. 01/01/1973 = 100. End-of-month	Datastream	TOTMKEF	•	•		•				
U.S. Nom. Eff. Exch. Rate	U.S. Effective Exchange Rate - Nominal - Narrow (27 economies)	BIS		•	•		•				
U.S. Federal Funds Rate	Effective Federal Funds Rate, Percent, Monthly, NSA	FRED	DFF				•				
U.S. 1Y Treasury Rate	1-Year Treasury Constant Maturity Rate, Percent, Monthly, NSA	FRED	DSG1				•	•	•		
U.S. 10Y Treasury Rate	10-Year Treasury Constant Maturity Rate, Percent, Monthly, NSA	FRED	DSG10				•				
S&P 500	S&P 500 COMPOSITE - PRICE INDEX. End-of-month	Datastream	S&PCOMP	•	•		•			•	•
Excess Bond Premium		Gilchrist and Zakrajšek (2012)					•				
VIX	Before 1986:1 VXO as reconstructed in Bloom (2009), monthly avg. From 1986:1 to 1990:11 VXO from FRED, end-of-month, close. From 1990:1 VIX from FRED, end-of-month, close.	Bloom (2009) before 1986:1; FRED	VXOCLS, VIXCLS	•	•		•	•	•	•	•
Commodity Prices (ex Energy)	World Bank Commodity Price Data, Non-energy index	World Bank Pink Sheet		•	•		•				
5Y5Y Fwd Infl. Exp. Rate	5-Year, 5-Year Forward Inflation Expectation Rate, Percent, Monthly, NSA	St. Louis Fed	T5YIFR							•	
5Y Exp. Inflation	5-Year Expected Inflation, Percent, Monthly, NSA	Cleveland Fed	EXPINF5YR								•
1Y Exp. Inflation	1-Year Expected Inflation, Percent, Monthly, NSA	Cleveland Fed	EXPINF1YR								•
1Y UoM Infl. Exp.	University of Michigan: Inflation Expectation, Percent, Monthly, NSA	University of Michigan	MICH								•
Industrial Production		OECD		•	•			•	•		
CPI		OECD		•	•			•	•		
Core CPI		OECD		•	•			•			
Stock Price Index		Datastream		•	•			•	•		
Exchange Rate		BIS		•	•			•	•		
Policy Rate		BIS						•	•		

Note: The table lists all variable used in the different models. The models are: (1) the six-variable VAR; (2) the 16-variable VAR for the global economy; (3) the models used to estimate the effects on the median advanced economy; (4) the models for the median emerging economy; (5) the model to estimate the response of the 5Y5Y forward inflation rate; (6) the model to estimate the response of the other three inflation expectations measures. The bottom section of the table lists the variables that have been collected for each of the 30 countries in the sample. *Logs* indicates logarithmic transformations. *RW* indicates assignment of a random walk prior vis-à-vis a white noise prior.

Table B.3: Information content of the Monthly Oil Market Reports

	(1)	(2)	(3)	(4)	(5)	(6)
GDP_{world}	-0.845	-0.601*				
	(-0.97)	(-1.68)				
GDP_{US}	0.251	0.345				
	(0.27)	(0.87)				
GDP^{rev}_{world}	3.995*	2.009*				
	(1.96)	(1.67)				
GDP^{rev}_{US}	-5.001**	-1.816				
	(-2.05)	(-1.26)				
Oil $\text{Demand}^{backcast}_{world}$	0.955		1.176*			
	(0.77)		(1.97)			
Oil $\text{Demand}^{nowcast}_{world}$	-0.995		-1.191**			
	(-0.71)		(-2.05)			
Oil $\text{Demand}^{backcast}_{OECD}$	-3.954*		-2.038*			
	(-1.93)		(-1.93)			
Oil $\text{Demand}^{nowcast}_{OECD}$	4.078*		1.935*			
	(1.94)		(1.70)			
Oil $\text{Demand}^{backcast,\,rev}_{world}$	-3.738		-2.932*			
	(-1.32)		(-1.89)			
Oil $\text{Demand}^{nowcast,\,rev}_{world}$	2.248		2.335**			
	(0.79)		(2.06)			
Oil $\text{Demand}^{backcast,\,rev}_{OECD}$	12.834*		6.014*			
	(1.71)		(1.86)			
Oil $\text{Demand}^{nowcast,\,rev}_{OECD}$	-3.205		-2.082			
	(-0.93)		(-1.53)			
Non-OPEC $\text{Supply}^{backcast}$	1.108			0.685		
	(0.94)			(1.36)		
Non-OPEC $\text{Supply}^{nowcast}$	-1.231			-0.731		
	(-1.05)			(-1.55)		
Non-OPEC $\text{Supply}^{backcast,\,rev}$	2.549			0.250		
	(1.26)			(0.20)		
Non-OPEC $\text{Supply}^{nowcast,\,rev}$	-0.264			1.607		
	(-0.16)			(1.58)		
OPEC $\text{Supply}^{nowcast}$	0.000				-0.000	
	(0.28)				(-0.76)	
OPEC Supply^{mom}	0.000				0.000	
	(0.13)				(0.42)	
OPEC $\text{Supply}^{nowcast,\,rev}$	0.000				0.000**	
	(1.65)				(2.21)	
OPEC $\text{Supply}^{mom,\,rev}$	0.000				0.000	
	(0.68)				(0.52)	
D/S $\text{Balance}^{backcast}$	0.989					-0.122
	(0.97)					(-0.48)
D/S $\text{Balance}^{nowcast}$	-0.711					-0.041
	(-1.25)					(-0.19)
D/S $\text{Balance}^{backcast,\,rev}$	-5.030					-0.153
	(-1.04)					(-0.43)
D/S $\text{Balance}^{nowcast,\,rev}$	0.145					0.253
	(0.17)					(0.51)
constant	-4.952	1.190**	7.242	2.868	3.266	4.692
	(-0.26)	(2.46)	(0.47)	(0.96)	(0.78)	(1.30)
R^2	0.380	0.070	0.132	0.148	0.021	0.035
F	91.603	7.506	5.769	4.475	1.299	8.494
p-value	0.000	0.000	0.000	0.003	0.279	0.000
N	76	76	76	76	76	76

Note: Projection of daily surprises in oil futures on the OPEC Monthly Oil Market Reports forecasts and revisions. (1) projection on all MOMR forecasts and revisions; (2) only forecasts of GDP; (3) only forecasts of global oil demand; (4) only forecasts of non-OPEC supply; (5) only forecasts of OPEC supply from secondary sources; (6) only forecasts of global demand-supply balance. Robust t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table B.4: Autoregressive component

	Residual	Fitted
$l = 1$	-0.395***	-0.164
	(-3.12)	(-1.60)
$l = 2$	0.058	0.962***
	(0.25)	(3.09)
$l = 3$	-0.075	0.374***
	(-0.43)	(3.87)
$l = 4$	-0.201	-0.027
	(-1.41)	(-0.15)
$l = 5$	-0.360	0.174
	(-1.60)	(0.70)
$l = 6$	-0.336*	-0.084
	(-1.91)	(-0.44)
$l = 7$	-0.298	0.283**
	(-1.58)	(2.04)
$l = 8$	0.039	0.326
	(0.25)	(0.90)
$l = 9$	-0.002	-0.077
	(-0.01)	(-0.78)
$l = 10$	-0.044	-0.536***
	(-0.41)	(-5.43)
$l = 11$	-0.230	0.402**
	(-0.47)	(2.32)
$l = 12$	-0.267	0.078
	(-1.56)	(0.55)
constant	0.198	-0.274
	(0.66)	(-1.26)
R^2	0.155	0.355
F	1.456	9.621
p-value	0.170	0.000
N	68	68

Note: Projection of the residual and fitted components of Equation (20) on their own lags. Robust t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

C Additional robustness checks

C.1 Evidence from forecast revisions around announcements

To lend more credibility to the information channel put forward in this paper, I look at how professional forecasters revise their GDP growth and inflation forecasts in a narrow interval around OPEC announcements. I show that, consistently with the information effect described in the paper, following several OPEC announcements, professional forecasters revise their forecasts of GDP growth and inflation in the same direction as the surprise in the price of oil futures.

Continuous Consensus Forecast (CCF), provided by Consensus Economics Inc., collects daily forecasts for GDP growth and inflation for several countries since 2011. The forecasts used in this exercise is a moving average, calculated each business day, of the latest 8+ qualified changed forecasts. This forecast concept, reflecting the consensus among forecasters, gives more weight to the forecasters that revised their forecast. I use the forecasts of GDP growth and inflation for the same year of the corresponding OPEC conference. Between 2011 and 2023 there are 48 OPEC announcements to which I can match forecast revisions.

Figure C.1 reports the results for the U.S. obtained using two alternative windows to compute the forecast revisions. Each panel reports the scatter of surprises in oil futures against GDP forecast revisions (blue crosses) and, if the revision in growth forecasts corresponds to a revision of the same sign in the inflation forecast – consistently with a demand shock – an orange circle is overlaid to the blue cross. In the left panel, the forecast revision is computed from the day before the OPEC announcement to the first business day successive to the announcement for which there is a revision in the forecast of GDP growth. For most observations this ranges from one to two days after the announcement, but for a few observations the window extends to 7 business days after the announcement. In the right panel the forecast revision is always computed between the day prior to and the day after the announcement. Under this more conservative definition of the event window, for several observations, the forecasts remain unrevised.

The mass of points in the first quadrant of the two panels indicate that there are several instances in which forecasters revised upwards their expectations of GDP growth in the U.S. following an OPEC announcement that led to an increase in the price of oil futures. For many of these, the upward revision in growth corresponded to an upward revision in inflation forecasts, consistently with a revision in demand conditions. Under the broader definition of the event window in the left panel, there are 28 OPEC events out of 48 in which forecasters revise their expectations of GDP growth in the same direction as the surprise in oil futures. For 16 of these there is a revision in inflation forecasts of the same sign. When using the more conservative window, forecasters revise their expectations of GDP growth in the same direction as the surprise in oil futures for 20 OPEC events. For 12 of these the revision in inflation has the same size.

Figure C.1: U.S. GROWTH FORECASTS REVISIONS AROUND OPEC ANNOUNCEMENTS

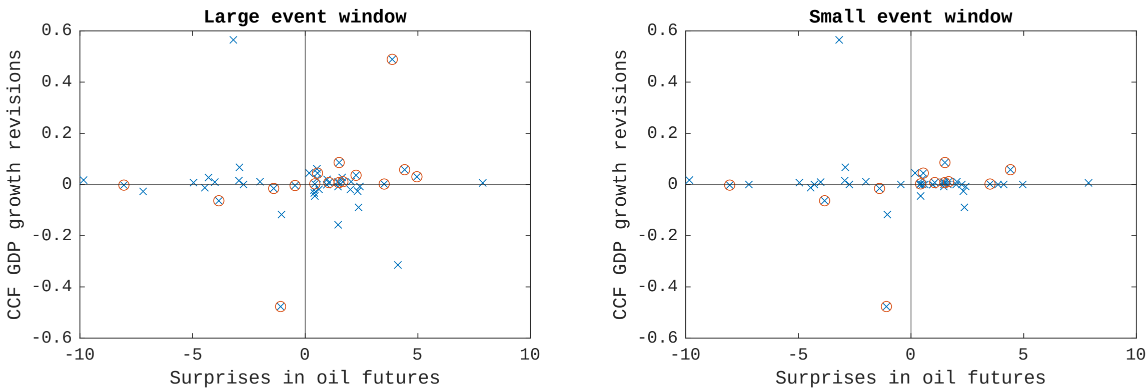

(a) Extended window

(b) Small window

Note: Scatterplots of surprises in oil futures against U.S. GDP growth forecast revisions for the current year around OPEC announcements (blue crosses). Orange circles indicate instances in which the revisions in growth and inflation forecasts have the same sign. Left panel: forecast revisions computed between the day before and the first business day successive to the OPEC announcement for which there is a revision in the CCF moving-average forecast of GDP growth. Right panel: forecast revision computed between the day before and the day after the announcement. OPEC announcements considered: from 08/06/2011 to 04/06/2023. Daily forecasts are sourced from Consensus Economics Inc., Continuous Consensus Forecasts.

C.2 Robust identification on different sample lengths

Figure C.2 shows that responses to the shocks identified with the robust instruments are consistent across samples and show no trace of puzzles in economic activity. In the main text, only the responses for the sample 1982:7–2023:5 are reported. Here it is shown that the responses obtained on the samples 1975:1–2023:5 and 1990:1–2023:5 are consistent to the baseline results.

Figure C.2: Robust identification on alternative samples

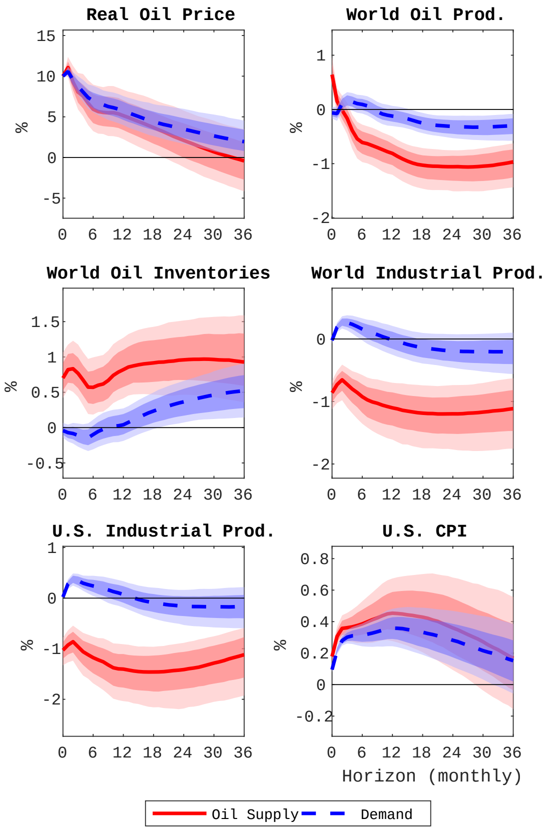

(a) Sample: 1975:1–2023:5

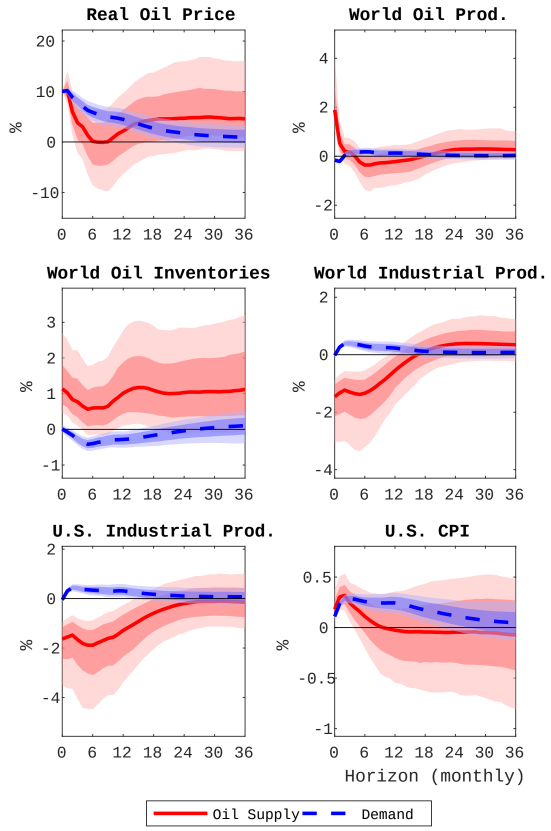

(b) Sample: 1990:1–2023:5

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Samples: 1975:1–2023:5 (left); 1990:1–2023:5 (right). Both proxies span the period 1983:7–2023:6.

C.3 Joint identification à la Mertens and Ravn (2013)

Here I identify the shocks jointly, using the two proxies. Unless one is willing to impose an assumption of zero cross-correlation between the proxies and the structural shocks, identification requires an additional restriction. I follow Mertens and Ravn (2013) by imposing a recursive structure to the shocks, ordering the oil shock first. The impulse responses to both shocks are in line with our baseline results in terms of sign, magnitude, and dynamics (Figure C.3).

Figure C.3: Joint identification à la Mertens and Ravn (2013)

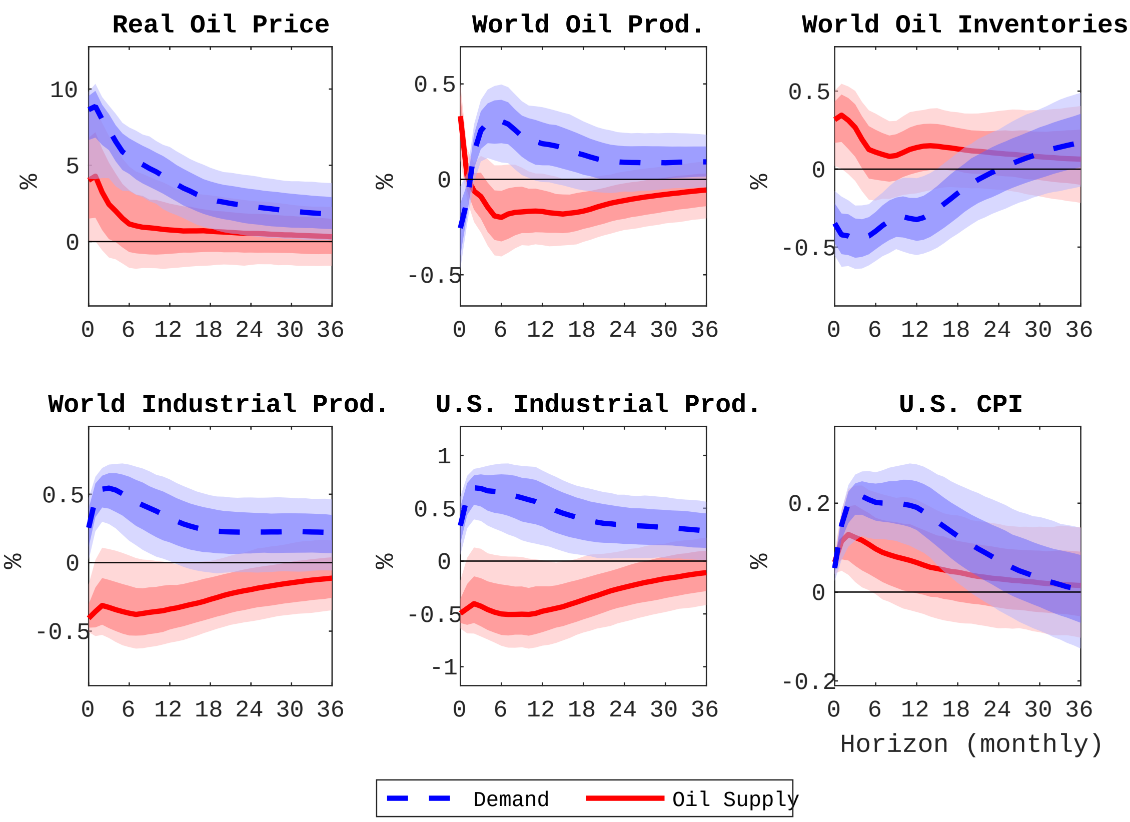

Note: Solid red: impulse responses to a one-standard-deviation oil supply news shock. Dashed blue: responses to a one-standard-deviation information shock. The shocks are jointly identified using the robust proxies and an additional recursive restriction. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1982:7–2023:5.

C.4 Identification by internal instruments

Adopting an internal Proxy-SVAR approach – whereby the instrument is included as the first of the endogenous variables in the VAR, and the covariance matrix is orthogonalized using

Figure C.4: Identification by internal instruments – full sample

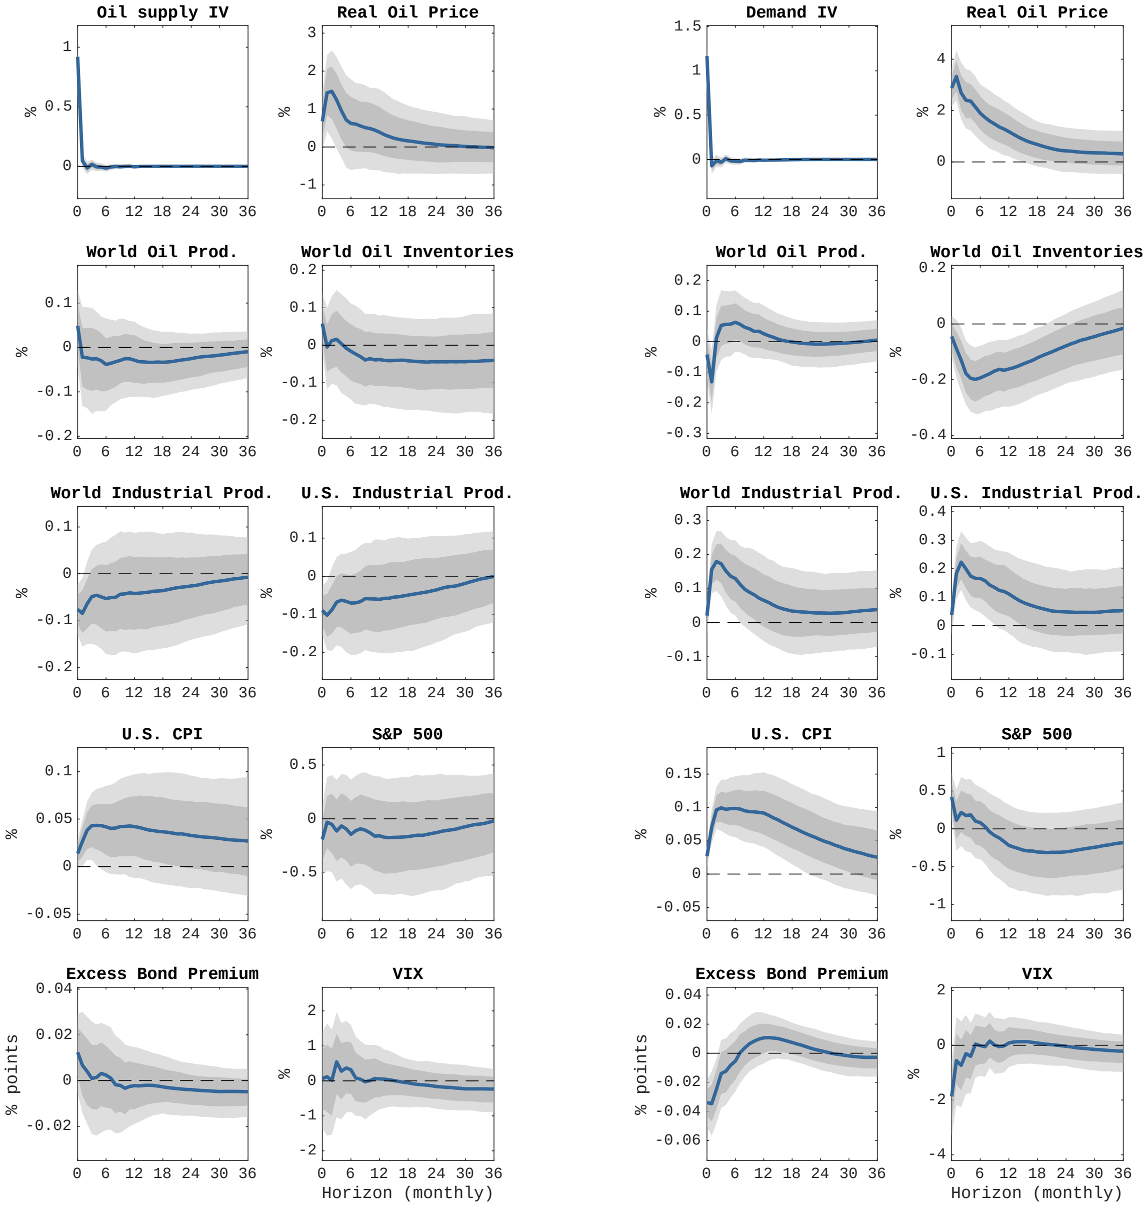

(a) Oil supply news shock

(b) Information shock

Note: Left: impulse responses to a one-standard-deviation oil supply news shock, identified using the robust instrument for oil news shocks as an internal instrument. Right: responses to a one-standard-deviation information shock, identified using the information component of the surprises as an internal instrument. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1983:7–2023:5.

Figure C.5: Identification by internal instruments – pre-pandemic

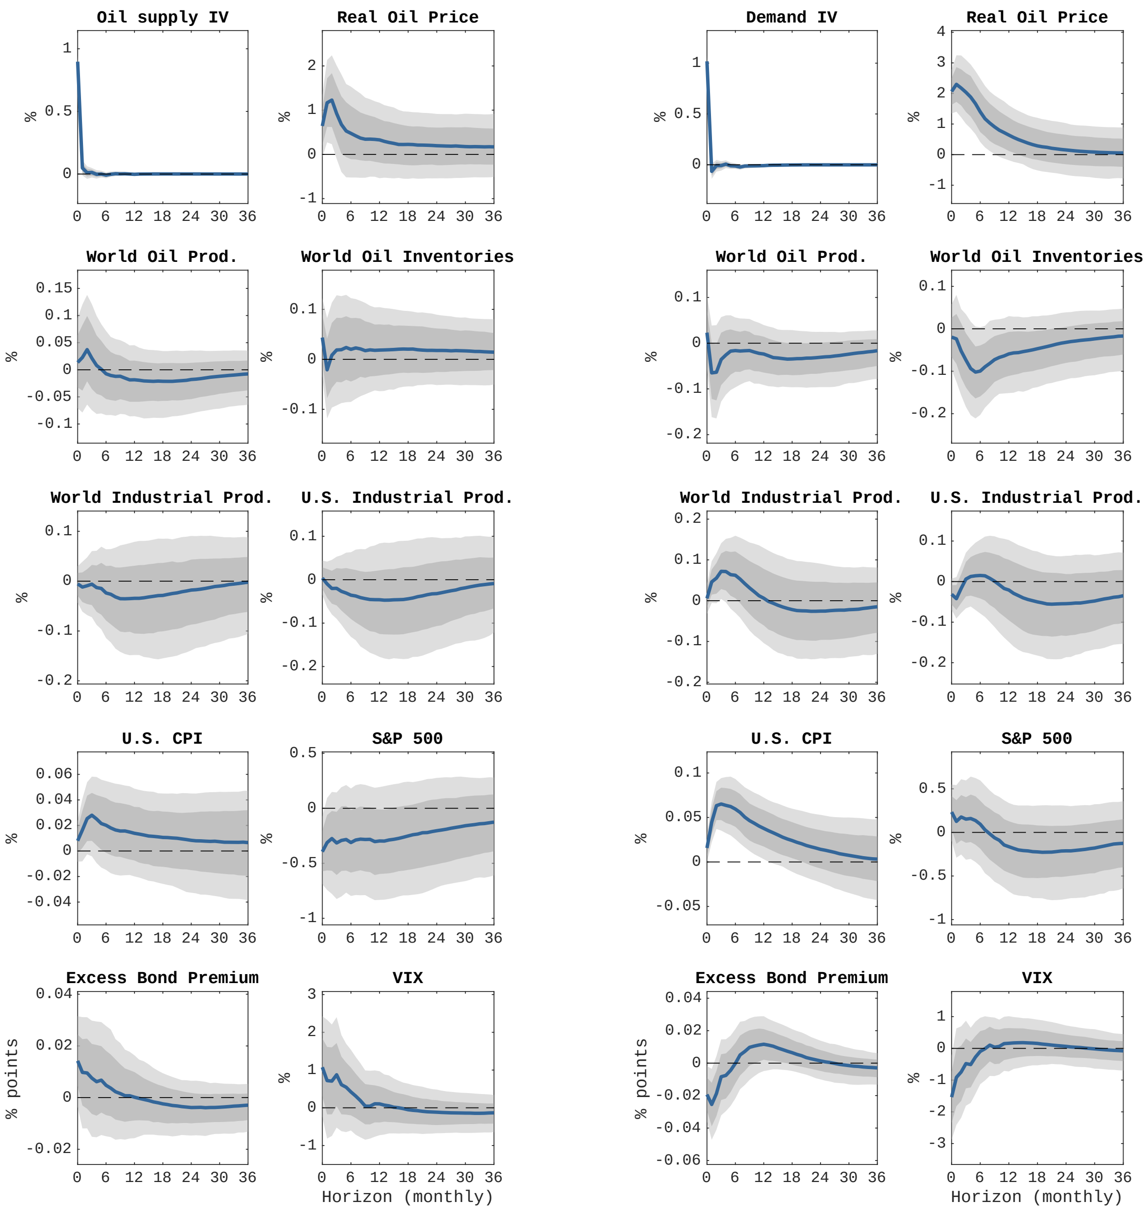

(a) Oil supply news shock

(b) Information shock

Note: Left: impulse responses to a one-standard-deviation oil supply news shock, identified using the robust instrument for oil news shocks as an internal instrument. Right: responses to a one-standard-deviation information shock, identified using the information component of the surprises as an internal instrument. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1983:7–2019:12.

a Cholesky decomposition – does not significantly alter the results.[42] Credibility regions are noticeably wider – particularly for the oil supply news shock – but the results are confirmed in both the full and pre-pandemic samples (Figures C.4 and C.5).

C.5 Dealing with informational insufficiency

To make sure that results are not confounding the problems stemming from informational insufficiency in the VAR with those deriving from the failure of the proxy exogeneity assumption, I re-estimate the effects of the shocks in an augmented VAR that includes financial variables, as suggested in Mori and Peersman (2024). They recommend augmenting the original 6-variables VAR of Känzig (2021) with the 1-year Treasury rate, the S&P 500, and the VIX. I report the results obtained using the instruments based on the high-frequency comovement between oil futures and stock prices and the results obtained using the alternative identification strategy that directly controls for the information set of OPEC. Results are not substantially different from those obtained from the 6-variable VAR, indicating that the endogeneity of the proxy remains a quantitatively important issue even when solving the issue of informational insufficiency. Figures C.6 and C.7 report the IRFs to the shocks identified on the full and pre-pandemic samples. Figure C.8 shows the IRFs to the shocks identified with the MOMR-based instruments.

[42]The instruments are included in the VAR sequentially, not jointly.

Figure C.6: Extended VAR – full sample

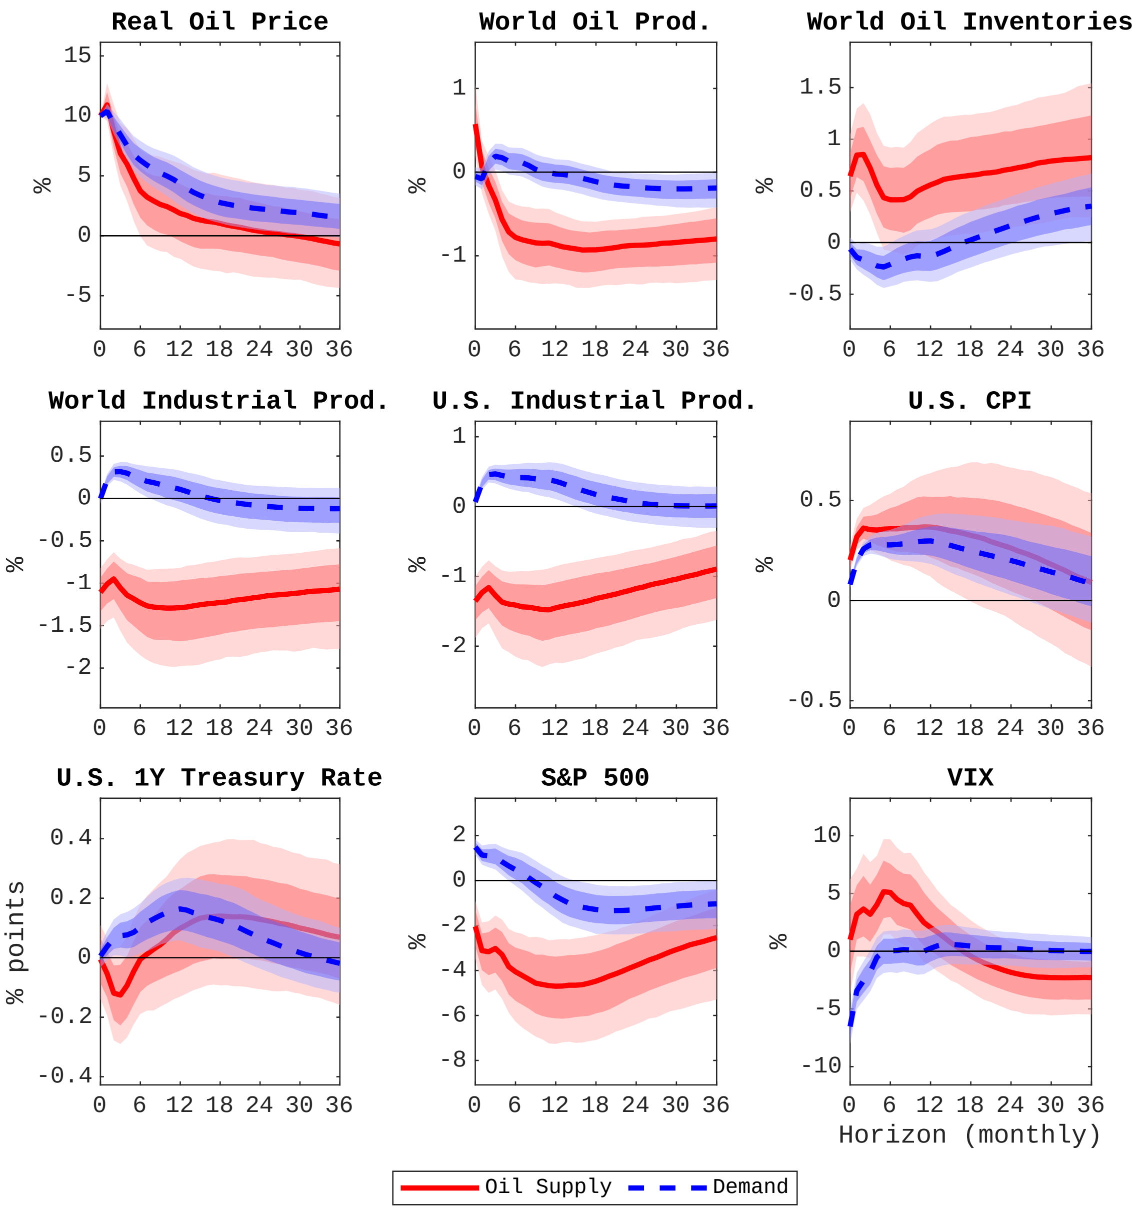

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies presented in Section 3.1.2. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1975:1–2023:5. IV sample: 1983:7–2023:6.

Figure C.7: Extended VAR – pre-Covid sample

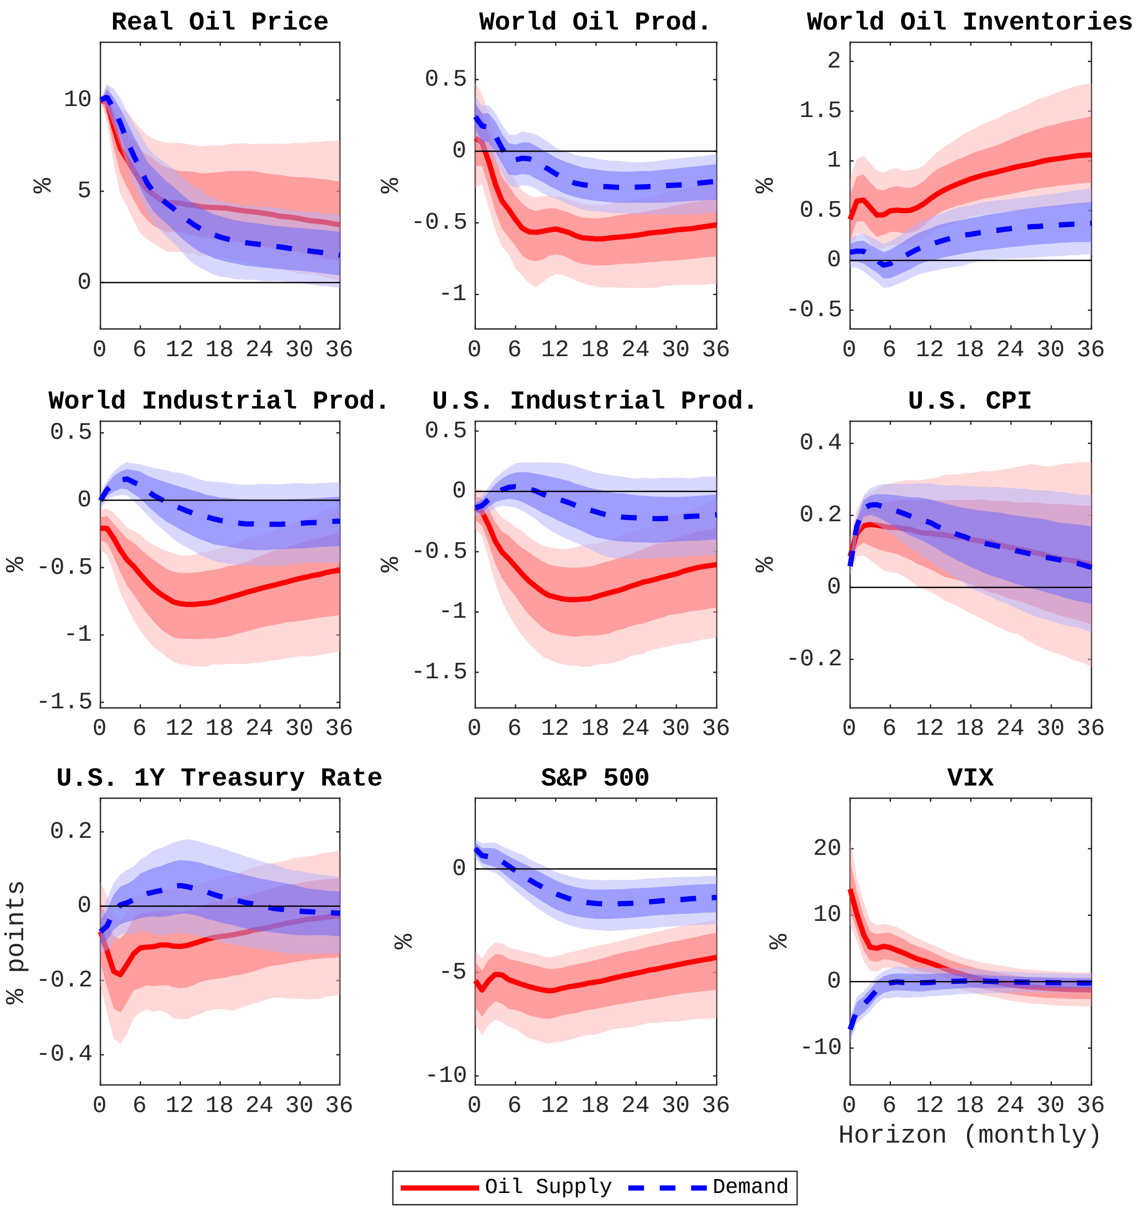

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies presented in Section 3.1.2. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1975:1–2019:12. IV sample: 1983:7–2023:6.

Figure C.8: Extended VAR – surprises orthogonalised wrt MOMRs

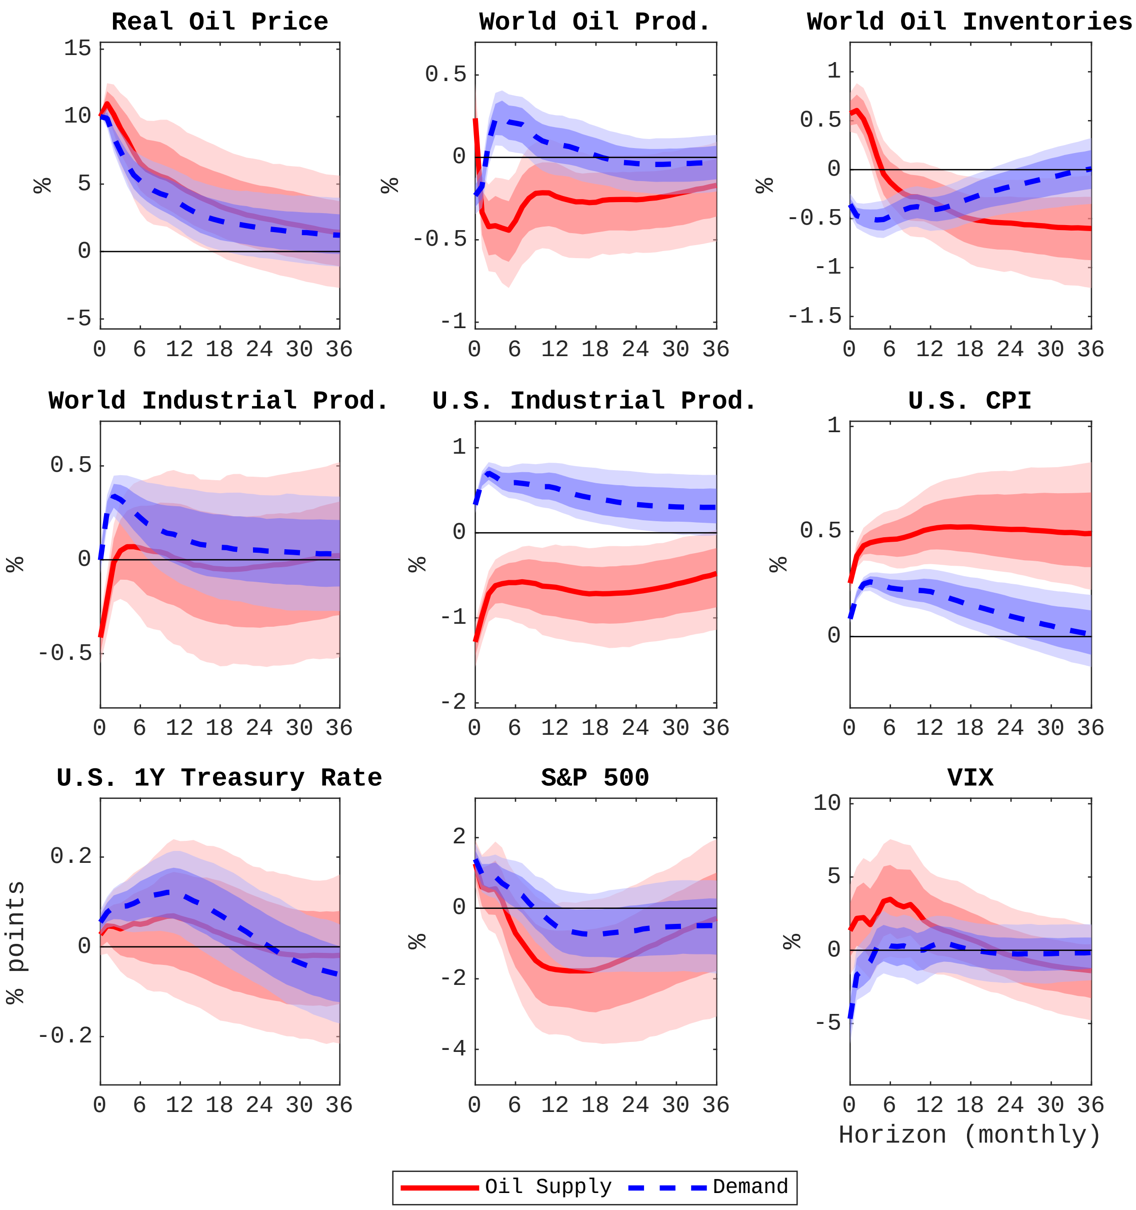

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies presented in Section 5.1. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1982:7–2023:5. IV sample: 2002:1–2021:3.

C.6 Forecast Error Variance and Historical Decompositions

I provide historical and forecast error variance (FEVD) decompositions based on the extended VAR model proposed by Mori and Peersman (2024) – which augments the 6-variable specification with the 1-year Treasury rate, the S&P 500, and the VIX. The cumulative historical contribution of the two shocks to the real price of oil is shown in Figure C.9. The FEVD for both shocks is reported in Figure C.10. Results are reported for the full sample (1975:1–2023:5); they are similar when estimated on the pre-pandemic sample.

Figure C.9: Historical decomposition of real oil price

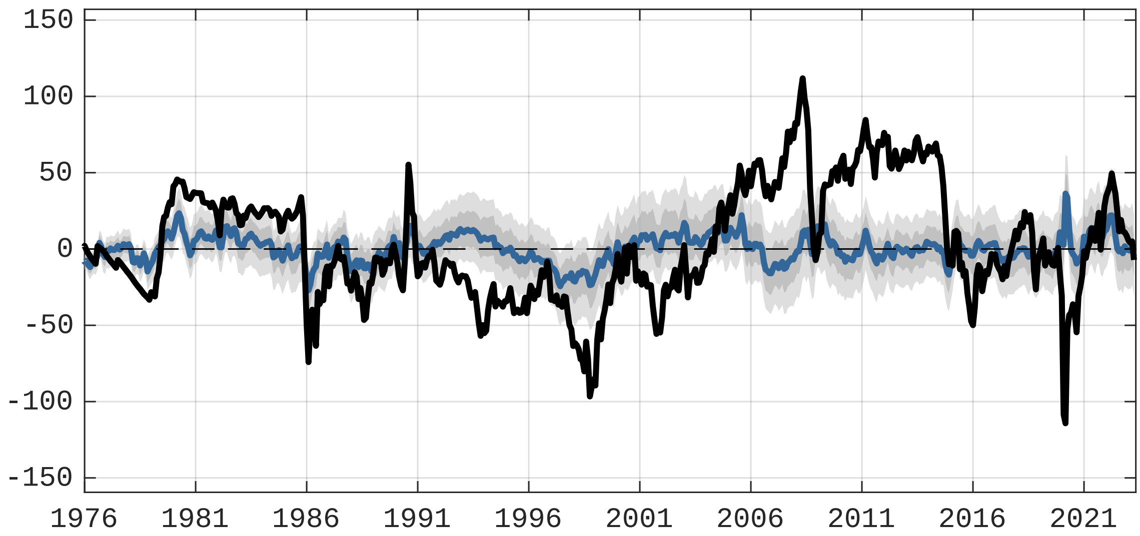

(a) Oil supply news shock

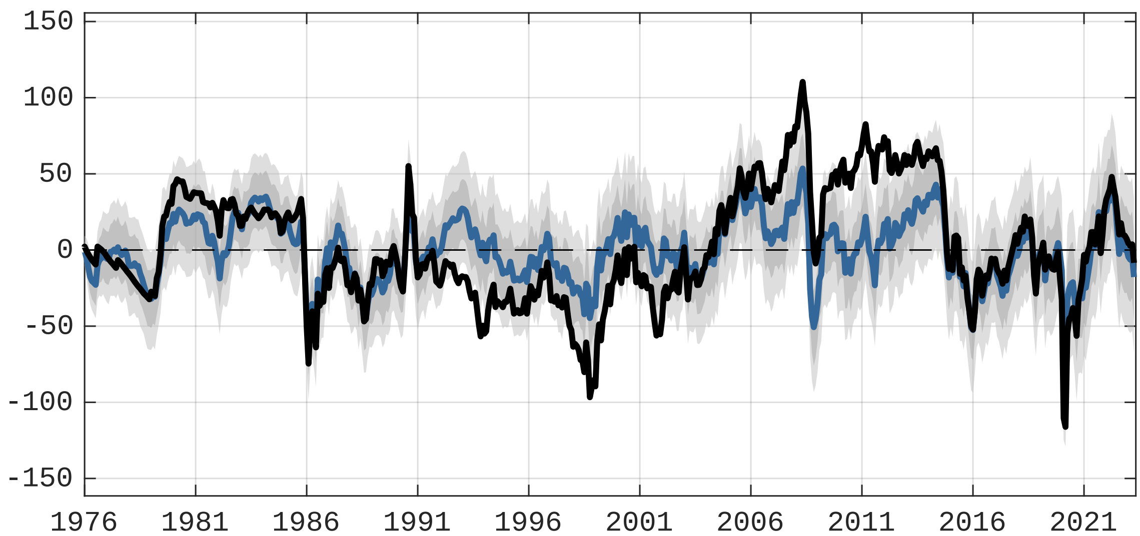

(b) Information shock

Note: Upper panel: cumulative historical contribution of oil supply news shocks to the real price of oil. Lower panel: cumulative historical contribution of information shocks to the real price of oil. Solid black: real price of oil (in percent deviations from the mean). Shaded areas represent 68% and 90% posterior coverage bands. The IRFs of the underlying model are presented in Figure C.6.

Figure C.10: Forecast Error Variance Decomposition

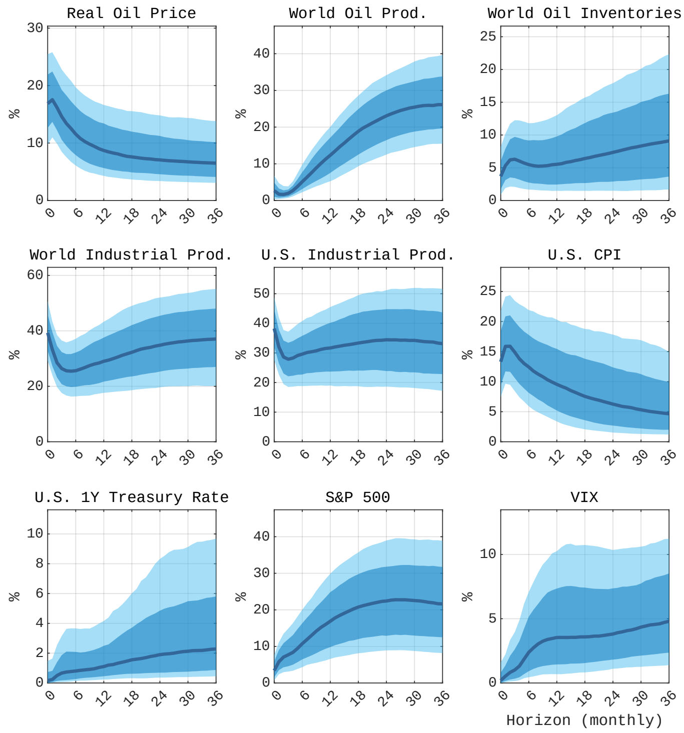

(a) Oil supply news shock

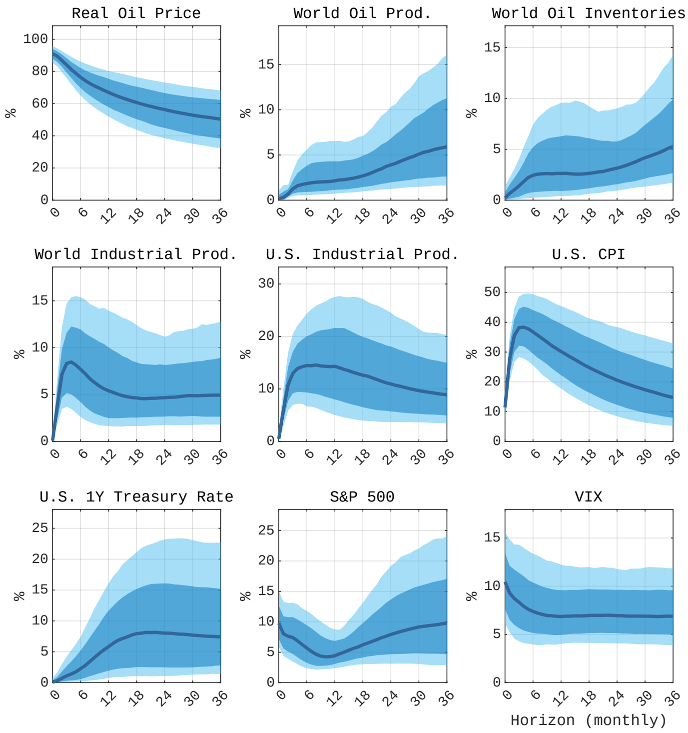

(b) Information shock

Note: Upper panel: share of forecast error variance explained by oil supply news shocks. Lower panel: share of forecast error variance explained by information shocks. Shaded areas represent 68% and 90% posterior coverage bands. The IRFs of the underlying model are presented in Figure C.6.

C.7 Controlling for low liquidity of WTI futures contracts

Kilian (2024) highlights two important issues with using contracts at distant maturities in the construction of the proxies: these contracts started trading later than March 1983 and, in the early part of the sample, were scarcely liquid. Contracts with maturities between the current month and 6 months began trading in March 1983. Contracts between 7 and 10 months began trading in November 1983. The 11-month contract started trading in December 1983, and the 12-month contract in April 1984. Känzig (2021) replaces the surprises in the prices of contracts that had not yet started trading with zeros, which may bias the estimation of the principal components.

To test whether this affects my results, I select a subset of maturities that were already trading from March 1983. Specifically, I consider maturities between 1 and 4 months, which were reasonably liquid even in the early sample. I re-estimate the first principal component of the surprises in the prices of these contracts and separate oil supply and demand shocks based on their comovement with the surprises in the S&P 500. The results are robust to identifying the shocks using these new instruments, which account for the low liquidity of WTI futures contracts in the 1980s and avoid imputing zeros for surprises in contracts that had not yet started trading (Figure C.11).

To be precise, Kilian (2024) recommends dropping the sample before April 1989 when constructing the proxies, as before that date trading in oil futures markets was limited to selected dates and maturities. Using only maturities between 1 and 4 months, I re-estimate the first principal component in the surprises only on OPEC-conference days posterior to April 1989. There are 133 such events in my sample, of which 70 identify an oil supply news shock, given their comovement with the S&P 500. The impulse responses obtained from a VAR estimated on the sample from April 1988 to June 2023 to the shocks identified with these proxies, which span the period 1988:4–2023:6, are reported in Figure C.12. The proxies suffer from instrumental weakness, as the robust F-statistics are 5.3 for the oil supply news shock and 5.5 for the demand shock, but the dynamic responses obtained match quite well the baseline results.

Figure C.11: Using only maturities between 1 and 4 months

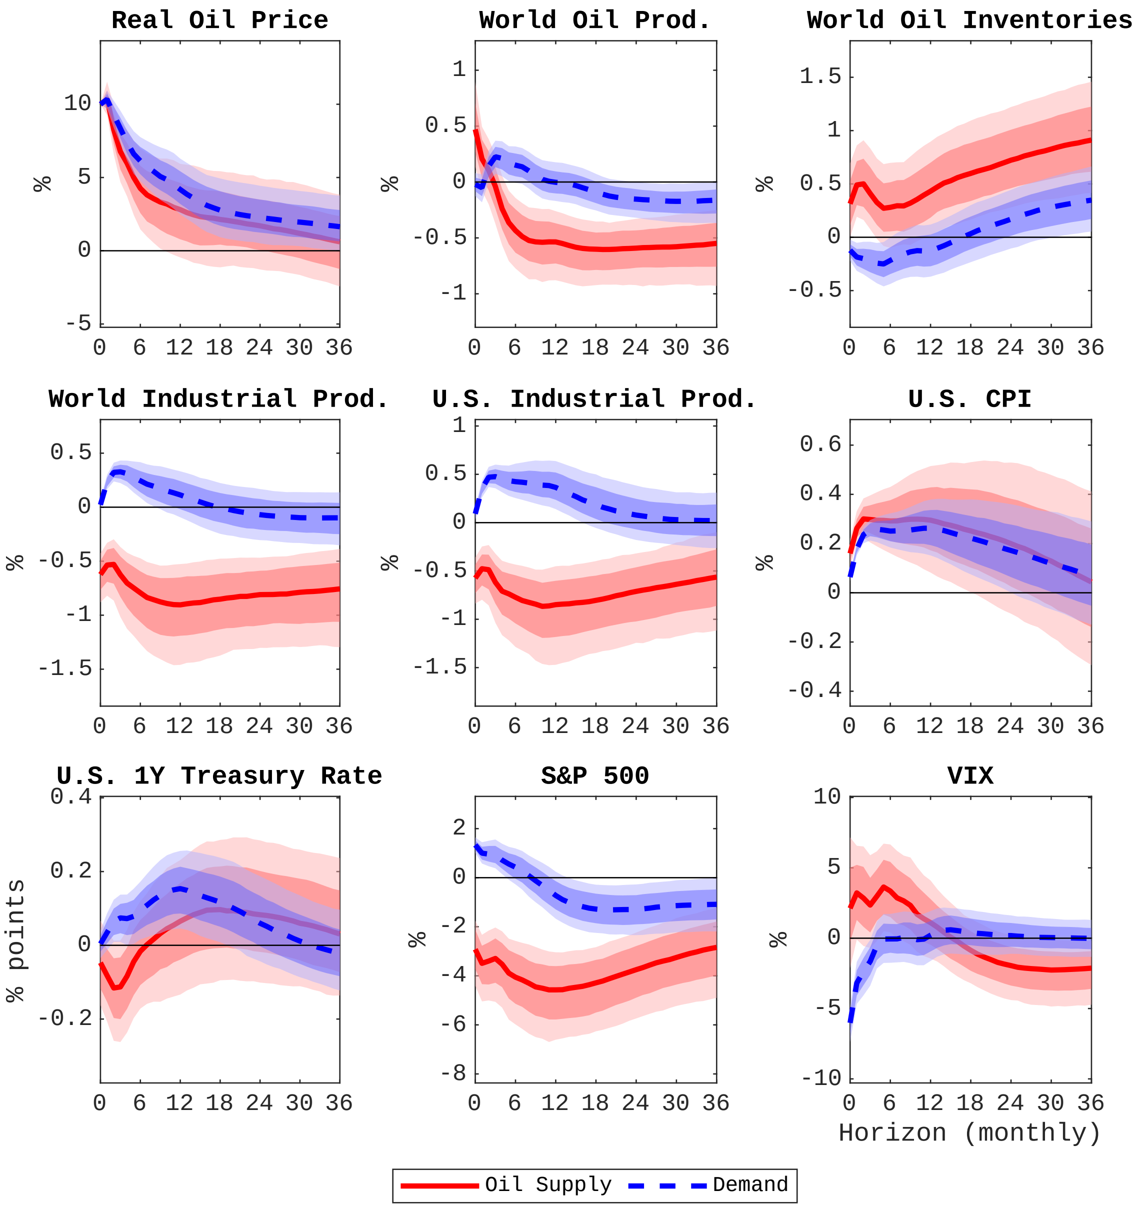

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies based on the comovement between the first principal component of the surprises in the price of WTI futures contracts at maturities between 1 and 4 months, and the surprises in the S&P 500, as described in Section C.7. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1975:1–2023:5. IV sample: 1983:7–2023:6.

Figure C.12: USING ONLY DATA FROM APRIL 1989 AND MATURITIES 1 TO 4 MONTHS

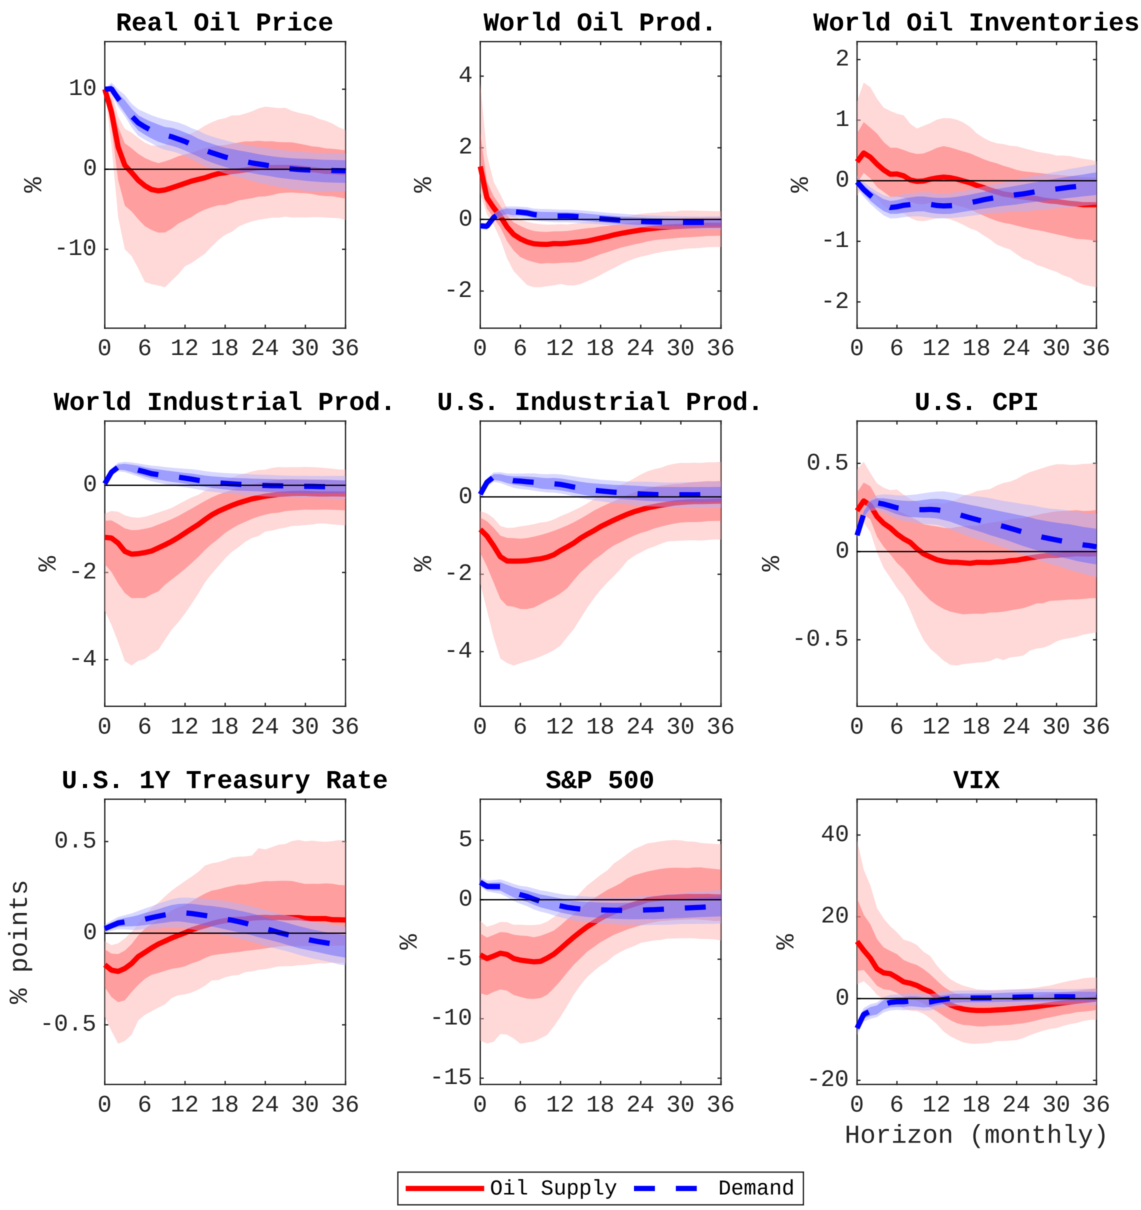

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies based on the comovement between the first principal component of the surprises in the price of WTI futures contracts at maturities between 1 and 4 months, and the surprises in the S&P 500, as described in Section C.7. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1988:4–2023:5. IV sample: 1989:4–2023:6.

C.8 Controlling for temporal aggregation bias

Kilian (2024) also proposes a strategy to address temporal aggregation bias in the construction of the surprises. Typically, no adjustment is made for whether a shock occurs closer to the beginning or the end of the month. This can be particularly problematic when the oil price included in the VAR is the average over the month. The implicit assumption is that shocks of the same size have the same effect on the monthly average oil price, regardless of their timing (see also Lee and Sekhposyan, 2024, for a discussion of different weighting schemes).

My results should be less affected by this issue, as I use the end-of-month WTI price starting from 1986.1.[43] Nevertheless, a robustness check in which I apply the reweighting procedure suggested in Kilian (2024) yields results that are similar to those in the main text (Figure C.13). The reweighting is applied to the instruments constructed using maturities between 1 and 4 months, in order to avoid the issues related to low liquidity discussed in Section C.7. First, the factor in the surprises is split based on its comovement with the surprises in the S&P 500; the resulting components are then reweighted accordingly.

[43] This series is spliced with the average monthly oil price before that date.

Figure C.13: Kilian (2024)'s correction for aggregation bias

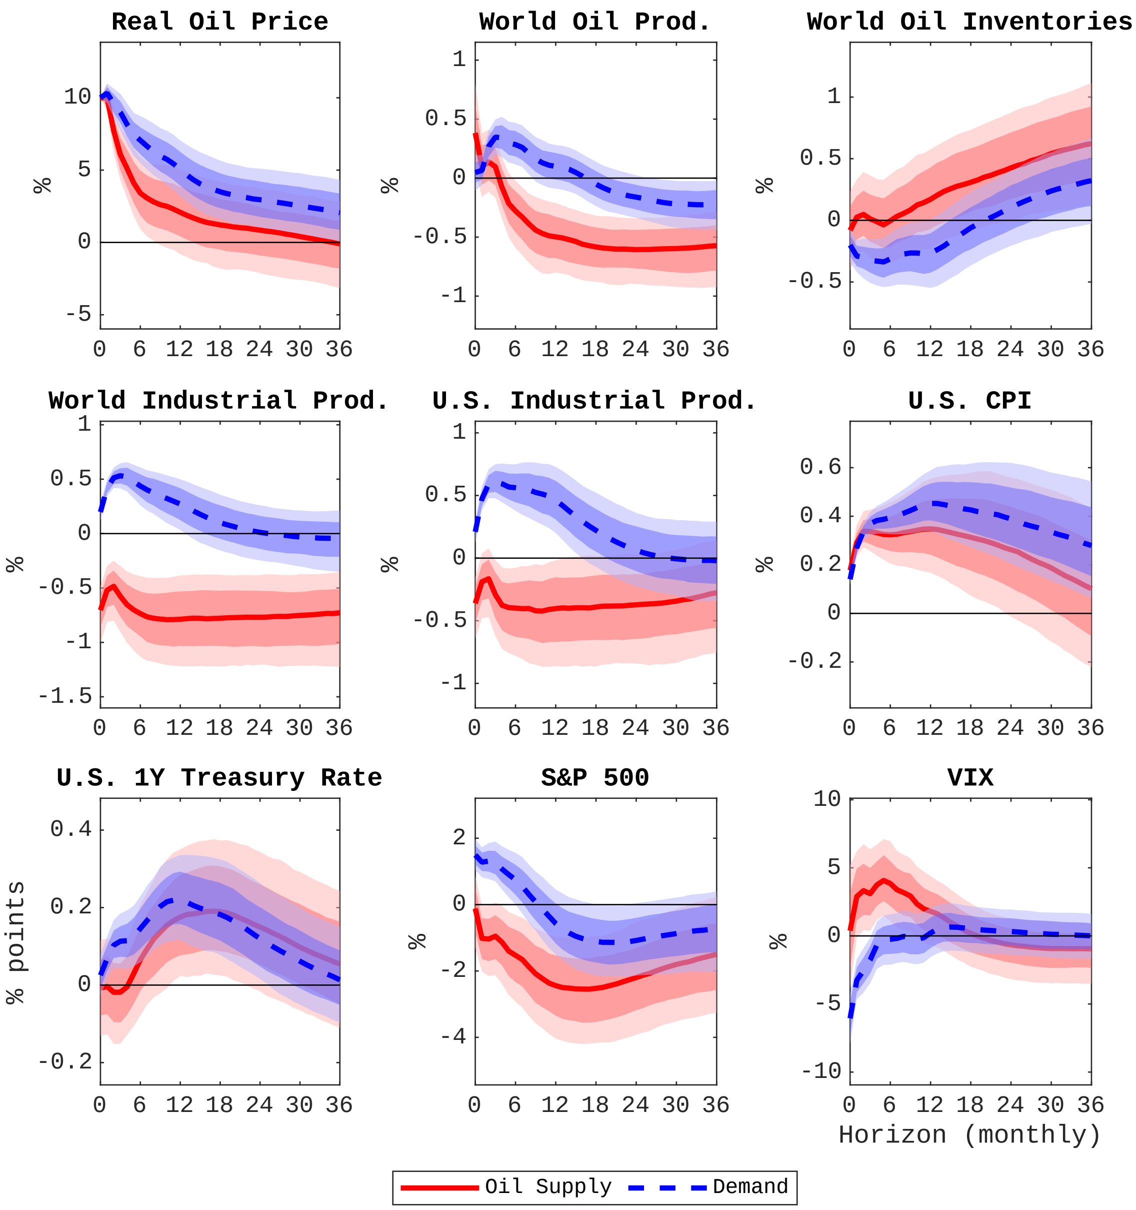

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies based on the comovement between the first principal component of the surprises in the price of WTI futures contracts at maturities between 1 and 4 months, and the surprises in the S&P 500, reweighted following the procedure proposed in Kilian (2024), as described in Section C.8. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 1975:1–2023:5. IV sample: 1983:7–2023:6.

C.9 Identification based on U.S. daily real business conditions

The high-frequency co-movement between oil futures and the stock price index can be used to disentangle the shocks because stock prices are a high-frequency proxy for economic activity. However, they are not the only available high-frequency measure of economic activity. Figure C.14 shows that the oil supply news and demand components in the surprises can be separated based on the co-movement of oil futures and the daily measure of U.S. real business conditions proposed by Aruoba et al. (2009). This approach delivers responses for U.S. production that are significantly contractionary for a negative oil supply news shock and expansionary for a positive information shock. However, residual traces of puzzles in the response of world industrial production are still visible.

Figure C.14: Identification based on Aruoba et al. (2009)

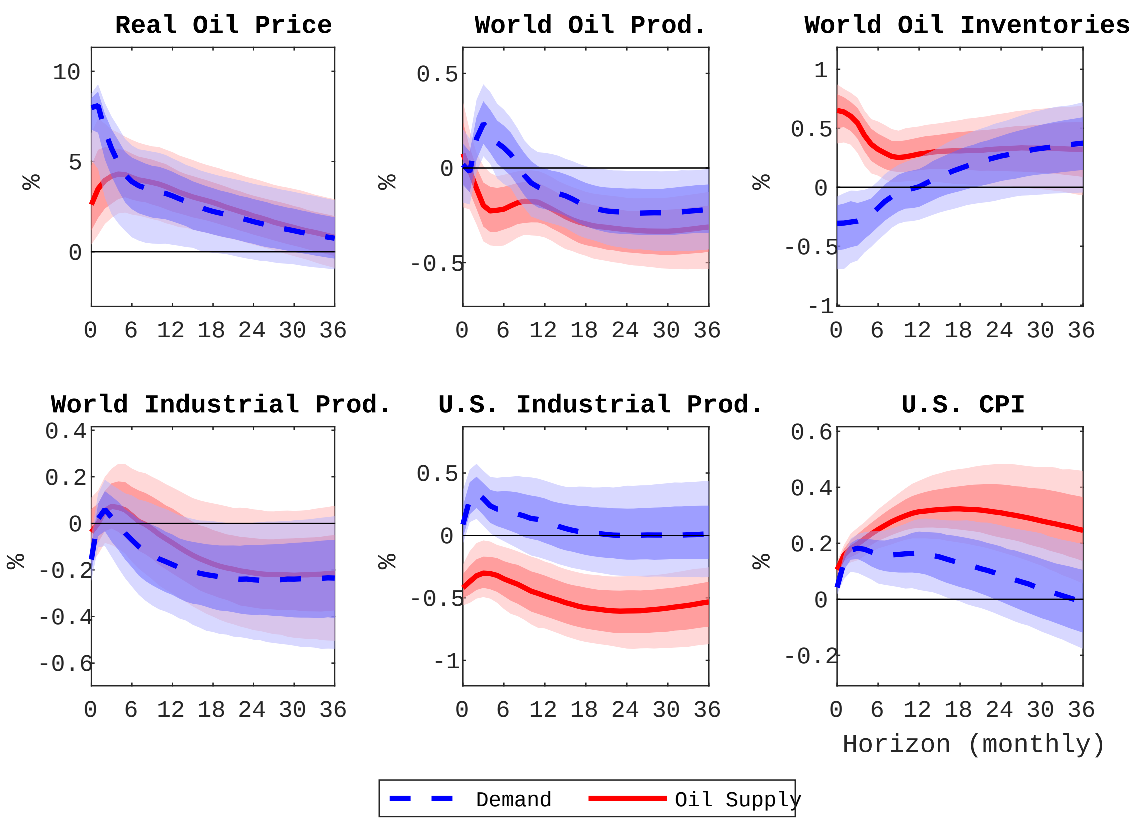

Note: Solid red: impulse responses to a one-standard-deviation oil supply news shock. Dashed blue: responses to a one-standard-deviation information shock. The shocks are jointly identified using the robust proxies, constructed based on the comovement on OPEC conference days between daily surprises in oil futures and changes in the daily indicator of U.S. real business conditions of Aruoba et al. (2009). BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1975:1–2023:5.

C.10 Removing market holidays

Another concern is that the inclusion of OPEC conferences that happened during market holidays, for which the surprise is computed on the first day of market reopening, might have an unduly large effect on the identification of the responses. Between 19/07/1983 and 04/06/2023 there are 151 announcement days. Of these, 29 are market holiday days. As a test that results are robust to the exclusion of announcements during market holidays, the robustness exercise in Section 5.3 is repeated removing the 29 observations that coincide with market holidays. The impulse responses obtained are in line with the baseline results.

Figure C.15: Removing market holidays

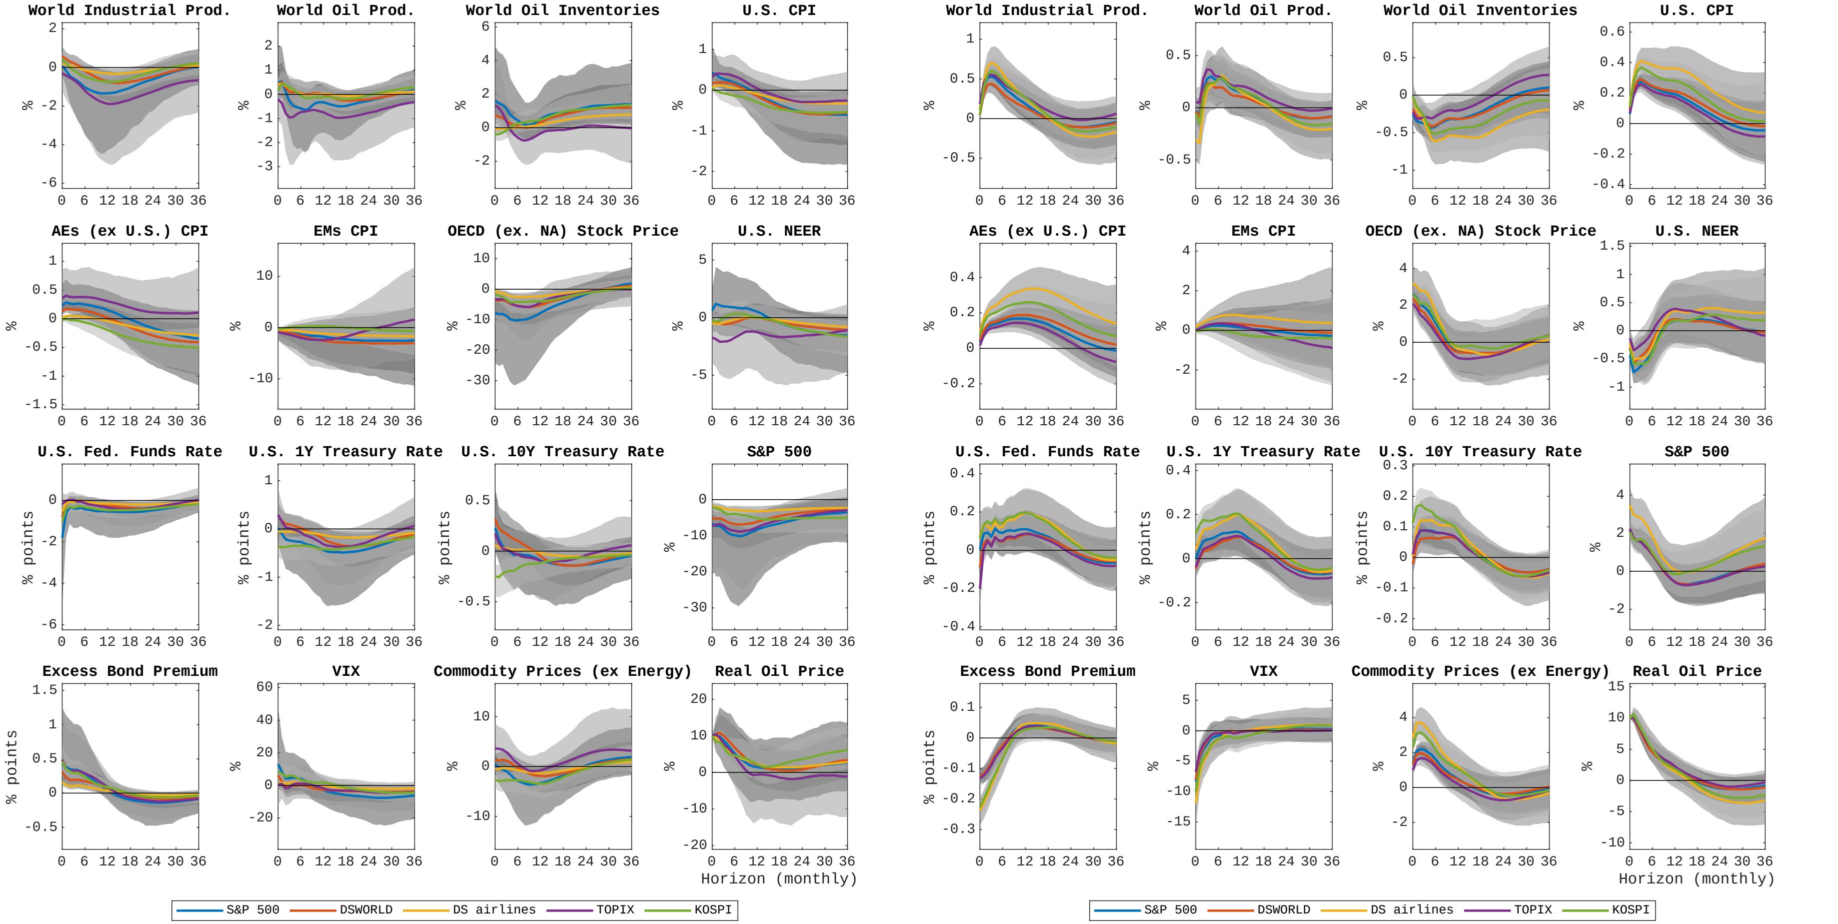

(a) Oil supply news shock

(b) Information shock

Note: Impulse responses obtained by dropping market holidays from the regression of the reduced-form VAR innovations on the instrument. Left panel: responses to an oil supply news shock. Right panel: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using robust proxies, constructed based on the comovement on OPEC conference days between daily surprises in oil futures and changes in different stock price indices, as in Figure 12. **Blue**: S&P 500; **Red**: DS World; **Yellow**: DS Airlines; **Purple**: TOPIX; **Green**: KOSPI. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1980:3–2023:5. All proxies span the period 1983:7–2023:6.

C.11 Removing most influential observations

The identification of the shocks is based on the regression of the reduced-form innovations of the VAR on the proxy and a constant. One might worry that a limited number of important events, in particular the largest spikes in Figure 2, might have an unduly large influence on the coefficient estimates that represent the impact of the shocks on the variables of interest. In this subsection it is shown that removing the 6 most influential observations (MIO) from the first-stage regression does not change the results, and indeed makes them even stronger.

Figure C.16 displays the responses to a negative oil supply news shock (left) and to a positive information shock (right) identified with the robust instruments based on the S&P 500 (solid red) and with the same instrument without the 6 MIOs (dashed blue). The 10 MIO for both shocks are listed in Table C.1. The MIOs are determined using the Stata function dfbeta, which provides an influence statistic. Dfbeta computes the difference in the coefficient estimate when a specific observation is included or excluded from the sample. Results for both shocks are robust to the exclusion of the 6 MIO, the precision of the responses does not change substantially, and the magnitude of the expansion following the information shock is even larger.

Figure C.16: Removing Most Influential Observations

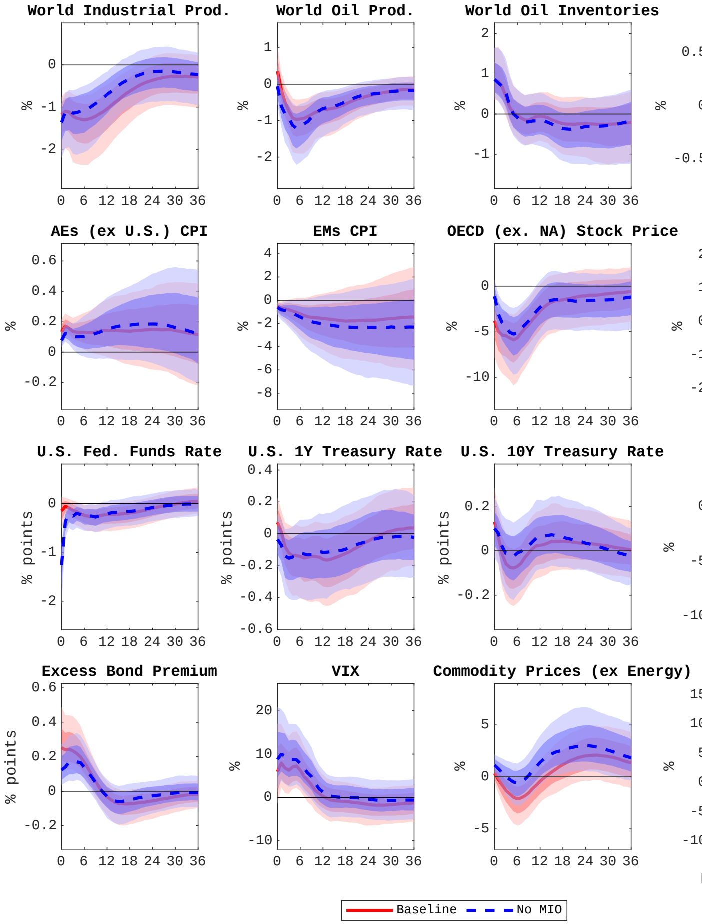

(a) Oil supply news shock

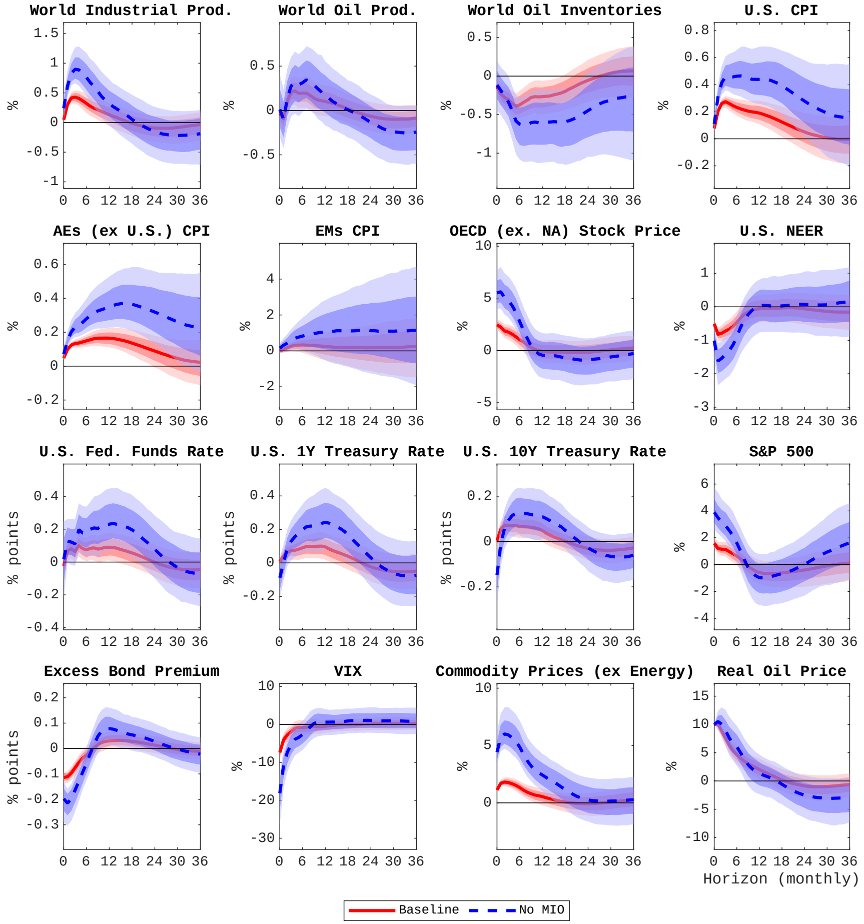

(b) Information shock

Note: Impulse responses obtained by removing the highest and lowest three most influential observations from the regression of the reduced-form VAR innovations on the instrument. Left panel: responses to an oil supply news shock. Right panel: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1980:3–2023:5.

Table C.1: Most Influential Observations

Information shock		Oil supply news shock	
Date	DFBETA	Date	DFBETA
2021:7	-1.076	2021:3	-0.333
2014:11	-0.417	2008:2	-0.291
1988:11	-0.209	2000:9	-0.275
2008:3	-0.206	2020:6	-0.244
1992:2	-0.201	2006:3	-0.193
2020:3	1.495	2016:11	0.720
1986:8	0.553	2001:11	0.368
1998:11	0.233	1986:12	0.278
2008:10	0.230	1987:12	0.224
2021:12	0.145	2022:10	0.137

Note: Most influential observations in the regression of the reduced-form VAR innovations corresponding to the real oil price equation on the instrument (left: for the information shock; right: for the oil supply news shock) and a constant. The proxies are based on the S&P 500. The innovations are from the system represented in Figure C.16, estimated on the sample 1982:7–2023:5. DFBETA measures the difference in the parameter estimate with and without the influential observation.

C.12 Alternative measures of industrial production

Figure C.17 shows that results are robust to using alternative measures of world industrial production. The baseline measure used throughout the paper is the OECD plus six index by Baumeister and Hamilton (2019). The alternative measures are the Dallas Fed world (excluding U.S.) industrial production and the OECD industrial production from the OECD Main Economic Indicators.

Figure C.17: ALTERNATIVE MEASURES OF INDUSTRIAL PRODUCTION

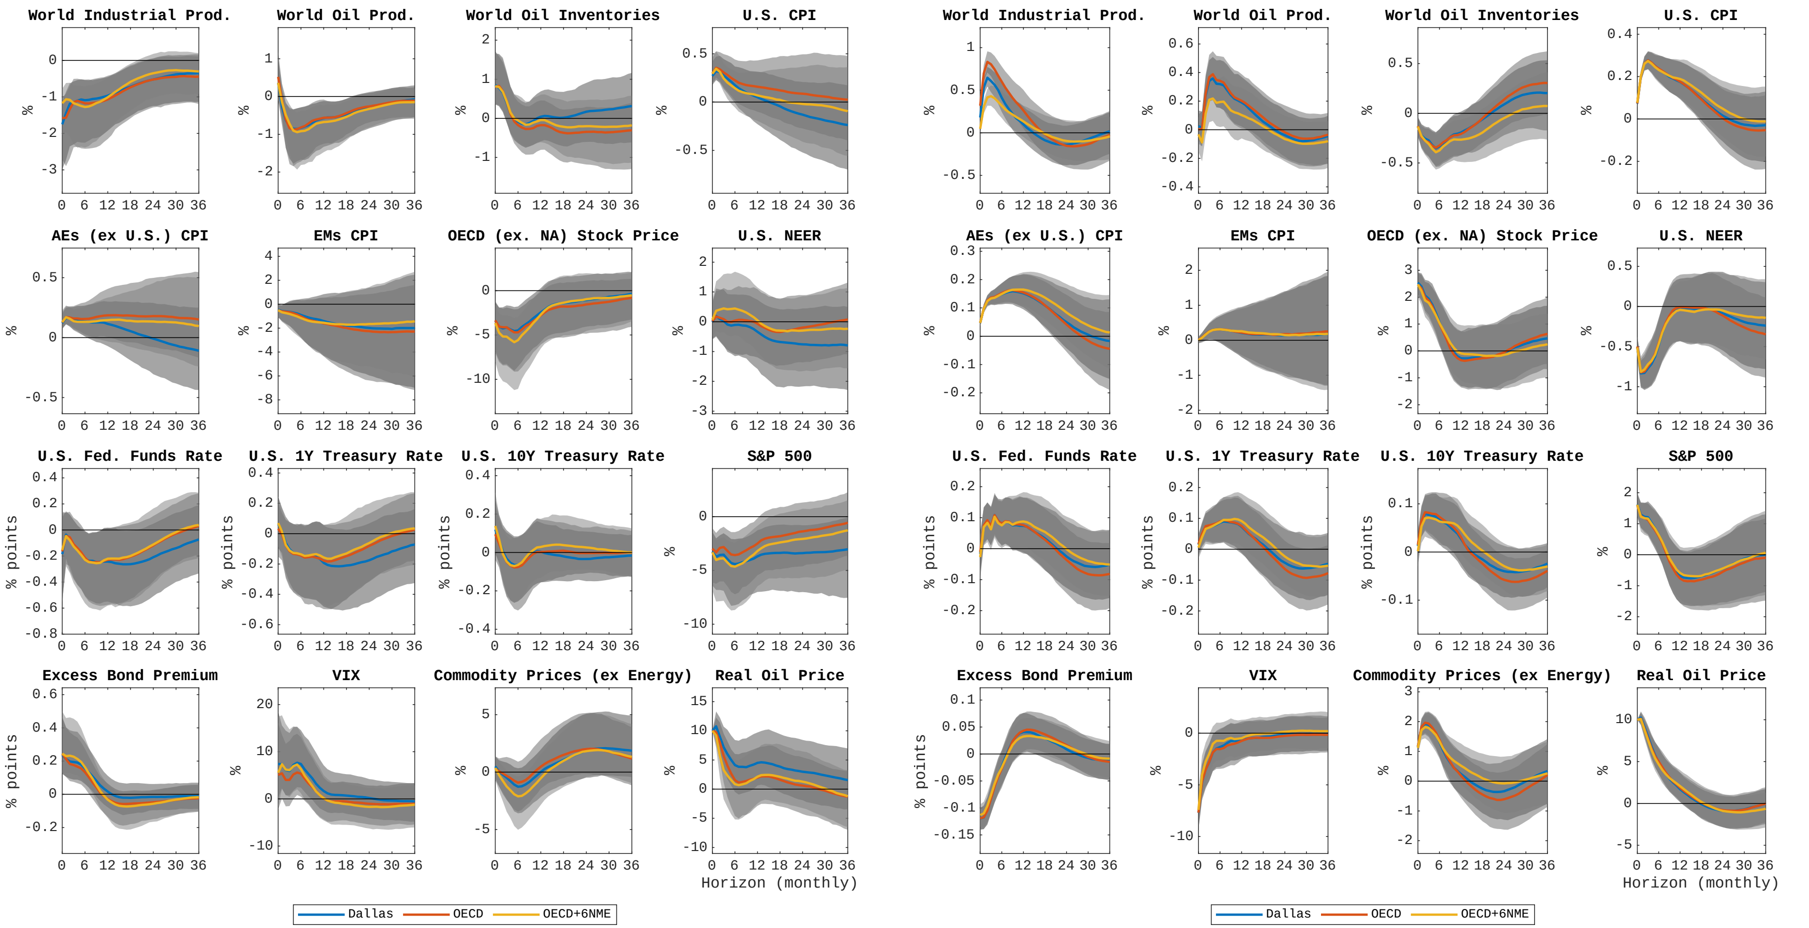

(a) Oil supply news shock

(b) Information shock

Note: Comparison among three alternative measures of World Industrial Production. Left panel: responses to an oil supply news shock. Right panel: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using robust proxies. **Blue**: World production excluding U.S., from the Dallas Fed; **Red**: OECD production, from the OECD Main Economic Indicators; **Yellow**: OECD production plus 6 major emerging markets, from Baumeister and Hamilton (2019). BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1980:3–2023:5.

C.13 Excluding the Covid-19 period

Results are robust to cutting the sample at December 2019, before the Covid-19 pandemic, or to reducing the informativeness of the Covid period in the estimation of the reduced-form parameters of the VAR by including dummy variables for the months between March 2020 and August 2020 by casting pandemic priors (Cascaldi-Garcia, 2022). This is arguably an important check as Covid-19 produced deep discontinuities in the time series used in the analysis. Figure C.18 shows that the impulse responses coincide with the baseline results. The same holds when identifying the shocks on a smaller system (Figure C.7). It is interesting to notice that 2020:3 and 2020:6 appear in Table C.1, listing the most influential observations, but they are not *the* most influential observations.

Figure C.18: Transmission to the Global Economy excluding Covid-19

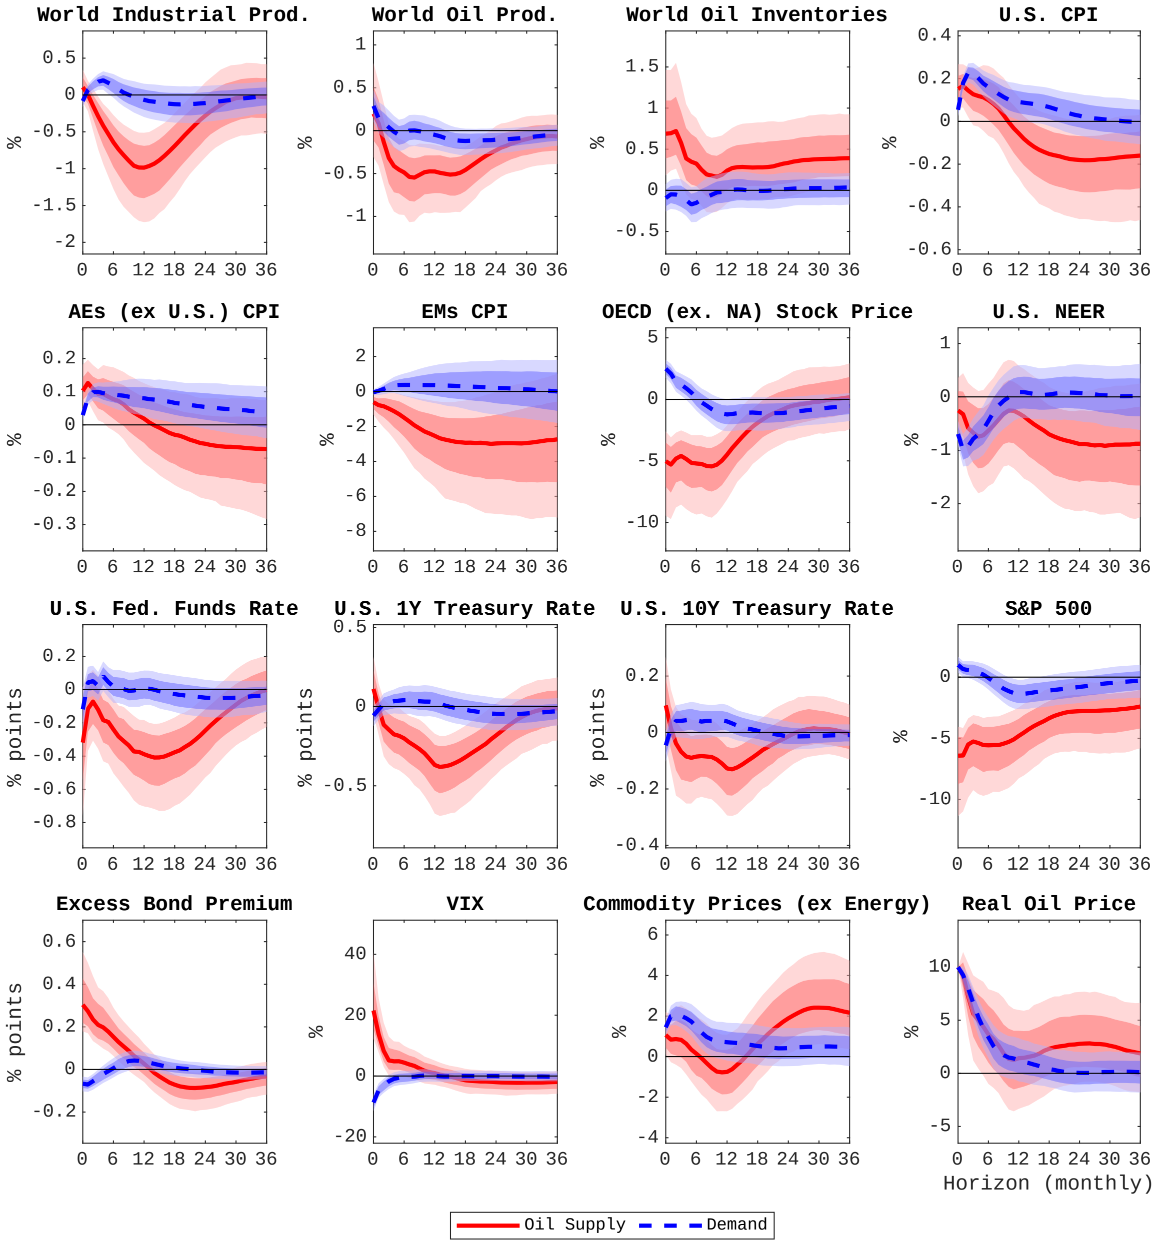

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1980:3–2019:12.

Figure C.19: Transmission to the Global Economy – Pandemic Priors

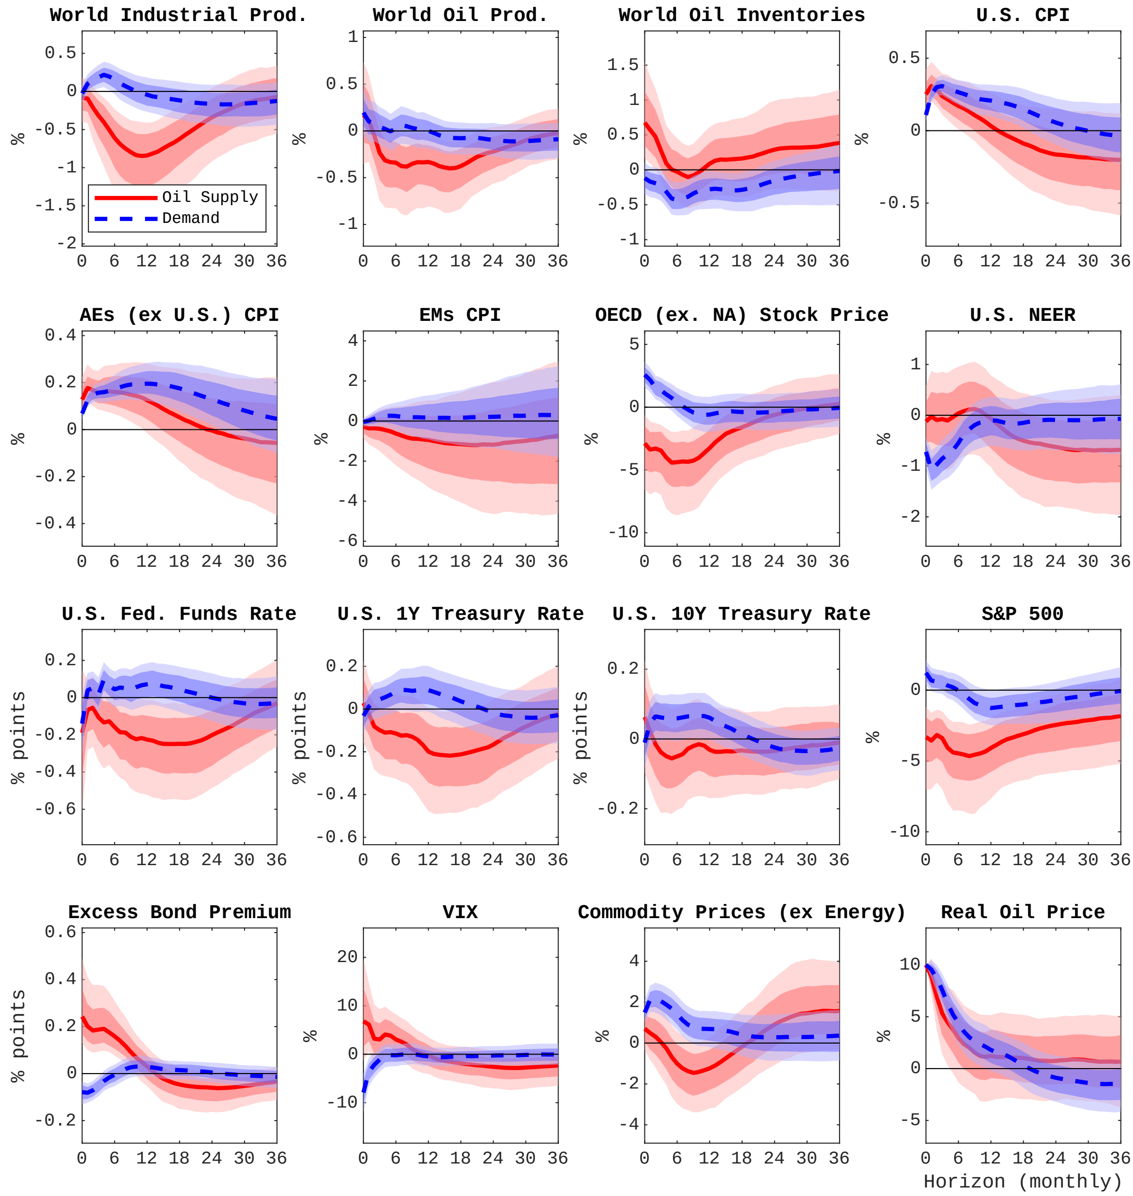

Note: Solid red: impulse responses to an oil supply news shock. Dashed blue: responses to an information shock. The informativeness of observations between 2020:3 and 2020:8 is scaled using pandemic priors (Cascaldi-Garcia, 2022). Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1980:3–2023:5.

C.14 Full IRFs from the bilateral-VARs exercise

This section reports the full impulse response functions of the exercise that estimates the effects of the shocks across 15 advanced and 15 emerging economies in Sections 4.5 and 4.6. Figure C.20 displays the full responses for the median advanced economy. Figure C.21 shows the responses for the median emerging economy.

Figure C.20: Transmission to advanced economies

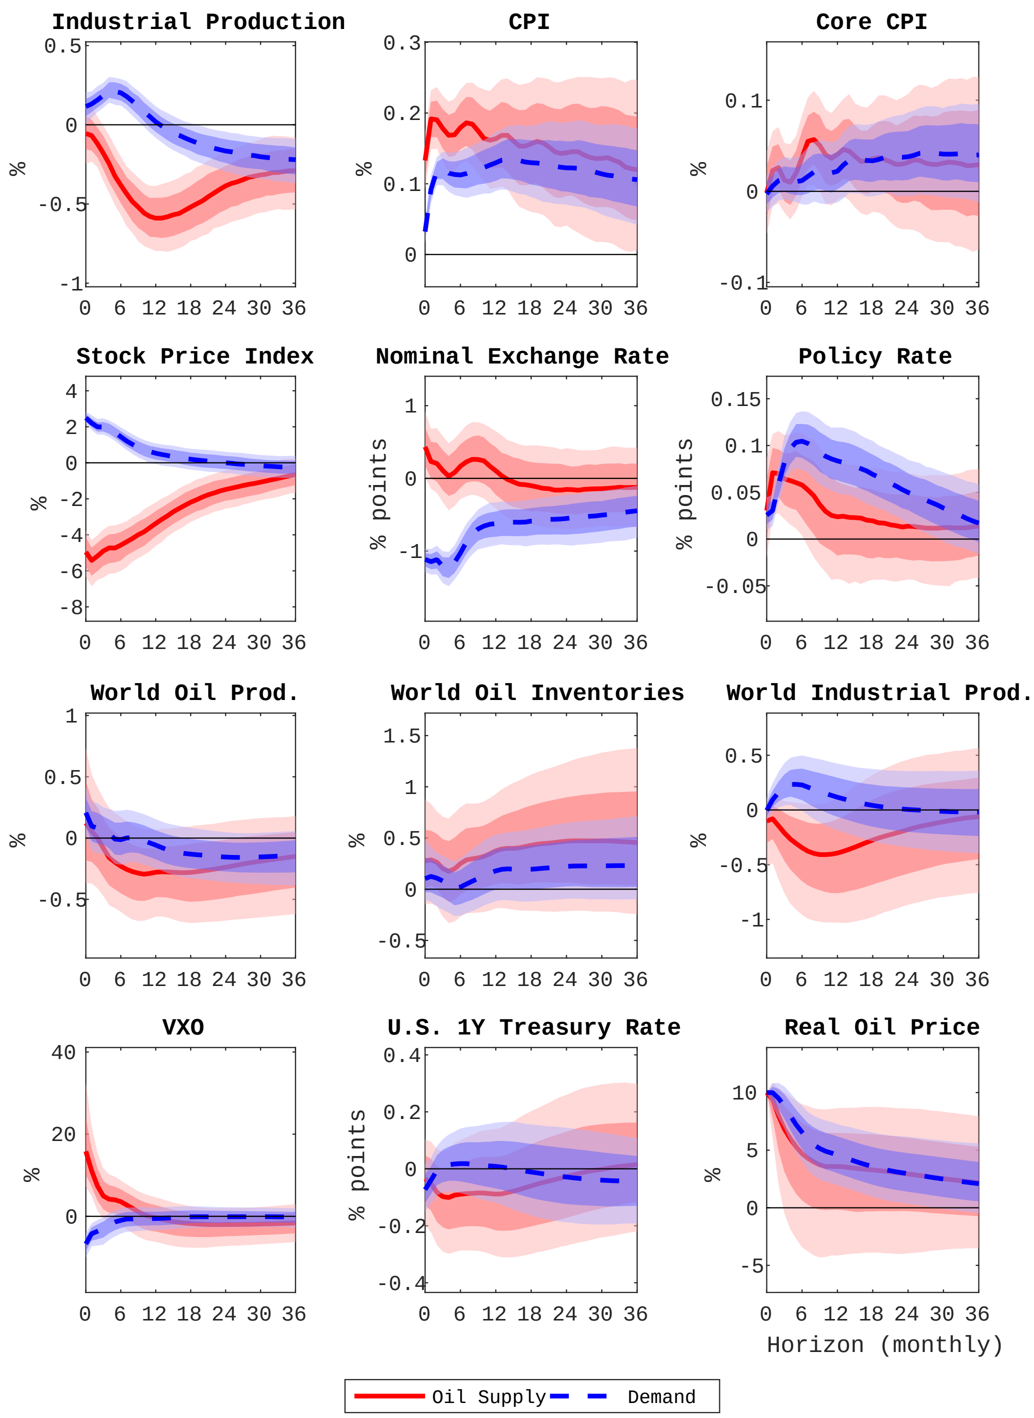

Note: Impulse responses for the median advanced economy. Solid red: responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1. Both proxies span the period 1983:7–2023:6.

Figure C.21: Transmission to emerging markets

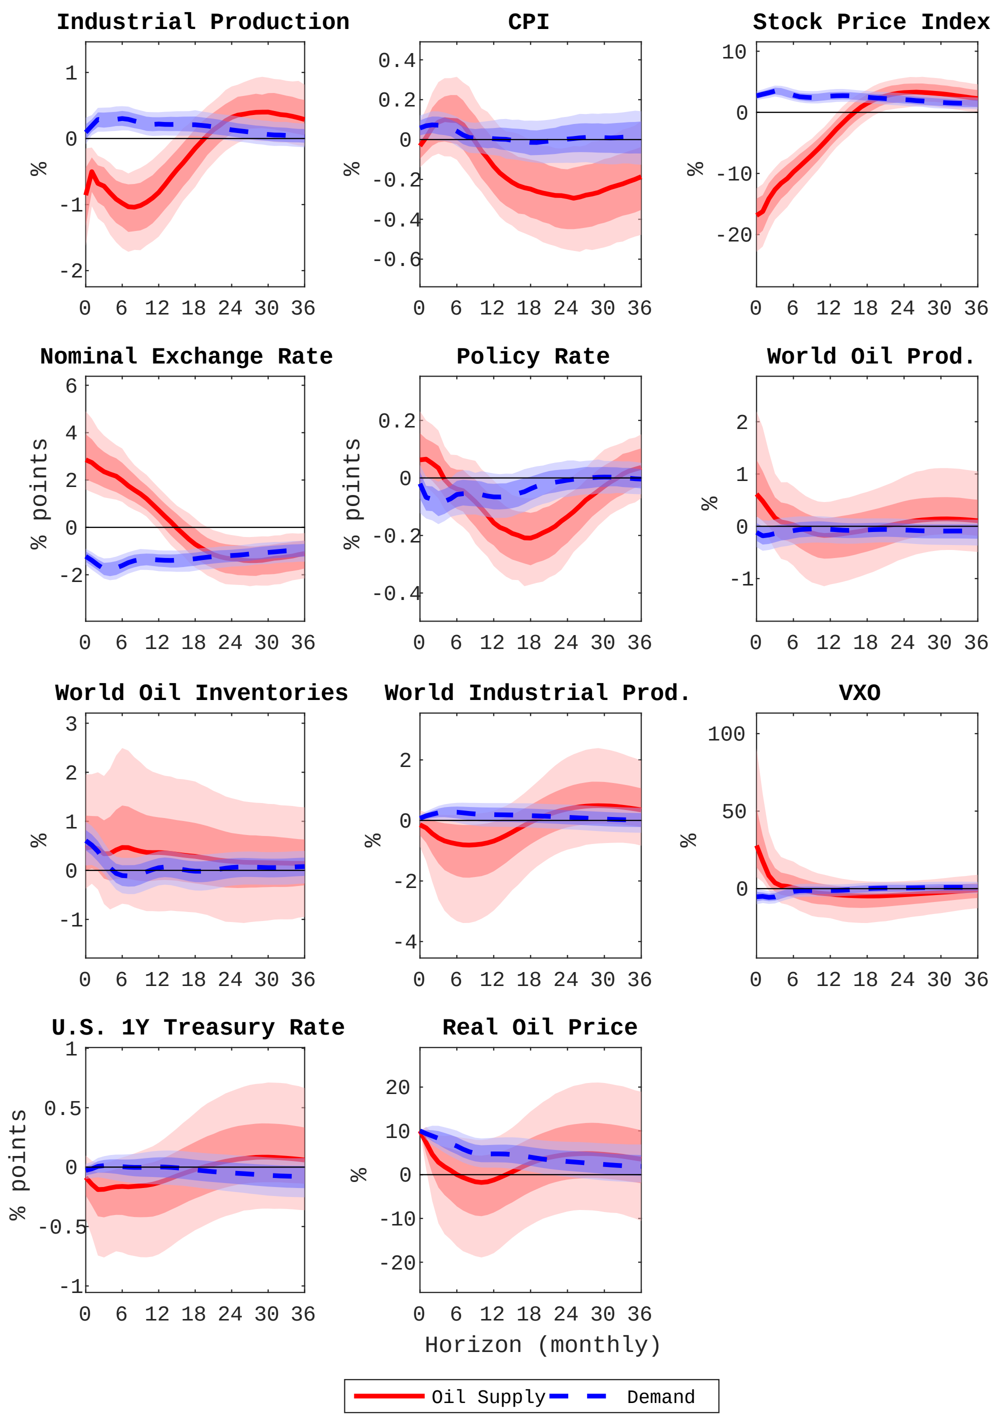

Note: Impulse responses for the median emerging market. Solid red: responses to an oil supply news shock. Dashed blue: responses to an information shock. Both shocks are normalised to induce a 10% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1. Both proxies span the period 1983:7–2023:6.

C.15 Country-specific responses of emerging markets

This section provides evidence on the homogeneity of results across the 15 emerging markets used for the analysis in Section 4.6. Figure C.22 displays the responses of the bilateral domestic/U.S. dollar exchange rate for all 15 countries to an oil supply news shock normalised to induce a 10% increase in real oil price. Figure C.23 shows the response of the stock price index to the shock.

Figure C.22: Country-specific responses of the Exchange Rate

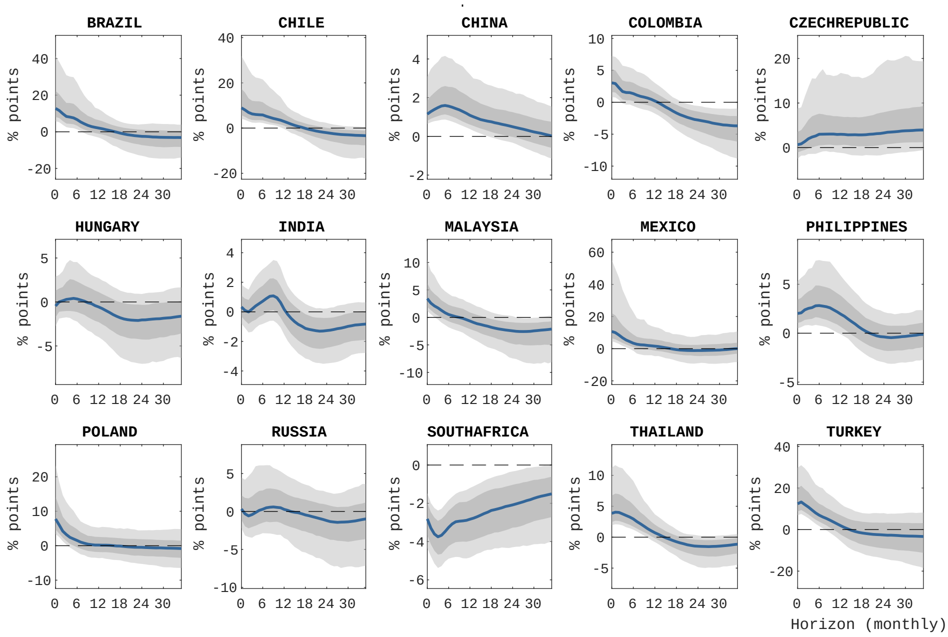

Note: Responses to an oil supply news shock normalised to induce a 10% increase in real oil price. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1.

Figure C.23: Country-specific responses of Stock Prices

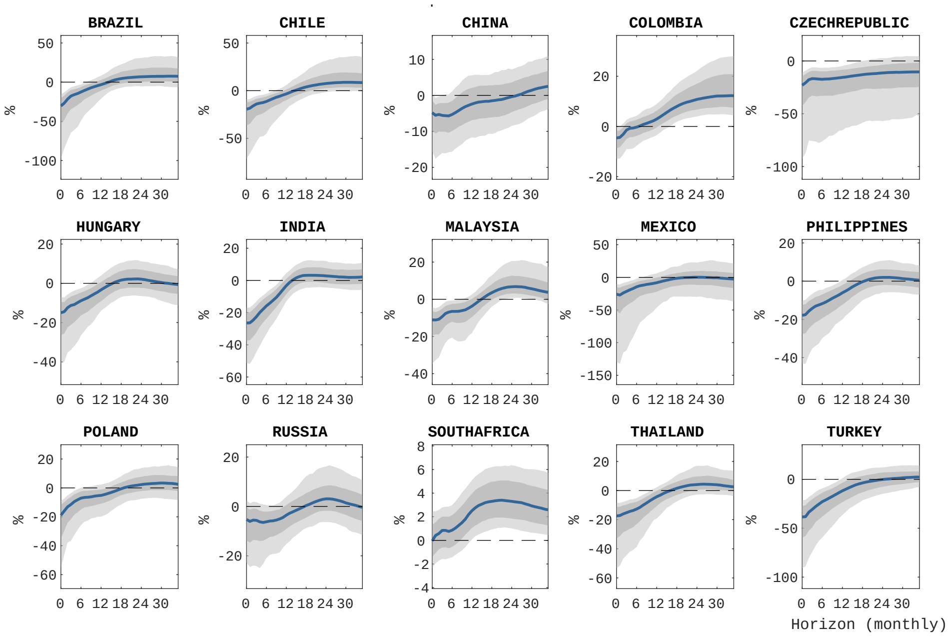

Note: Responses to an oil supply news shock normalised to induce a 10% increase in real oil price. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1.