Identification of Expectational Shocks in the Oil Market using OPEC Announcements

Riccardo Degasperi

June 2023 No: 1464

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

Identification of Expectational Shocks in the Oil Market using OPEC Announcements

Riccardo Degasperi*
University of Warwick

This version: 13 November 2021
Click here for the latest version

Abstract

Surprises in the price of oil futures computed on the day of OPEC announcements have been used as an exogenous measure of shifts in market beliefs about future oil supply to identify shocks to oil supply expectations. I show that these surprises capture not only revisions in market expectations about oil supply, but also revisions in expectations about oil demand. This conflation of supply and demand components invalidates the use of the surprises as an exogenous measure of shocks to oil supply expectations. I show that imposing an additional restriction on the sign of the co-movement between surprises in oil futures and changes in stock prices within the day of the OPEC announcement disentangles the two underlying shocks. Accordingly, I derive two robust instruments for the identification of shocks to oil supply and demand expectations that combine the surprises with this additional sign restriction, and I test them on a set of empirical specifications modelling the oil market and the global economy. A negative shock to oil supply expectations has deep and long-lasting stagflationary effects on global economic conditions that are stronger and more immediate than previously reported. These effects represent a challenge for monetary authorities that seek to stabilise both prices and output. I show that information effects are a prominent feature of the oil market and need to be accounted for when estimating the effects of shocks to oil supply expectations.

Keywords: Oil supply expectations, Information frictions, OPEC announcements, High-frequency identification, External instruments, International transmission

JEL Classification: C3, E3, Q4

*Department of Economics, University of Warwick, Social Sciences Building, Coventry, West Midlands CV4 7AL, UK. Email: R.Degasperi@warwick.ac.uk Web: www.riccardo-degasperi.com

I am indebted to Giovanni Ricco and Ivan Petrella for their invaluable guidance. I also would like to thank Dan Bernhardt, Apurav Yash Bhatiya, Christine Braun, Natalie Chen, Anna Cieslak, Roger Farmer, James Fenske, Simon Hong, Matteo Iacoviello, Diego Känzig, Dennis Novy, Roberto Pancrazi, Claudia Rei, Federico Rossi, Fabrizio Venditti, Marija Vukotić, and all the participants to the Warwick Macro & International Workshop. All errors are my own.

1 Introduction

How oil price fluctuations affect the global economy is an important question that has kept economists and policymakers busy for at least the past 40 years (Hamilton, 1983; Baumeister and Kilian, 2016). Quantifying these effects is challenging because oil prices and oil price expectations are endogenous and respond to both supply and demand conditions. To analyse the effect of movements in oil prices, one has to identify the shocks that initiate them (Kilian, 2009). Oil supply shocks are of particular interest because they depress economic activity and increase prices, posing a serious challenge to monetary authorities seeking to contemporaneously stabilise prices and output (Darby, 1982; Bernanke et al., 1997). 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). Shocks to oil supply expectations can have powerful effects on global economic activity so it is important to understand their transmission and quantify their magnitude.

Announcements of the Organization of the Petroleum Exporting Countries (OPEC) about their output quota decisions can help solve the identification problem, in that they allow us to isolate unexpected variation in future supply conditions. The role of OPEC as the key player in the global oil market means that OPEC decisions 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).[1] Consequently, 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 a shock to oil supply expectations. Of course, OPEC decisions on production quotas are themselves a response to global economic conditions. However, under some conditions, the surprises can be interpreted as revisions in oil spot price expectations that are a direct consequence of changes in markets' beliefs about oil supply.

[1]OPEC produces approximately 40% of the world's total crude oil and retains 80% of global crude oil reserves (OPEC, 2021). Despite a recent increase in the 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 99.7 mmb/d in 2019, and is expected to grow to 104.1 mmb/d in 2026 (IEA, 2021).

The conditions under which this is true are three. First, the event window within which the surprises are computed is narrow enough that contamination from other shocks is reduced to a minimum. Second, OPEC and the markets share the same information and promptly price all available information in their decisions. Third, risk premia do not change within the event window.

I argue that the second condition, in general, does not hold: OPEC has an information advantage relative to the markets regarding future oil supply. As a consequence, OPEC announcements not only reveal deviations from expected oil supply, but also clarify markets' expectations of oil demand, and might additionally 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 oil demand, invalidating the identification assumption. I show this point formally with a model of information frictions in the oil market. The intuition is the following. Assume that markets do not perfectly observe the strength of the economy. Moreover, assume that OPEC possesses private information about future oil supply that is relevant to the pricing decisions of the markets. Financial markets use the price of oil as a public signal to forecast the state of the economy, where potential deviations from expected oil supply act as noise (Hu and Xiong, 2013; Venditti and Veronese, 2020). Since markets do not know whether a movement in oil price is due to noise or to fundamentals, in equilibrium, they place less weight on the forecasted component than they would under full information (Coibion and Gorodnichenko, 2015). An OPEC announcement, by revealing information about oil supply conditions, reduces the noise in the public signal, causing a revision in the weight that markets place on the forecasted component. This is captured in the oil price revision and gives rise to the identification problem. Any identification strategy based on surprises in oil futures alone would then conflate shocks to expectations of oil supply *and* demand, and cannot be used to inform policy. Indeed, the dynamic responses obtained with this identification strategy present various puzzles, as the underlying shocks have opposite effects on macroeconomic aggregates.

Importantly, for this result to hold true, OPEC does not need to possess an information advantage over the markets about demand conditions. 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. Conversely, it is reasonable to assume that OPEC knows more than the markets about oil supply because, ultimately, OPEC decides how much crude oil to produce. For the markets, it is difficult to forecast OPEC production decisions for at least two reasons. First, production quotas are the outcome of closed-doors 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 somewhat be gauged by the comments and other forms of 'soft information' released after the conferences.

I propose a method to disentangle the demand and supply components in the surprises to obtain 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 and stock prices in a narrow window around OPEC announcements. This co-movement is informative because a shock to oil demand expectations moves both oil futures and stock prices in the same direction, while a shock to oil supply expectations moves them in opposite directions. The sign of the co-movement enables a clean identification of the shocks. Using this additional restriction on the sign of the co-movement it is possible to obtain two robust high-frequency instruments: one to identify shocks to oil supply expectations and one for shocks to oil demand expectations. In practice, I separate the surprises in oil futures into two based on their co-movement with stock prices. The instrument for shocks to oil supply expectations consists of only those surprises that induce a negative co-movement with stock prices. The instrument for shocks to oil demand expectations uses the surprises that induce a positive co-movement with stock prices.

Figure 1 offers a preview of the main results. It shows the impulse responses of world industrial production to a shock to oil supply expectations identified using the daily surprises in the price of oil futures on OPEC announcement days as an instrument (left panels) and the two shocks to oil supply and oil demand expectations identified using the robust instruments developed in this paper (right panels). The responses are obtained from a six-variable VAR of the global oil market estimated on the sample

Figure 1: Comparison between alternative identification strategies

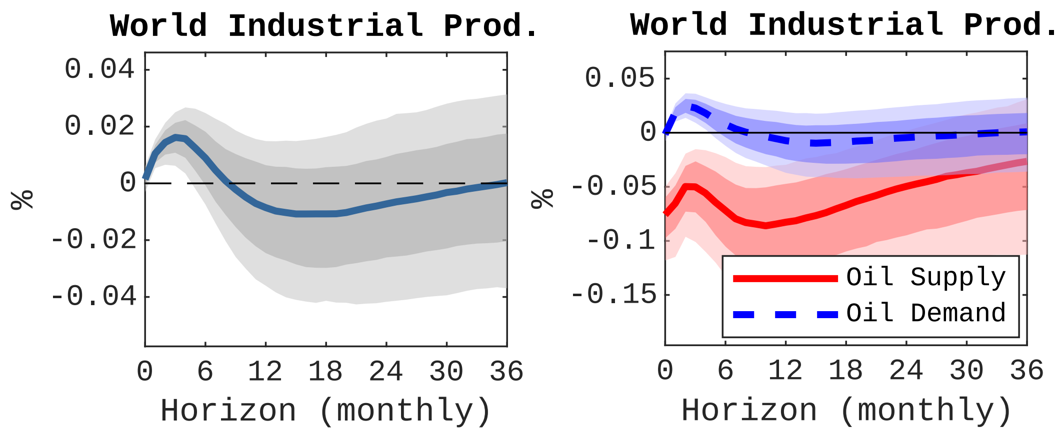

Note: Left panel – Impulse responses of world industrial production to a shock to oil supply expectations identified using daily surprises in oil futures on OPEC conference days. Right panel – Solid red: responses to a shock to oil supply expectations identified using the robust proxy for oil supply shocks. Dashed blue: responses to a shock to oil demand expectations identified using the robust proxy for demand shocks. All shocks are normalised to induce a 1% increase in real oil price. 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: 1984:1–2019:12. Both proxies span the period 1984:1–2021:1.

1984:1–2019:12. In both cases, shocks are normalised to induce a 1% increase in the real oil price.

A shock to oil supply expectations that increases oil prices is supposed to have a contractionary effect on industrial production. However, when using the surprises as an instrument, world industrial production rises for at least 6 months after the shock hits and never contracts. This is a puzzling result and exemplifies the issues that emerge when basing identification on the surprises in oil futures alone. On the other hand, the robust instruments based on the high-frequency co-movement of oil futures and stock prices give predictions that are consistent with theory (right panels). A shock to oil supply expectations that raises the price of oil has an unambiguous contractionary effect on world industrial production. World industrial production contracts by 0.1% following a shock to oil supply expectations normalised to increase real oil price by 1%. Conversely, a shock to oil demand expectations that raises the price of oil causes 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, I find that the effects of shocks to oil supply expectations 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 a shock to oil supply expectations that raises the price of oil, stock markets contract, financial uncertainty grows, and credit spreads widen. The credit channel acts as an amplification mechanism for the effects of the shock on the real economy.

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).[2] This is an important contribution because it shows that the expectational component of oil price movements can have macroeconomic effects as important as actual shocks to oil supply.

I show that surprises in oil futures cannot be used 'as is' to identify expectational shocks and propose a way of separating 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 one exploits the high-frequency co-movement between high-frequency surprises 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 information set of the central bank

[2] 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.

(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.

My paper fits in the large and well-established literature that studies the macroeconomic effects of oil price shocks. One key takeaway from this literature is that to determine the effects of fluctuations in oil prices, the underlying sources of the price movement have to be identified (Kilian, 2009). A heated point of debate is about 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 (Kilian and Park, 2009; Kilian and Murphy, 2014; Bjørnland et al., 2018; Baumeister and Hamilton, 2019; Caldara et al., 2019). Recent 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. Venditti and Veronese (2020) 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. 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 oil demand. I also provide the rationale on which the identification strategy proposed in this paper is based. In Section 3 I describe the methodology adopted to disentangle the two shocks to the expectations of oil supply and demand. 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 shock to oil supply expectations are sample-dependent and give rise to output puzzles. Second, the robust instruments developed to identify the two shocks to oil supply and oil demand expectations solve the puzzles. Third, an analysis of the transmission of the shocks to the global economy and to advanced and emerging economies shows that the shocks have powerful effects on the global economy. In Section 5 I provide important 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 oil price that follow an OPEC announcement contain both supply and demand components and therefore cannot be used as such to identify the effects of a shock to oil supply expectations on the economy. In the model, agents observe the price of oil, but do not know with certainty what is the contribution of oil supply shocks and aggregate productivity to determine the price. They only observe a noisy signal of economic conditions and only know the distribution of oil supply shocks. As a consequence, in equilibrium they place less weight on the productivity component of oil price than they would under full information.[3] The OPEC announcement reveals information about the oil supply shock and, indirectly, about economic conditions. This is captured in the oil price revision and gives rise to the identification problem. An upward price revision can be due to the revelation of economic conditions or to a negative oil supply shock.

The model also provides an additional restriction that allows us to disentangle the two underlying shocks. After the announcement, the revision in the aggregate demand for oil moves in the same direction of the price revision following an oil demand shock, but the two revisions move in opposite directions following an oil supply shock. The aggregate demand for oil determines how much agents produce and can be proxied in

[3] This is in line with the standard result in the literature on information rigidities. Since agents do not know whether the movement in a variable they are forecasting is due to noise or to fundamentals, the average expectation adjusts by less than the movement in the forecasted variable (Coibion and Gorodnichenko, 2015).

the empirical exercise by a stock price index or a daily indicator of economic activity.[4] This suggests an identification strategy based on the sign of the co-movement of oil price and stock price revisions. A positive co-movement identifies the demand shock, while a negative co-movement identifies the oil supply shock.

2.1 Model Setting

The model builds on the informational frictions model of commodity markets of Sockin and Xiong (2015). 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. The number of islands is normalised to 1. 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 according to a profit-maximising rule.

The representative agent on island i has preference over both foreign and domestic goods. Her utility function is given by

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

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

Representative firms on all islands produce according to a decreasing-returns-to-scale

[4] To satisfy the identification assumption, the indicator of aggregate production must be measured at the same frequency of the event window. In this case, the frequency is daily. Ideally, one would like to use a daily index of industrial production. But industrial production is usually measured at monthly frequency. Therefore, I assume that a general stock price index is a good representation of daily movements in industrial production. For a detailed discussion, see Section 3.

production function that uses oil as input:

$$Y_i = AX_i^{\phi}, \tag{2}$$

where $\phi \in (0, 1]$, A is the productivity that is common to all islands, and X_i is the oil input.[5] Productivity A is a random variable with a log-normal distribution,

$$\log A \sim \mathcal{N}\left(\bar{a}, \tau_A^{-1}\right). \tag{3}$$

Given the complementarity in consumption, the demand for the domestic good depends on the production of its 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 economy. Ultimately, A determines oil demand.

Importantly, producers observe A only after production has taken place at $t = 2$ and only know the parameters of its distribution. Moreover, at $t = 1$, producers receive a private signal about A,

$$s_i = \log A + \varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}\left(0, \tau_s^{-1}\right), \tag{4}$$

that they use 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. Producers also use the price of oil, which is always observable, as a public signal to form their expectations about aggregate productivity.

OPEC supplies oil according to

$$\log X_S = k \log P_X + \xi, \qquad \xi \sim \mathcal{N}\left(\bar{\xi}, \tau_{\xi}^{-1}\right), \tag{5}$$

where P_X is the price of oil, ξ represents unexpected deviations from the supply rule (i.e. ξ is an oil supply shock), and k is a positive scalar.[6] OPEC observes ξ while the producers only know the parameters of its distribution.

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

[6]OPEC's supply rule can be easily micro-founded by assuming convex labour costs (see Sockin and Xiong, 2015).

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. One can derive the equilibrium conditions for the oil market at $t = 1$. Oil price is a log-linear function of productivity A and the oil supply shock ξ:

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

where $h_A > 0$ and $h_\xi > 0$. It is clear from Eq. (6) that, from the perspective of the producers, the oil price serves as a public signal for $\log A$ and the supply shock 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. 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 Appendix, Section A.

Now we can focus on what happens when an OPEC announcement reveals information about oil supply. Assume that the OPEC announcement reveals the oil supply shock 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 information. 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.

First, observe that the revision in the price of oil that follows the announcement is given by

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

where the coefficients on aggregate productivity $\log A$ and the oil supply shock ξ are

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

Eq. (11) shows that the revision in the price of oil following the announcement can be due both to an oil supply shock and an oil demand shock. Therefore, surprises in the price of oil cannot be used as an exogenous measure of shocks to oil supply expectations.

Second, observe that the model provides an additional restriction that can be used to disentangle the shocks. The revision in aggregate demand for oil is given by

$$\log X_S' - \log X_S = \tilde{l}_A \log A + \tilde{l}_\xi \xi + \tilde{l}_0, \tag{12}$$

where the coefficients on aggregate productivity $\log A$ and the oil supply shock ξ are

$$\tilde{l}_A = \frac{k}{1+k\,(1-\phi)} - kh_A > 0,$$
$$\tilde{l}_\xi = \frac{-k\,(1-\phi)}{1+k\,(1-\phi)} - kh_\xi > 0.$$

As seen from Eq. (11), a positive oil demand shock and a negative oil supply shock move the oil price in the same direction. However, they move the aggregate demand for oil in opposite directions, as seen from Eq. (12). This forms 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 acts as a proxy for production decisions – identifies the oil demand shock, while a negative co-movement identifies the oil supply shock.

3 Methodology and Data

I estimate the impulse responses of a wide set of global macroeconomic aggregates to shocks to oil supply and demand expectations. I consider three separate 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). Obtaining exogenous instruments for the identification of the shocks to oil supply and demand expectations 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 high-frequency 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 in a 30-minute window 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 oil markets 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 are available for a relatively short span, limiting the sample size.

3.1.1 Construction of the Daily Surprises in Oil Futures

The daily surprises in the price of oil futures are obtained in two steps. First, I compute the daily change in the price of West Texas Intermediate (WTI) crude futures contracts with maturity from 1 to 12 months ahead on the closing days of OPEC conferences. These contracts have been traded on the New York Mercantile Exchange (NYMEX) since March 1983. Following Känzig (2021), I use contracts with maturity from 1 month to 12 months, which are the most liquid (Alquist and Kilian, 2010). Second, I estimate the first principal component in the term structure of the daily surprises. As the price for the 12-month contract is first listed in late December 1983, the sample used in the estimation of the first principal component spans the period from January 1984 to January 2021.

The conference dates are obtained from the OPEC website. OPEC conferences bring together delegations from each member country and are held at least twice a year.[7] 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. 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 quota decisions are observed by markets on the last day of the conference.

Between January 1984 and January 2021 there were 143 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. For months when there are more than one OPEC announcements, the monthly surprise is the sum of the daily surprises. In months when there is no OPEC announcement, the monthly surprise is set to zero.

[7] OPEC was founded in 1960 by Iran, Iraq, Kuwait, Saudi Arabia, and Venezuela. Currently, it counts 13 members. In addition to the founding members, the following countries are part of OPEC accession date in parentheses): Algeria (1969), Angola (2007), 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), and Ecuador (1973–1992; 2007–2020).

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 stock prices identifies a shock to oil demand expectations, while a negative co-movement identifies a shock to oil supply expectations. 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 shock to oil supply expectations. The revisions that induce a surprise in the stock price index of the same sign are stored in the instrument for the shock to oil demand expectations. Of the 143 announcement days that occur between 01/01/1984 and 05/01/2021, 69 are classified as events revealing the shock to oil supply expectations, and the remaining 74 are classified as events revealing the shock to oil demand expectations. The surprises are then aggregated at monthly frequency as follows. For both instruments, in months when there is 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 when there is 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 FTSE 100 (an index of the 100 companies listed on the London Stock Exchange by capitalisation), 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).[8] 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 shocks to oil supply expectations is represented by the blue bars. The instrument for shocks to oil demand expectations 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. 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

[8] The S&P 500, DS World, and DS Airlines are sourced from Datastream. All other stock price series are sourced from Bloomberg.

Figure 2: Robust instruments

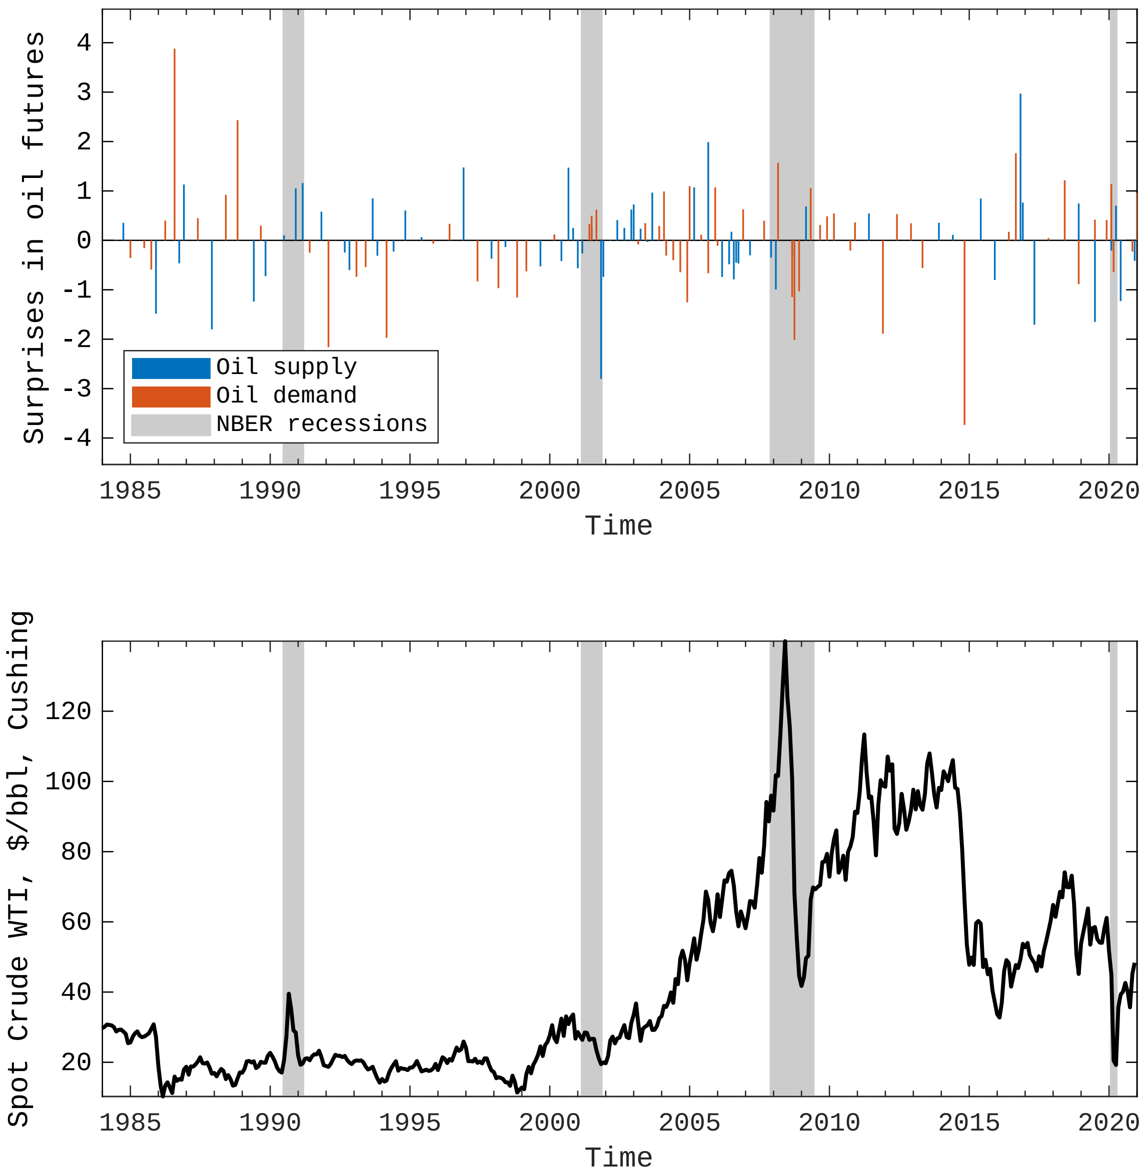

Note: Upper panel: the robust instruments for shocks to oil supply expectations (blue) and to oil demand expectations (red) obtained by separating the surprises in oil futures on the basis of their co-movement with the daily changes in stock prices. Lower panel: nominal spot crude WTI price in dollars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

than the density in non-announcement days.[9] 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%. Non-

[9]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.

Figure 3: Surprises on announcement and non-announcement days

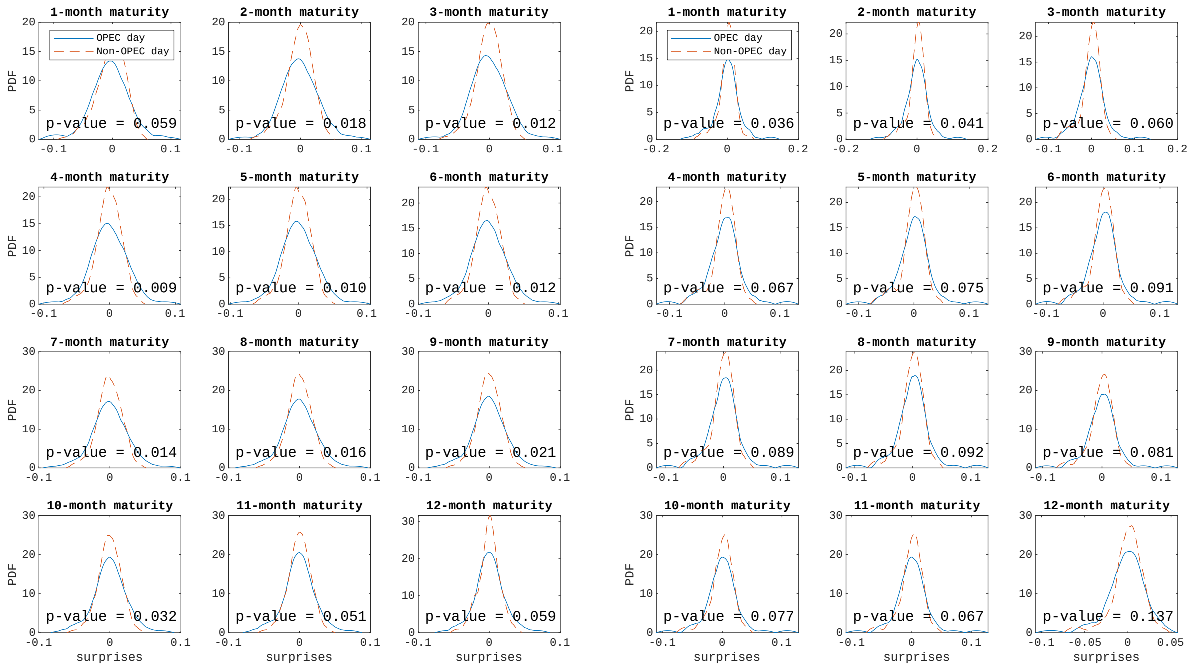

(a) Negative co-movement (b) Positive co-movement

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. Announcement and non-announcement groups contain the same number of observations.

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.

3.2 Identification via 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). 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 $\boldsymbol{e}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ 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 $\boldsymbol{u}_t$.[10] 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 instrumented variable (in this application, the real price of oil) 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:

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

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

[10] 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 and potentially controversial identifying restrictions (see for instance Giacomini et al., 2021).

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 u_t^1 induces a unitary response in the instrumented variable. 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 (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 of variable j in equation i at lag ℓ and δ_i is either 1 for variables in levels or 0 for rates.[11] 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.

[11] This reflects the idea that variables characterised by high mean reversion are best represented by a white noise process rather than a random walk.

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 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 in oil supply and demand expectations 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 is in line with 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). In the case of advanced economies, the model includes 12 endogenous variables, as described in the data section below. The variables included are 11 in the case of emerging markets, as core CPI is not available for some countries in the sample. The confidence regions for the impulse responses of each model are estimated using a standard Gibbs sampler.

Second, for each country-specific variable in the VAR, the structural impulse responses are stacked across countries and the median across countries at each horizon is computed. For global variables, the structural impulse responses are simply stacked 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% confidence 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.

Six-variable VAR. 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 1% 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 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 (excluding the U.S.) and for emerging markets, and a stock price index for OECD economies. The system also includes two main bilateral exchange rates (Euro and Pound Sterling per U.S. Dollar) and two commodity price indices (one provided by Refinitiv and the other by the Commodity Research Bureau, now *cmdty* by Barchart). Finally, five variables for the U.S. economy are included: the 1-year and 10-year constant maturity treasury rates, S&P 500, excess bond premium by Gilchrist and Zakrajšek (2012), and the CBOE VXO index (that prior to 1986:1 is reconstructed following Bloom, 2009). The excess bond premium is a measure of corporate credit spreads and captures risk appetite in the corporate bond market. The VXO is a measure of volatility in the S&P 100 and captures uncertainty in the financial markets.

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.

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

This section presents the results of the paper. First, simply using the daily surprises in the price of oil futures as an instrument to identify shocks to oil supply expectations gives rise to puzzles. Second, the puzzle disappears when the two underlying shocks are identified using the robust instruments based on the high-frequency co-movement of oil futures and stock prices. A shock to oil supply expectations that raises the price of oil has an unambiguous contractionary effect on world industrial production. Conversely, a shock to oil demand expectations that raises the price of oil causes an expansion in economic activity. Third, the robust instruments can be used to identify the effects of the shocks on larger models that better capture features of the global economy. 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, major exchange rates, credit conditions, and uncertainty measures. Moreover, they induce an endogenous response by monetary policy authorities worldwide.

4.1 Price Revisions as Instrument Give Rise to Puzzles

The results obtained using the daily surprises are not robust to changes to the length of the sample. In particular, shortening the estimation sample gives rise to output puzzles by which a shock to oil supply expectations 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 shocks to the expectations of oil supply and demand. 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. A shock to oil supply expectations that increases the price of oil is supposed to reduce economic activity.

Figure 4 documents the lack of robustness to the sample size and the presence of puzzles. The left panel replicates the main results of Känzig (2021) using the extended series of surprises in oil futures.[12] The shock is normalised to induce a 1% 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 completely alters the results. The right panel shows the impulse responses to the same shock identified on the samples 1984:1–2019:12 (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 after impact. U.S. production appears to expand, but the increase is significant at the 10% level only for the sample starting in 1990. World oil inventories and U.S. CPI show an expansion, but the dynamics differ substantially from those on the longer sample. Clearly, the responses obtained on the shorter samples do not conform to the description

[12] The series of surprises used by Känzig (2021) spans the period 1983:4 to 2017:12. 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, Bayesian methods are used that in general deliver smoother confidence regions.

Figure 4: Sample dependence under non-robust identification

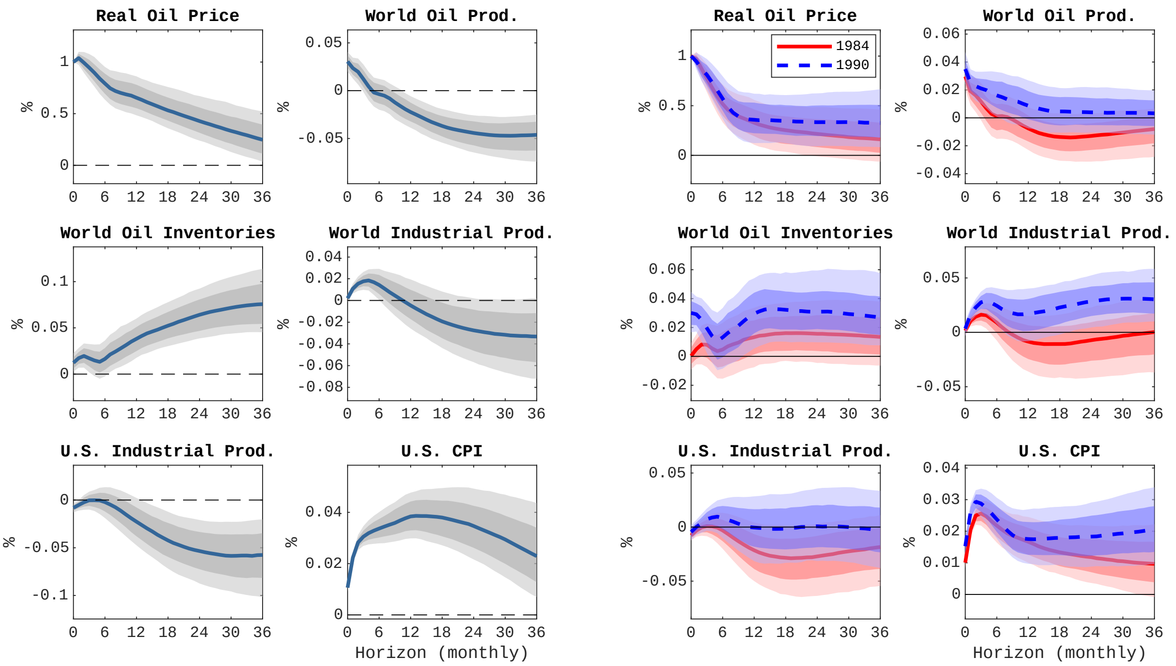

(a) Replica for Känzig (2021)

(b) Shorter samples

Note: Impulse responses to a shock to oil supply expectations normalised to induce a 1% increase in real oil price. The shock is identified using an extension of the proxy provided in Känzig (2021) which spans the period 1984:1–2021:1. Left panel: BVAR(12) estimated on the sample 1975:1–2019:12. Right panel: BVAR(12) estimated on the samples 1984:1–2019:12 (solid red) and 1990:1–2019:12 (dashed blue). Shaded areas represent 68% and 90% posterior coverage bands.

of a shock to oil supply expectations and represent a puzzle.

4.2 Puzzles Disappear using Robust Instruments

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 shock to oil supply expectations has an unambiguous contractionary effect on world industrial production. Conversely, a positive shock to oil demand expectations 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 1984:1–2020:9, but results are robust to changing the sample size, as can be seen in Figure C.1.

A shock to oil supply expectations normalised to increase the oil price by 1% causes a 0.1% contraction in world industrial production and a 0.2% contraction in U.S. industrial production on impact. This is true for all three estimation samples considered. The

Figure 5: CORRECT RESPONSES UNDER ROBUST IDENTIFICATION

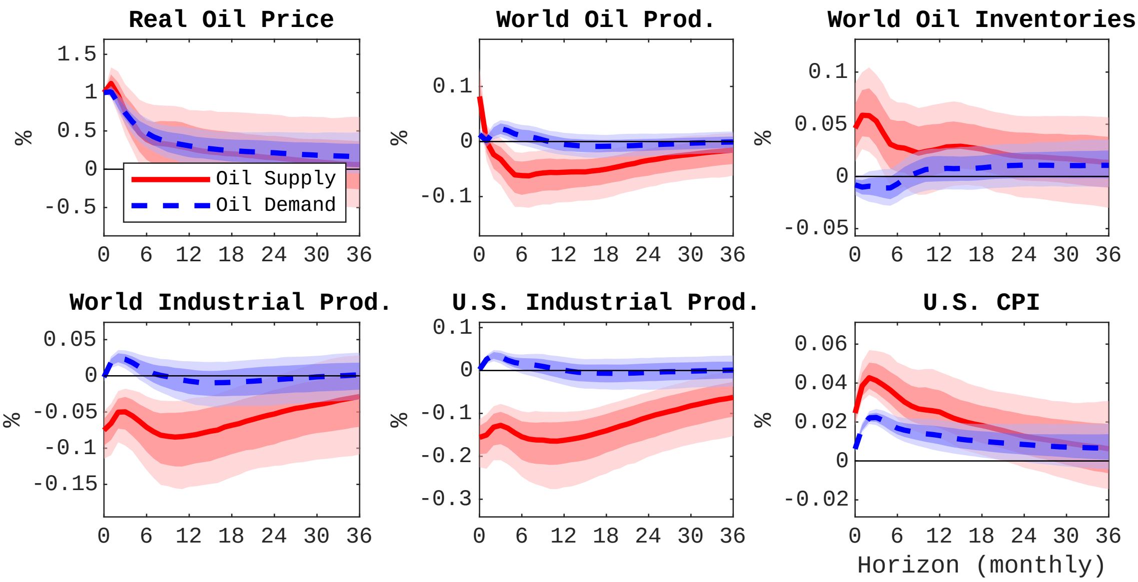

Note: Solid red: impulse responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1984:1–2020:9. Both proxies span the period 1984:1–2021:1.

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, consistently with the lag in the implementation by OPEC of the announced quota cuts. The trough is at around -0.1%, 6 months after the shock. U.S. CPI expands on impact and continues to increase, reaching a peak at +0.04% around 3 months from impact. World oil inventories also expand. In the baseline setting, the peak is reached at +0.05%, roughly 3 months from impact. However, the dynamics for this variable are slightly different across samples, which is consistent with the idea that oil inventories should be 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 pressure on economic activity.

A shock to oil demand expectations normalised to increase the oil price by 1% causes a 0.03% expansion in both world and U.S. industrial production that lasts around 6 months before reverting to trend. The effects of a shock to oil demand expectations are one order of magnitude smaller than the effects of the other shock. However, this explains why the responses identified with the surprises that mix the two shocks are unstable and puzzles emerge. World oil inventories are also contracting by 0.01%, consistently with increased demand for oil. U.S. CPI and world oil production are increasing persistently, with a peak response of roughly +0.02% 3 months since 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 a shock to oil supply expectations, and for at least 18 months for a shock to oil demand expectations.

4.3 Strength of the Instruments

The robust instrument for shocks to oil demand expectations is strong. However, the robust instrument for the shock to oil supply expectations is only weakly correlated with the instrumented variable. 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 surprises in oil futures (i.e. the instrument proposed by Känzig, 2021), (2) the robust instrument for the shock to oil supply expectations, and (3) the robust instrument for the shock to oil demand expectations. Moreover, results are presented using residuals obtained from the six-variables VAR estimated on two different samples: 1975:1–2020:9 and 1984:1–2020:9.

Stock et al. (2002) recommends a threshold for the F statistic of 10 or above to rule out weak instrument problems. The instrument for oil demand expectations is not weak according to this criterion ($F = 22.53$ on the baseline sample). However, the instrument for oil supply expectations is weak ($F = 1.527$ on the baseline sample). The relatively high F statistic obtained when using the surprises as instrument (column 1, $F = 20.30$) is consistent with the evidence that the surprises are conflating two different shocks, as contamination from other structural shocks can inflate the F statistics.

Weak instruments are problematic because they compromises 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.,

Table 1: First-stage F statistics

	1975:1–2020:9			1984:1–2020:9		
	(1)	(2)	(3)	(1)	(2)	(3)
Coefficient	3.670***	1.452	5.054***	3.442***	1.463	4.663***
N	441	441	441	429	429	429
F	22.36	1.454	25.67	20.30	1.527	22.53
F Robust	14.44	2.495	14.00	12.58	2.674	11.41

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 1984:1 respectively) and for three different instrumental variables. (1) is the non-robust proxy of Känzig (2021), extended to span 1984:1–2020:1, (2) is the robust proxy for shocks to oil supply expectations, (3) is the robust proxy for shocks to oil demand expectations. F Robust allows for heteroscedasticity. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2021).[13] Although the dynamic responses shown in Figure 5 do not fully account for the weak instrument problem, Section 5.2 implements an identification strategy that is robust to weak instruments. This identification strategy combines sign restrictions and high-frequency data following Jarociński and Karadi (2020). The results validate the impulse responses obtained using the weak instrument to identify the shock to oil supply expectations, indicating that the problem might not be so severe as to invalidate the analysis.

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 shock to oil supply expectations has strong and long-lasting contractionary effects on the global economy. Contrarily, a positive shock to oil demand expectations has a relatively short-lived expansive effect that is smaller in magnitude.

[13]The Proxy VAR algorithm employed here proceeds in two steps. First, the reduced-form VAR parameters are estimated with Bayesian methods. Second, the identification relies on a frequentist regression of the VAR innovations on the instrument. In this sense, the confidence regions of the charts shown so far do not account for the weakness of the instrument, as we are not drawing from the posterior of the parameter that represents the impact of the shock on the instrumented variable.

Figure 6: Transmission to the Global Economy

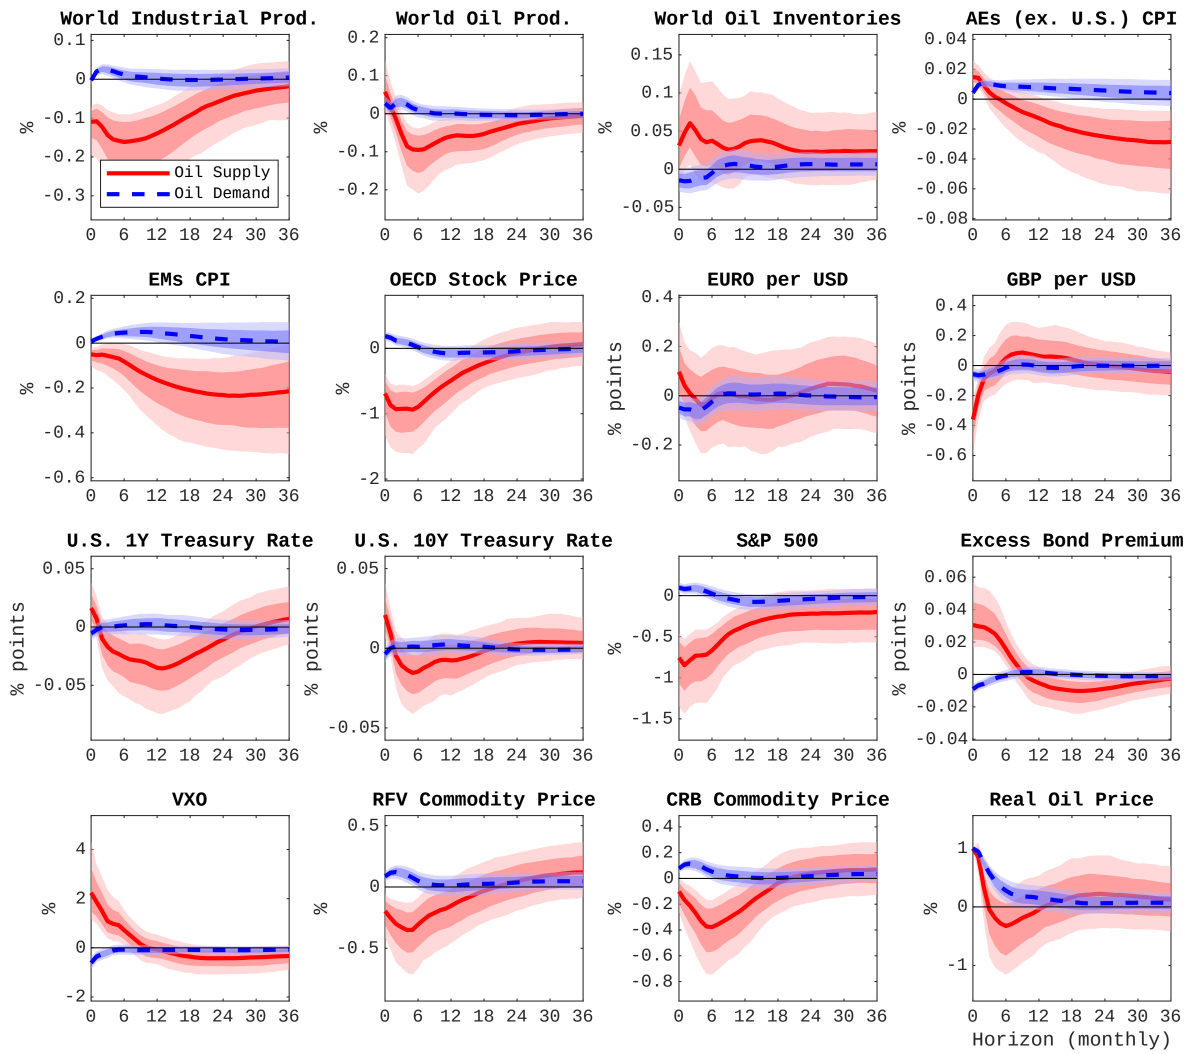

Note: Solid red: impulse responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1980:3–2020:8. Both proxies span the period 1984:1–2021:1.

Consistently with the results of the six-variable VAR, world industrial production and world oil production contract by roughly 0.1%, while oil inventories expand by 0.05%, following a shock to oil supply expectations normalised to increase real oil price by 1%. In addition to the results from the smaller VAR, the shock to oil supply expectations contracts OECD and U.S. stock prices by 1%, 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 VXO). Importantly, the shock causes a contraction in the price of other commodities, indicating that the new instrument is unlikely to be contaminated by other shocks.

The shock also appears to depress aggregate consumer price indices in both advanced (AEs) and emerging (EMs) economies. This would make sense if the contraction in economic activity was so strong to offset the inflationary effect of the shock. However, this is not consistent with the results obtained in the next two subsections and might just be an artefact of the aggregation used to construct the variables.

Following a shock to oil demand expectations that increases real oil price by 1%, 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 a news shock with a monetary easing. The monetary stimulus normally transmits along the yield curve to higher maturities, indicating that the shocks do not affect U.S. risk premia. The shock also increases the price of other commodities, consistent with an upward revision in demand expectations.[14]

4.5 Transmission to Advanced Economies

This section complements the result on global aggregates by analysing the transmission of the two shocks in a set of country-specific VARs. This approach allows us to study the effects of the shocks at a disaggregated level and compare the responses of specific groups of countries. This section will focus on advanced economies, while 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 (see Degasperi et al., 2021, for further details). The upper six variables model the median economy while the lower six variables are global controls. Table B.1 reports the sample size used for each country. The sample for each country in the analysis excludes the Covid-19 pandemic. Results are in general consistent with the ones presented in the previous section. A negative shock to oil supply expectations has a long-lasting

[14] Figure C.6 in the Appendix excludes the year 2020 to avoid the coronavirus period. The sample there spans 1984:1–2019:12. Results are virtually indistinguishable.

contractionary effect on industrial production and stock prices of the median advanced economy. The domestic exchange rate depreciates vis-à-vis the U.S. dollar. A positive shock to oil demand expectations causes an expansion in the median advanced economy. Industrial production and stock prices rise, and the exchange rate appreciates.

An important difference between the results displayed here and those for the global aggregates of the previous section is the response of CPI. Both shocks are now clearly inflationary. Also core CPI is increasing, indicating that the inflationary effect of the shock spills over to non-energy sectors. As mentioned above, the contemporaneous contraction in economic activity and higher inflation caused by a negative shock to oil supply expectations put central banks 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 the central bank faces a challenging trade-off between even higher inflation or an even deeper recession. There is an indication of loosening, but the response is not significant. Instead, in the case of a positive shock to oil demand expectations, which has a positive effect on both industrial production and prices, the direction of monetary policy is unambiguous. Indeed, the median advanced economy tightens monetary conditions.

4.6 Transmission to Emerging Markets

Figure 8 shows the median-group responses of emerging markets to the two shocks. Results are strikingly similar to those of advanced economies. A negative shock to oil supply expectations has deep and persistent recessionary effects on the median emerging market. Following a shock that induces a 1% increase in the real oil price, industrial production contracts by 0.1%, stock prices drop by 1.5%, and the currency depreciates by 0.3% vis-à-vis the U.S. dollar.

CPI tends to increase after the shock, up to 6 to 12 months later, but then reverts back to trend and turns negative. Similarly, the response of the policy rate displays an initial tightening that becomes a loosening after roughly 6 months from impact. The monetary authority of the median emerging market appears to favour price stabilisation to the detriment of economic activity. However, the response of the policy rate masks a high degree of underlying heterogeneity across countries. Heterogeneity in the response

Figure 7: Transmission to Advanced Economies

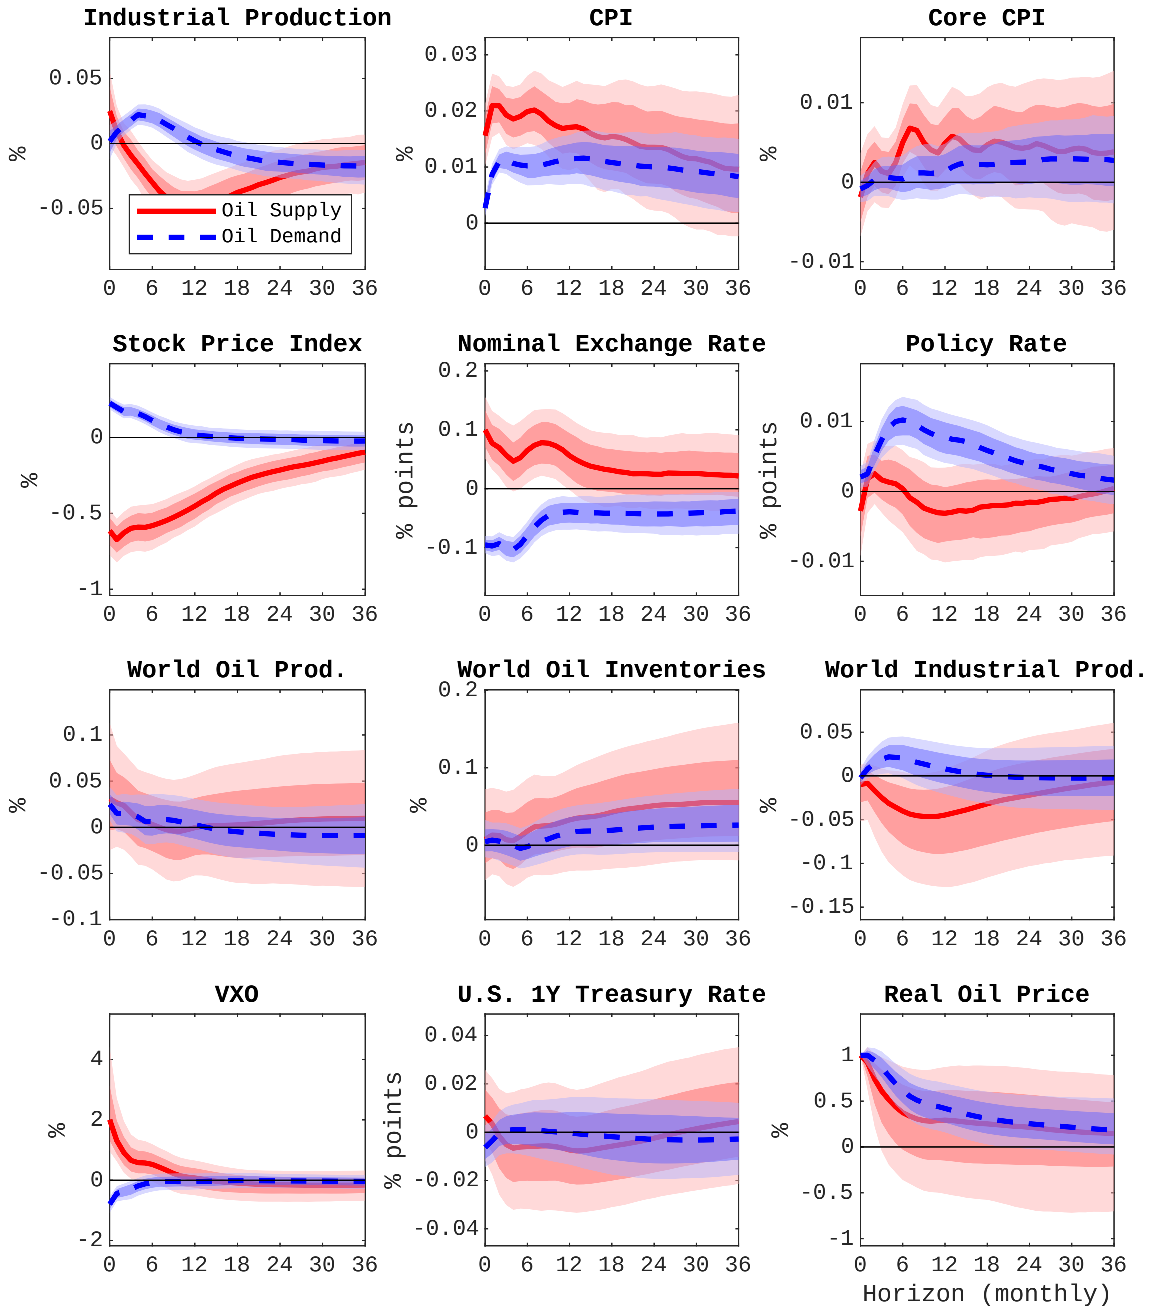

Note: Impulse responses for the median advanced economy. Solid red: responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1. Both proxies span the period 1984:1–2021:1.

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

Figure 8: Transmission to Emerging Markets

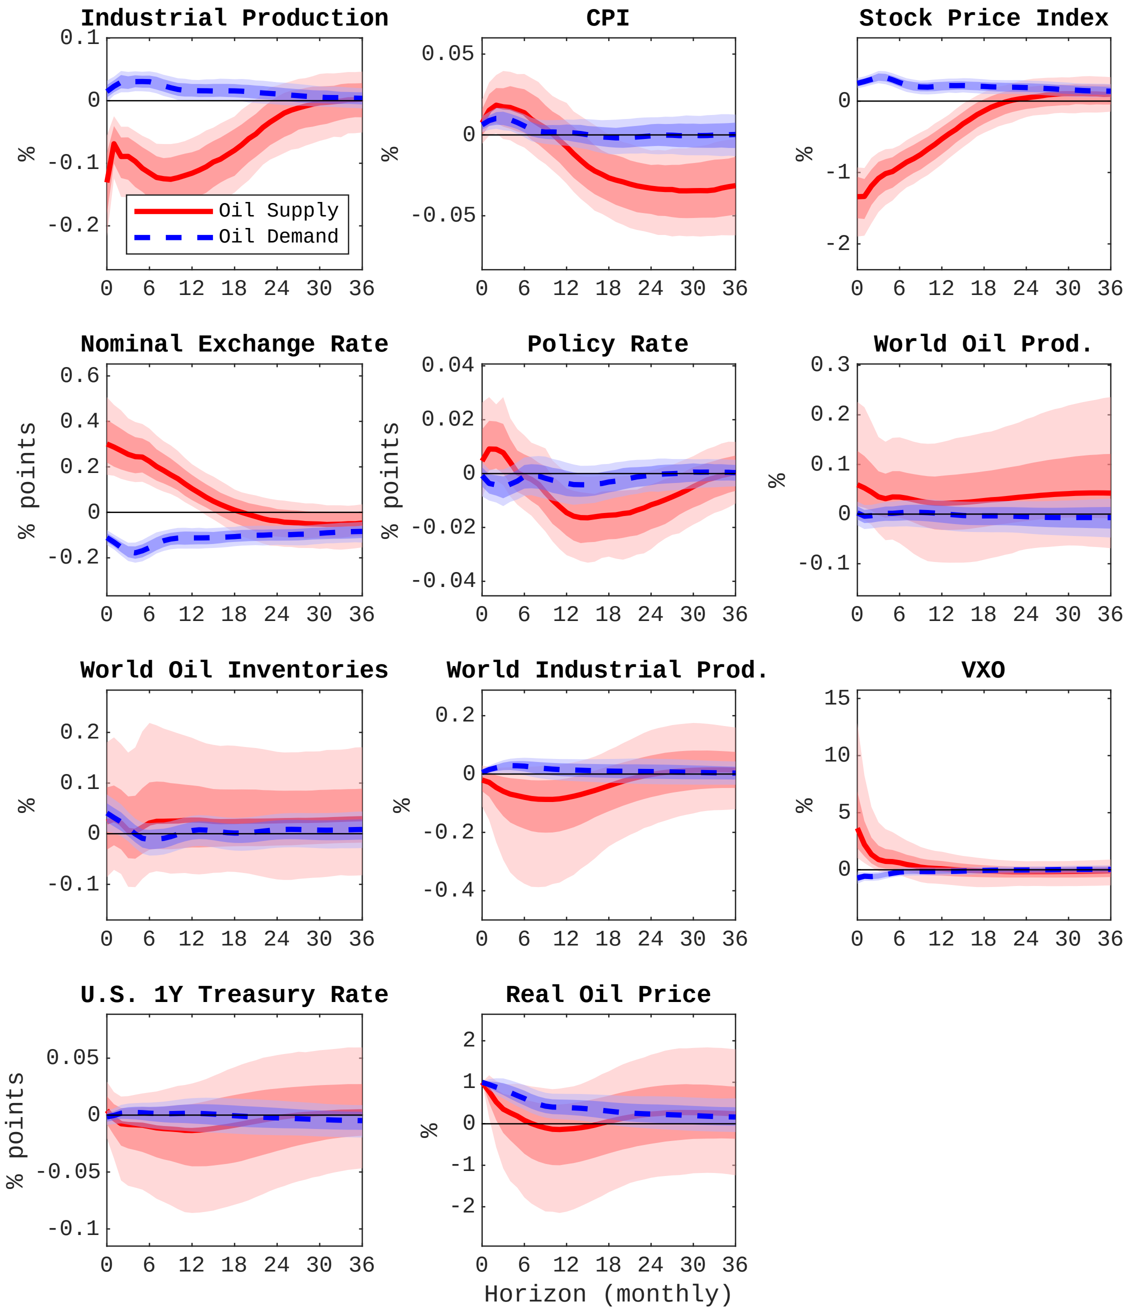

Note: Impulse responses for the median emerging market. Solid red: responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1. Both proxies span the period 1984:1–2021:1.

to a positive shock to oil demand expectations 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 insignificant.

Importantly, the responses show a remarkable degree of homogeneity across the 15 emerging markets in the sample. Following a negative shock to oil supply expectations, the response of industrial production is unambiguously negative for all countries, with the exception of China (Figure C.7). Except for South Africa, all domestic currencies depreciate against the dollar (Figure C.8). Stock prices unambiguously contract for all 15 EMs (Figure C.9). Table B.1 reports the sample size used for each country. In this case as well, the sample excludes the Covid-19 pandemic.

5 Robustness

This section provides additional support for the baseline results, showing that they are robust to a number of modifications. First, qualitatively similar results can be obtained by using an alternative methodology to separate supply and demand components in the surprises in oil futures. This method directly controls for the information set of OPEC. Second, I show that an alternative identification strategy based on sign restrictions that is robust to the weak instrument problem affecting the proxy for shocks to oil supply expectations 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 changes in a set of alternative stock price indices.

5.1 Directly Controlling for Information Effects

The second methodology to disentangle the shocks to demand and oil supply expectations consists in directly controlling for the potential informational asymmetry between OPEC and financial markets (Miranda-Agrippino and Ricco, 2021; Romer and Romer, 2000). The main idea behind this approach is that the surprises in oil futures are not only due to deviations of OPEC from its policy rule, but are also due to revisions in markets' beliefs about global economic conditions. In other words, the surprises contain information about OPEC's assessment of global economic conditions that the markets use to update their own beliefs. Conditional on finding a sufficient statistic for OPEC's information set, a regression of the surprises on this measure would separate the informational component (the fitted values) from the policy 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 concepts that I can consistently obtain for the 2001–2021 period are five: (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.[15]

This methodology consists of two steps. The first step regresses the surprises in oil futures on the information in the MOMRs at OPEC conference frequency. The residuals of this regression correlate with the shock to oil supply expectations, while the fitted part correlates with the shock to oil demand expectations. The model is the following:

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

where the subscript d indexes the day of the announcement, $F_d x_y$ represents the MOMR 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 to 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 to 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). This implies that revisions of expectations might be autocorrelated and might contain inform-

[15]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 since 2001 and there is no reason why markets should not have incorporated their informational content into their pricing decisions. However, MOMRs were not available publicly prior to 2001 and became easily accessible via the OPEC website only recently.

Table 2: Informational content of the Monthly Oil Market Reports

	(1)	(2)	(3)	(4)	(5)	(6)
R^2	0.283	0.047	0.130	0.074	0.014	0.015
F	12.368	3.075	2.776	2.165	4.658	0.346
p-value	0.000	0.021	0.010	0.081	0.002	0.846
N	83	83	83	83	83	83

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.

ation 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 instruments for shocks to oil supply expectations or the instrument for the shock to oil demand expectations, both at monthly frequencies, according to the dependent variables selected.[16]

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% 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 (6), the null of joint non-significance is rejected at 10%.[17]

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

[17] The significance of the individual coefficients is not important in this context, as the objective is simply maximising the fit. However, the full set of estimates is reported in Table B.3.

Table 3: Autoregressive component

	Residual	Fitted
R^2	0.165	0.362
F	1.352	17.067
p-value	0.219	0.000
N	66	66

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 results also show that informational 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 demand and oil supply components in the surprises are consistent with those obtained using the baseline identification strategy. Figure 9 displays the impulse responses to the two shocks. A negative shock to oil supply expectations contracts world and U.S. industrial production, contracts world oil production, and expands U.S. CPI. The contraction in economic activity is somewhat smaller relative to the baseline results. However, the responses support the baseline results and show that both methodologies can be effectively employed to separate the shocks in the surprises in oil futures. The estimation sample is reduced to the period from January 2000 to March 2020, to better match the length of the instruments (2002:1–2020:3). This might explain the differences in the dynamic responses between Figure 9 and Figure 5.

5.2 Identification by Sign Restrictions

As a robustness exercise, the shocks are identified using the approach combining high-frequency identification and sign restrictions proposed by Jarociński and Karadi (2020). This approach does not suffer from the weak instrument problem that affects the proxy for shocks to oil supply expectations. However, the interpretation of the results differ

Figure 9: ALTERNATIVE INFORMATION-ROBUST IDENTIFICATION

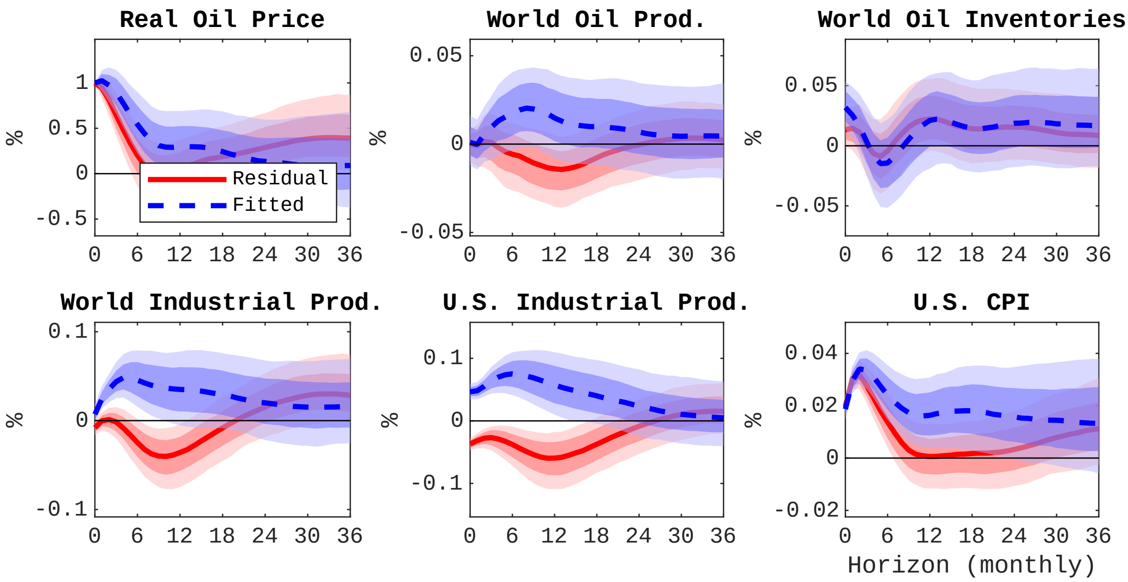

Note: Solid red: impulse responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations presented in Section 5.1. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. VAR sample: 2000:1–2020:3. IV sample: 2002:1–2020:3.

from the baseline case in a fundamental way. Identification based on sign restriction set identifies all models consistent with the restrictions. The confidence bands obtained with this method therefore cannot be thought of as representing solely the uncertainty about parameter estimates.

Nonetheless, the evidence from this alternative identification strategy matches the baseline results. Figures 10a and 10b report the responses of the six-variable VAR to the shocks to oil supply and demand expectations, respectively. World and U.S. industrial production contract by 0.1% and 0.2%, respectively, following a negative shock to oil supply expectations that increases real oil price by roughly 1%. World oil production appears to contract, although the bands include zero. Contrarily, U.S. CPI and world oil inventories do not respond.

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

Figure 10: Identification by sign restrictions

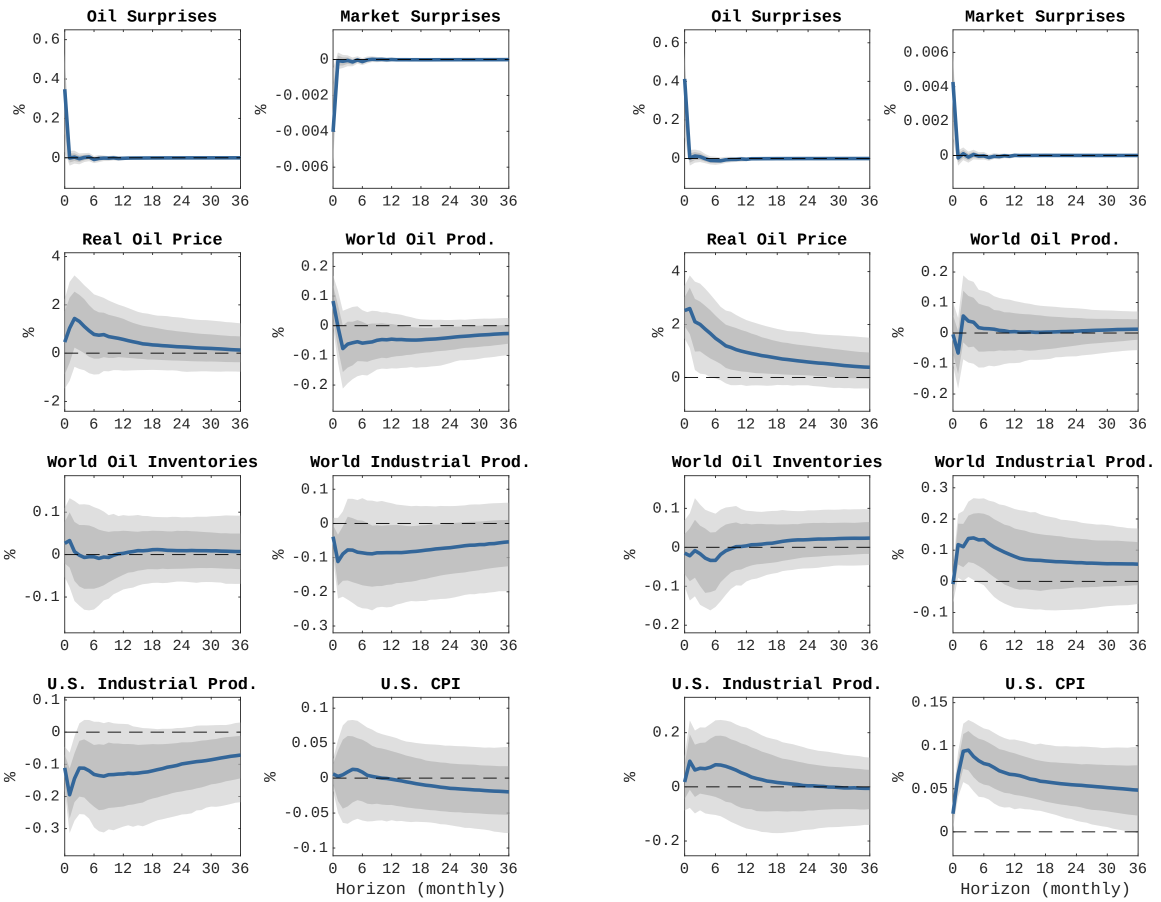

(a) Shock to oil supply expectations (b) Shock to oil demand expectations

Note: Left: impulse responses to a shock to oil supply expectations. The shock moves daily surprises in oil futures prices and daily changes in the S&P 500 on OPEC conference days in opposite directions. Right: responses to a shock to oil demand expectations. The shock moves daily surprises in oil futures prices and daily changes in the S&P 500 on OPEC conference days in the same direction. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1984:1–2020:9.

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 FTSE 100, the TOPIX, and the KOSPI. A quick look at Figure 11 shows that using alternative stock price indices does not alter the results.

Figure 11: Alternative stock price indices

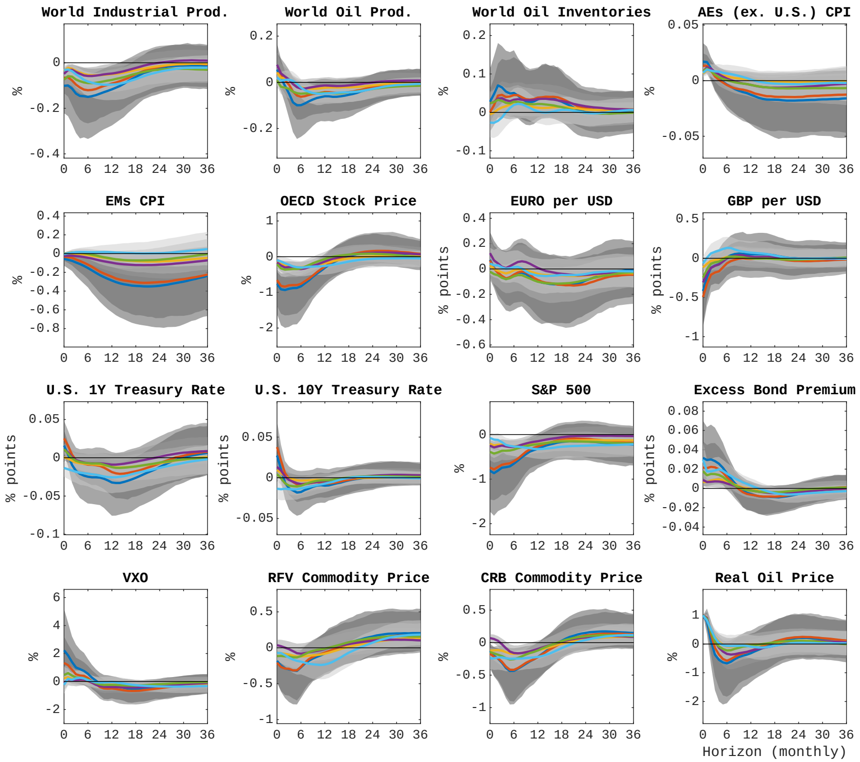

(a) Shock to oil supply expectations

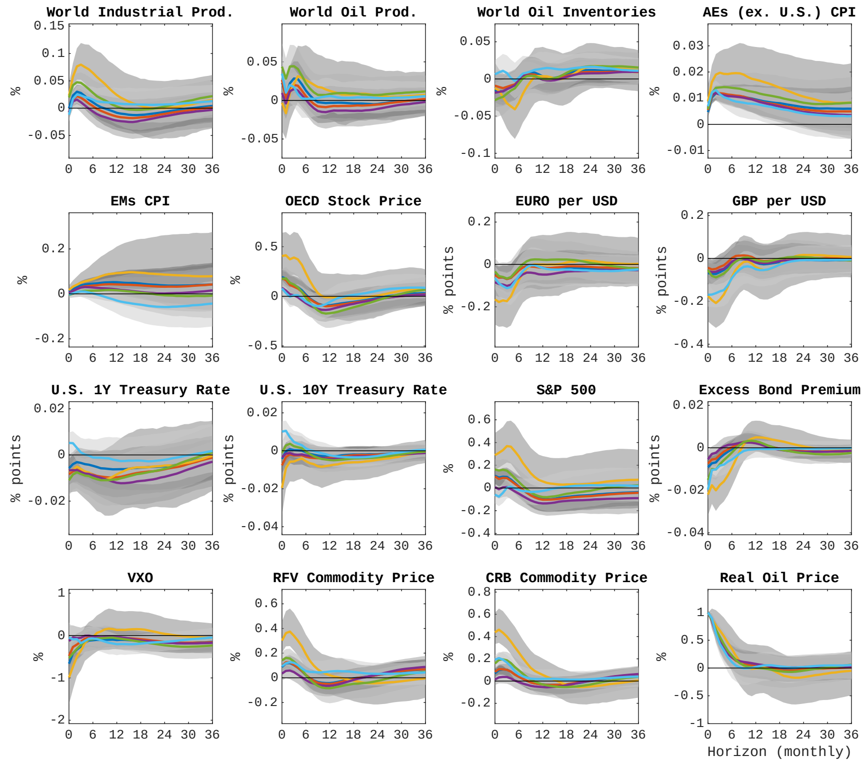

(b) Shock to oil demand expectations

Note: Left panel: impulse responses to a shock to oil supply expectations. Right panel: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% 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**: FTSE 100; **Green**: TOPIX; **Cyan**: KOSPI. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1980:3–2020:8. All proxies span the period 1984:1–2021:1.

5.4 Additional Robustness Checks

Section C in the Appendix presents some additional robustness checks. The results are summarised as follows. Results obtained with the robust instruments do not suffer from sample-dependence (Figure C.1). 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.2). Removing from the sample OPEC conferences that happened during market holidays does not alter the results (Figure C.3). Removing the most influential observations from the first-stage regression also does not alter the results (Figure C.4). Results are robust to using alternative measures of world industrial production (Figure C.5). Finally, results are robust to cutting the Covid-19 period from the sample (Figure C.6).

6 Conclusion

This paper studies the macroeconomic effects of shocks to oil supply expectations by exploiting institutional features of OPEC. It contributes to the literature in three distinct ways. First, it identifies 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 expected oil supply and expected 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 oil demand expectations, cannot be trusted. Indeed, dynamic responses obtained by identifying the shock to oil supply expectations using the surprises in oil futures alone present output puzzles. A shock to oil supply expectations that increases the price of oil appears to have an *expansionary* effect on the global economy. This is in opposition with the well-established theoretical result that a negative oil supply shock has a recessionary impact on macroeconomic conditions.

Second, the paper provides 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. This co-movement is informative because shocks to oil demand expectations move both oil futures and stock prices in the same direction, while shocks to oil supply expectations move them in opposite directions. This additional restriction on the sign of the co-movement allows to obtain two robust high-frequency instruments: one to identify shocks to oil supply expectations and one for shocks to oil demand expectations. 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 instrument the surprises in oil futures.

Third, having obtained exogenous instruments to identify the shocks to oil supply and demand expectations, the paper shows that these shocks have powerful effects on a large set of macroeconomic aggregates. A negative shock to oil supply expectations has deep and long-lasting contractionary effects on the global economy. World industrial production and U.S. industrial production contract by 0.1% and 0.2% respectively following a shock to oil supply expectations normalised to increase real oil price by 1%. CPI increases by roughly 0.03%. World oil production also contracts by roughly 0.1%, while oil inventories expand by 0.05%. The shock also contracts OECD and U.S. stock prices by 1%, 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 VXO). 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 shock to oil demand expectations, which has a positive effect on both industrial production and prices. In this instance the direction of monetary policy is unambiguous. A positive shock to oil demand expectations has a relatively short-lived and mild expansionary effect. Following a shock to oil demand expectations that increases real oil price by 1%, real activity, CPI, world oil production, and equity prices expand sharply, while credit spreads narrow and uncertainty subsides. A shock to oil demand expectations normalised to increase the oil price by 1% causes a 0.03% expansion in both world and U.S. industrial production. Instead, U.S. CPI and world oil production increase persistently, with a peak response of roughly +0.02% 3 months since the shock.

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.

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.

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 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, "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.

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.

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.

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

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.

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.

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.

_ **and Cheolbeom Park**, "The Impact of Oil Price Shocks on the U.S. Stock Market," *International Economic Review*, nov 2009, *50* (4), 1267–1287.

_ **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.

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 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.

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 2021," *Organization of the Petroleum Exporting Countries*, 2021, *56.*

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.

Venditti, Fabrizio and Giovanni Furio Veronese, "Global Financial Markets and Oil Price Shocks in Real Time," *SSRN Electronic Journal*, 2020.

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

A Model details and derivations

A.1 Market clearing conditions

An equilibrium satisfies the following market clearing conditions. At $t = 2$, in each pair of islands, markets are cleared for each good,

$$\begin{aligned} C_i + C_j^* &= AX_i^{\phi}, \\ C_i^* + C_j &= AX_j^{\phi}. \end{aligned} \tag{A.1}$$

At $t = 1$, the producers' aggregate demand for oil equals the supply,

$$\int_{-\infty}^{\infty} X_i \left(s_i, P_X\right) d\Phi \left(\varepsilon_i\right) = X_S \left(P_X\right), \tag{A.2}$$

where producers' aggregate demand for oil is integrated over the noise in their private signals.

A.2 Equilibrium in the goods market

At $t = 2$, households solve the following maximisation problem,

$$\max_{C_i, C_i^*} \left(\frac{C_i}{1-\eta}\right)^{1-\eta} \left(\frac{C_i^*}{\eta}\right)^{\eta}, \tag{A.3}$$

subject to the constraint

$$P_i C_i + P_j C_i^* = P_i Y_i. \tag{A.4}$$

The first-order conditions for the problem are:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial C_i} &= \left(\frac{C_i^*}{C_i}\right)^{\eta} \left(\frac{1-\eta}{\eta}\right)^{\eta} = \lambda_i P_i, \\ \frac{\partial \mathcal{L}}{\partial C_i^*} &= \left(\frac{C_i}{C_i^*}\right)^{1-\eta} \left(\frac{\eta}{1-\eta}\right)^{1-\eta} = \lambda_i P_j, \end{aligned}$$

which, combined, give

$$\frac{P_j}{P_i} = \frac{\eta}{1-\eta} \frac{C_i}{C_i^*}. \tag{A.5}$$

Substituting the optimality condition in the budget constraint one obtains

$$C_i = (1-\eta) Y_i, \tag{A.6}$$

and by the market clearing conditions one obtains

$$C_j^* = \eta Y_i,$$
$$C_j = (1-\eta) Y_j,$$
$$C_i^* = \eta Y_j.$$

Normalising $\lambda_i = 1$ one obtains

$$P_i = \left(\frac{C_i^*}{C_i}\right)^{\eta} \left(\frac{1-\eta}{\eta}\right)^{\eta} = \left(\frac{\eta Y_j}{(1-\eta) Y_i}\right)^{\eta} \left(\frac{1-\eta}{\eta}\right)^{\eta} = \left(\frac{Y_j}{Y_i}\right)^{\eta}, \tag{A.7}$$

which shows that the home good is more valuable when the trading partner j produces more. This feature implies the existence of strategic interactions between the producers of island i and j.

A.3 Demand for oil

The representative firm on island i solves

$$\max_{X_i} \mathbb{E}\left[P_i Y_i | \mathcal{I}_i\right] - P_X X_i, \tag{A.8}$$

which can be rewritten as

$$\max_{X_i} \mathbb{E}\left[A P_i X_i^{\phi} | s_i, P_X\right] - P_X X_i. \tag{A.9}$$

In equilibrium, marginal revenue equals marginal cost, giving the following optimality condition,

$$\phi \mathbb{E}\left[A P_i | s_i, P_X\right] X_i^{\phi-1} = P_X, \tag{A.10}$$

where we can bring X outside the expectation operator because it is a choice variable. Substituting P_i with the equilibrium price one can obtain the oil demand function of goods producer i:

$$X_i = \left\{ \frac{\phi \mathbb{E}\left[A X_j^{\phi\eta} | s_i, P_X\right]}{P_X} \right\}^{1/(1-\phi(1-\eta))}. \tag{A.11}$$

Producer i's demand for oil, which in turn determines i's production, depends on the expected production of the trading partner j.

A.4 Equilibrium in the oil market

Assume that the oil price and producer i's demand for oil take the following log-linear relationships:

$$\log P_X = h_A \log A + h_\xi \xi + h_0, \tag{A.12}$$

$$\log X_i = l_s s_i + l_P \log P_X + l_0, \tag{A.13}$$

and define the following sufficient statistic for the information contained in the oil price:

$$z \equiv \frac{\log P_X - h_\xi \bar{\xi} - h_0}{h_A} = \log A + \frac{h_\xi}{h_A}\left(\xi - \bar{\xi}\right). \tag{A.14}$$

Using the projection theorem, producer i's expectation of aggregate productivity is given by

$$\mathbb{E}\left[\log A | s_i, \log P_X\right] = \mathbb{E}\left[\log A | s_i, z\right] = \frac{1}{\frac{h_A^2}{h_\xi^2}\tau_\xi + \tau_s + \tau_A}\left[\tau_A \bar{a} + \tau_s s_i + \frac{h_A^2}{h_\xi^2}\tau_\xi z\right]. \tag{A.15}$$

The variance, by the projection theorem, is

$$\begin{aligned} Var\left[\log A | s_i, \log P_X\right] &= \mathbb{E}\left[\left(\log A - \mathbb{E}\left[\log A | s_i, \log P_X\right]\right)^2 | s_i, z\right] \\ &= \mathbb{E}\left[\left(\log A\right)^2 | s_i, z\right] - \mathbb{E}\left[\log A | s_i, z\right]^2 \\ &= \left(\frac{h_A^2}{h_\xi^2}\tau_\xi + \tau_s + \tau_A\right)^{-1}. \end{aligned}$$

Now log-transform (A.11)

$$\log X_i = \frac{1}{1 - \phi\left(1 - \eta\right)}\left\{\log \phi + \log \mathbb{E}\left[A X_j^{\phi\eta} | s_i, P_X\right] - \log P_X\right\} \tag{A.16}$$

We can rewrite

$$
\begin{aligned}
\mathbb{E}\left[AX_j^{\phi\eta}|s_i,P_X\right] &= \mathbb{E}\left[\exp\left(\log A+\phi\eta\log X_j\right)|s_i,z\right]\\
&= \mathbb{E}\left[\exp\left(\log A+\phi\eta\left(l_s s_j+l_P\log P_X+l_0\right)\right)|s_i,z\right]\\
&= \exp\left(\phi\eta\left(l_0+l_P\log P_X\right)\right)\cdot\mathbb{E}\left[\exp\left(\left(1+\phi\eta l_s\right)\log A+\phi\eta l_s\varepsilon_j\right)|s_i,z\right]\\
&= \exp\left\{\phi\eta\left(l_0+l_P\log P_X\right)+\left(1+\phi\eta l_s\right)\mathbb{E}\left[\log A|s_i,z\right]\right.\\
&\left.+\frac{1}{2}\mathbb{V}\left[\left(1+\phi\eta l_s\right)\log A+\phi\eta l_s\varepsilon_j|s_i,z\right]\right\}\\
&= \exp\left\{\phi\eta\left(l_0+l_P\log P_X\right)+\left(1+\phi\eta l_s\right)\mathbb{E}\left[\log A|s_i,z\right]+\frac{\left(1+\phi\eta l_s\right)^2}{2}\mathbb{V}\left[\log A|s_i,z\right]\right.\\
&\left.+\frac{\left(\phi\eta l_s\right)^2}{2}\mathbb{V}\left[\varepsilon_j|s_i,z\right]+\left(1+\phi\eta l_s\right)\left(\phi\eta l_s\right)\mathbb{C}\left[\log A,\varepsilon_j|s_i,z\right]\right\}\\
&= \exp\left\{\phi\eta\left(l_0+l_P\log P_X\right)+\frac{1+\phi\eta l_s}{\frac{h_A^2}{h_\xi^2}\tau_\xi+\tau_s+\tau_A}\left[\tau_A\bar{a}+\tau_s s_i+\frac{h_A^2}{h_\xi^2}\tau_\xi z\right]\right.\\
&\left.+\frac{\left(1+\phi\eta l_s\right)^2}{2}\left(\frac{h_A^2}{h_\xi^2}\tau_\xi+\tau_s+\tau_A\right)^{-1}+\frac{\left(\phi\eta l_s\right)^2}{2}\tau_s^{-1}\right\}
\end{aligned}
$$

where the fourth equality uses

$$
\mathbb{E}\left[\exp\left(aY\right)|X\right]=\exp\left(a\mathbb{E}\left[Y|X\right]+\frac{1}{2}a^2\mathbb{V}\left[Y|X\right]\right) \tag{A.17}
$$

After substituting this in (A.16), the coefficients of the expression obtained can be matched to the conjecture in (A.13). Sockin and Xiong (2015) shows how to combine it with the market clearing conditions to obtain the coefficients of (A.12) and (A.13):

$$
h_A=-\frac{(1-\phi)b+(1-\phi(1-\eta))\tau_s^{-1}\tau_\xi b^3}{1+k(1-\phi)}>0, \tag{A.18}
$$

$$
h_\xi=-\frac{1-\phi+(1-\phi(1-\eta))\tau_s^{-1}\tau_\xi b^2}{1+k(1-\phi)}<0, \tag{A.19}
$$

$$
h_0=\frac{1}{1+k(1-\phi)}\log\phi-\frac{1-\phi(1-\eta)}{1+k(1-\phi)}b\tau_s^{-1}\left(\tau_A\bar{a}-b\tau_\xi\bar{\xi}\right) \tag{A.20}
$$

$$
+\frac{1}{2}\frac{1-\phi(1-\eta)}{1+k(1-\phi)}\left(\left(\frac{1-\phi+\phi^2\eta^2}{1-\phi(1-\eta)}+\phi\eta\right)b-1\right)\tau_s^{-1}b, \tag{A.21}
$$

$$
l_s=-h_\xi^{-1}h_A>0, \tag{A.22}
$$

$$
l_P=k+h_\xi^{-1}, \tag{A.23}
$$

$$
l_0=\left(k-l_P\right)h_0-\frac{1}{2}l_s^2\tau_s^{-1}, \tag{A.24}
$$

where

$$b = \left(\frac{\tau_\xi^{-1}\tau_s}{2(1-\phi(1-\eta))}\right)^{1/3} \sqrt[3]{-1+\sqrt{1+\frac{4}{27}\left(\frac{\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}\right)^{-2}\left(\tau_A+\frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)^3}}$$
$$+\left(\frac{\tau_\xi^{-1}\tau_s}{2(1-\phi(1-\eta))}\right)^{1/3} \sqrt[3]{-1-\sqrt{1+\frac{4}{27}\left(\frac{\tau_\xi^{-1}\tau_s}{1-\phi(1-\eta)}\right)^{-2}\left(\tau_A+\frac{1-\phi}{1-\phi(1-\eta)}\tau_s\right)^3}} < 0. \tag{A.25}$$

A.5 Equilibrium under perfect information

Under perfect information each producers private signal becomes irrelevant as A is observable. Since the producers now share the same information about A, they must have the same expectations about their future trading partners' production decisions. Then

$$X_i = X_j = \left(\frac{\phi A}{P_X}\right)^{\frac{1}{1-\phi}}, \tag{A.26}$$

from which one can derive the equilibrium conditions

$$\log P_X' = \frac{1}{1+k(1-\phi)}\log A - \frac{1-\phi}{1+k(1-\phi)}\xi + \frac{1}{1+k(1-\phi)}\log\phi, \tag{A.27}$$
$$\log X_S' = \frac{k}{1+k\,(1-\phi)}\log A + \frac{1}{1+k\,(1-\phi)}\xi + \frac{k}{1+k\,(1-\phi)}\log\phi. \tag{A.28}$$

A.6 Surprises in oil price and aggregate demand for oil

The revision in the price of oil caused by the OPEC announcement, which reveals ξ and consequently turns the model into one of perfect information, is

$$\begin{aligned}\log P_X' - \log P_X &= \left[\frac{1}{1+k(1-\phi)} - h_A\right]\log A - \left[\frac{1-\phi}{1+k(1-\phi)} - h_\xi\right]\xi \\ &\quad + \frac{1}{1+k(1-\phi)}\log\phi - h_0 \\ &= \tilde{h}_A \log A - \tilde{h}_\xi \xi + \tilde{h}_0,\end{aligned}$$

where

$$\tilde{h}_A = \frac{1}{1+k(1-\phi)} - h_A > 0, \tag{A.29}$$

$$\tilde{h}_\xi = \frac{1-\phi}{1+k(1-\phi)} - h_\xi > 0. \tag{A.30}$$

The revision in the aggregate demand for oil following the announcement is given by

$$\begin{aligned}
\log X'_S - \log X_S &= \left[\frac{k}{1+k\,(1-\phi)} - (l_s + l_P h_A)\right]\log A + \left[\frac{1}{1+k\,(1-\phi)} - l_P h_\xi\right]\xi \\
&\quad + \frac{k}{1+k\,(1-\phi)} - l_0 - l_P h_0 - \frac{1}{2} l_s^2 \tau_s^{-1} \\
&= \left[\frac{k}{1+k\,(1-\phi)} - k h_A\right]\log A + \left[\frac{-k\,(1-\phi)}{1+k\,(1-\phi)} - k h_\xi\right]\xi + \frac{k}{1+k\,(1-\phi)} - k h_0 \\
&= \tilde{l}_A \log A + \tilde{l}_\xi \xi + \tilde{l}_0,
\end{aligned}$$

where

$$\tilde{l}_A = \frac{k}{1+k\,(1-\phi)} - k h_A > 0, \tag{A.31}$$

$$\tilde{l}_\xi = -\frac{k\,(1-\phi)}{1+k\,(1-\phi)} - k h_\xi > 0. \tag{A.32}$$

The inequality in (A.30) holds because $h_\xi < 0$. To see why the inequality in (A.32) holds, consider what happens when the informativeness of the private signal increases and the model tends to its perfect-information benchmark. One can show that

$$\lim_{\tau_s \to \infty} b = -\frac{1}{1-\phi},$$

which implies that both (A.18) and (A.19) converge to their perfect-information values:

$$\lim_{\tau_s \to \infty} h_A = \frac{1}{1+k\,(1-\phi)},$$

$$\lim_{\tau_s \to \infty} h_\xi = -\frac{1-\phi}{1+k\,(1-\phi)}.$$

The numerator of $|h_\xi|$ in (A.19) is positive and larger than $1-\phi$, which implies that $h_\xi < -\frac{1-\phi}{1+k(1-\phi)}$. By substituting (A.25) into (A.18) one can obtain

$$h_A = \frac{1 + \tau_A \tau_s^{-1}(1-\phi(1-\eta))b}{1+k(1-\phi)}.$$

Since $b < 0$, one must have that $h_A < \frac{1}{1+k(1-\phi)}$. This shows that the inequalities in (A.29) and (A.31) hold.

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)
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 Consumer Price Index Inflation, 2005 = 100	Dallas Fed, Global Economic Indicators		•	•		•		
EMs CPI	Headline Consumer Price Index Inflation, 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	•	•		•		
EURO per USD	Nominal bilateral exchange rate. End-of-month	BIS, USD Exchange rate		•			•		
GBP per USD	Nominal bilateral exchange rate. End-of-month	BIS, USD Exchange rate		•			•		
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)					•		
VXO	CBOE S&P 100 Volatility Index: VXO, Index, Monthly, NSA	Bloom (2009) before 1986:1; FRED	VXOCLS	•			•	•	•
RFV Commodity Price		Datastream	RECMDTY	•	•		•		
CRB Commodity Price	cmdty BLS Commodity Price Ind	Datastream	CRBSPOT	•	•		•		
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. The bottom section of the table list 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: Informational content of the Monthly Oil Market Reports

	(1)	(2)	(3)	(4)	(5)	(6)
GDP_{world}	-0.322	-0.245				
	(-0.93)	(-1.45)				
GDP_{US}	0.249	0.120				
	(0.61)	(0.74)				
GDP^{rev}_{world}	0.712	1.082*				
	(0.80)	(1.83)				
GDP^{rev}_{US}	-1.124	-0.600				
	(-1.32)	(-1.28)				
Oil $\text{Demand}^{backcast}_{world}$	1.031		0.837**			
	(1.67)		(2.43)			
Oil $\text{Demand}^{nowcast}_{world}$	-0.790		-0.819**			
	(-1.33)		(-2.44)			
Oil $\text{Demand}^{backcast}_{OECD}$	-1.662**		-1.090**			
	(-2.05)		(-2.55)			
Oil $\text{Demand}^{nowcast}_{OECD}$	1.830**		1.112**			
	(2.19)		(2.38)			
Oil $\text{Demand}^{backcast,\,rev}_{world}$	-0.847		-0.728			
	(-0.76)		(-0.99)			
Oil $\text{Demand}^{nowcast,\,rev}_{world}$	0.572		0.441			
	(0.51)		(0.59)			
Oil $\text{Demand}^{backcast,\,rev}_{OECD}$	3.432		2.537**			
	(1.67)		(2.23)			
Oil $\text{Demand}^{nowcast,\,rev}_{OECD}$	-0.942		-0.337			
	(-0.70)		(-0.40)			
Non-OPEC $\text{Supply}^{backcast}$	0.508			0.420*		
	(0.64)			(1.89)		
Non-OPEC $\text{Supply}^{nowcast}$	-0.781			-0.398*		
	(-1.53)			(-1.89)		
Non-OPEC $\text{Supply}^{backcast,\,rev}$	0.463			0.241		
	(0.51)			(0.44)		
Non-OPEC $\text{Supply}^{nowcast,\,rev}$	-0.315			0.044		
	(-0.46)			(0.10)		
OPEC $\text{Supply}^{nowcast}$	0.000				-0.000	
	(0.11)				(-0.22)	
OPEC Supply^{mom}	-0.000				0.000	
	(-0.09)				(0.67)	
OPEC $\text{Supply}^{nowcast,\,rev}$	0.000**				0.000***	
	(2.01)				(3.83)	
OPEC $\text{Supply}^{mom,\,rev}$	0.000				0.000	
	(0.75)				(0.27)	
D/S $\text{Balance}^{backcast}$	0.151					-0.116
	(0.24)					(-0.90)
D/S $\text{Balance}^{nowcast}$	-0.326					0.074
	(-1.61)					(0.62)
D/S $\text{Balance}^{backcast,\,rev}$	-0.995					0.057
	(-0.87)					(0.19)
D/S $\text{Balance}^{nowcast,\,rev}$	0.133					-0.093
	(0.48)					(-0.50)
constant	-8.192	0.577**	-1.765	-0.759	0.340	1.274
	(-1.12)	(2.06)	(-0.31)	(-0.69)	(0.23)	(1.02)
R^2	0.283	0.047	0.130	0.074	0.014	0.015
F	12.368	3.075	2.776	2.165	4.658	0.346
p-value	0.000	0.021	0.010	0.081	0.002	0.846
N	83	83	83	83	83	83

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
$l = 1$	-0.275	0.098
	(-0.68)	(0.35)
$l = 2$	0.541*	1.234***
	(1.78)	(8.36)
$l = 3$	-0.079	0.384**
	(-0.36)	(2.12)
$l = 4$	-0.199	-0.076
	(-1.30)	(-0.26)
$l = 5$	-0.171	0.112
	(-0.80)	(0.71)
$l = 6$	-0.218	0.039
	(-1.42)	(0.24)
$l = 7$	-0.282	0.273
	(-1.61)	(1.32)
$l = 8$	-0.104	-0.157
	(-0.27)	(-0.45)
$l = 9$	0.002	-0.248
	(0.01)	(-1.19)
$l = 10$	0.045	-0.113
	(0.27)	(-0.42)
$l = 11$	-0.113	0.235
	(-0.28)	(0.97)
$l = 12$	-0.324**	-0.049
	(-2.35)	(-0.27)
constant	0.012	-0.014
	(0.10)	(-0.21)
R^2	0.165	0.362
F	1.352	17.067
p-value	0.219	0.000
N	66	66

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

Figure C.1: Robust identification on alternative samples

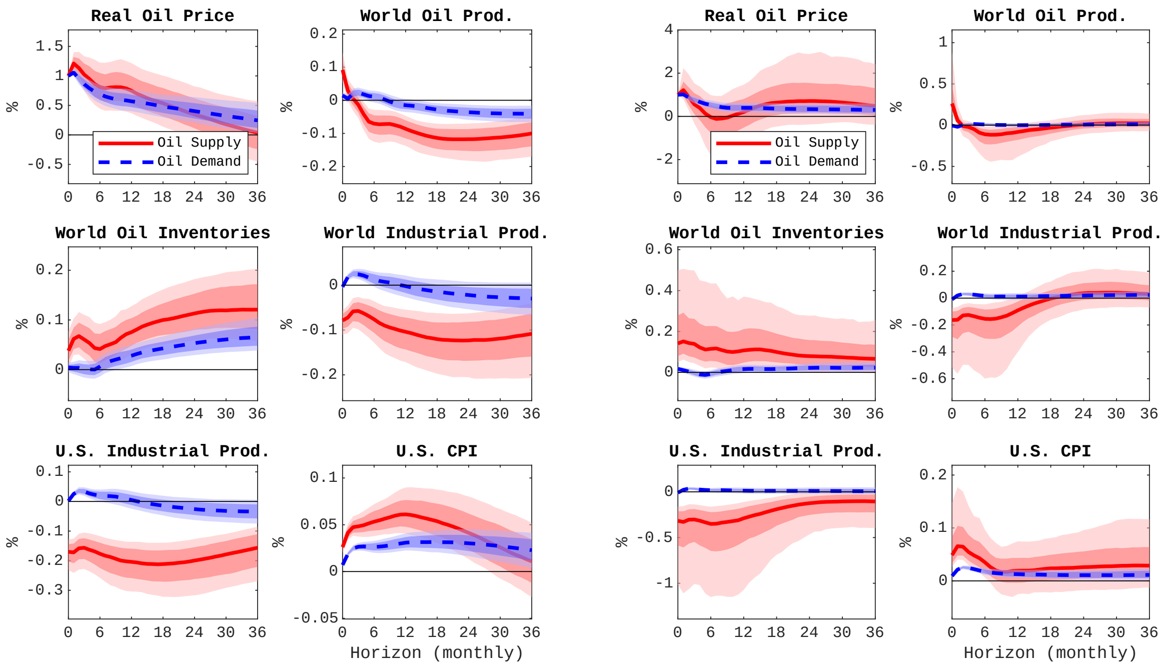

(a) Sample: 1975:1–2020:9 (b) Sample: 1990:1–2020:9

Note: Solid red: impulse responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Samples: 1975:1–2020:9 (left); 1990:1–2020:9 (right). Both proxies span the period 1984:1–2021:1.

C Additional Robustness Checks

C.1 Robust identification on different sample lengths

Figure C.1 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 1984:1–2020:9 are reported. Here it is shown that the responses obtained on the samples 1975:1–2020:9 and 1990:1–2020:9 are consistent to the baseline results.

C.2 Identification based on U.S. daily real business conditions

The high-frequency co-movement between oil futures and stock price index can be used to disentangle the shocks in the surprises in oil futures because stock prices are a high-frequency proxy for economic activity. However, they are not the only high-frequency measure of eco-

Figure C.2: Identification based on Aruoba et al. (2009)

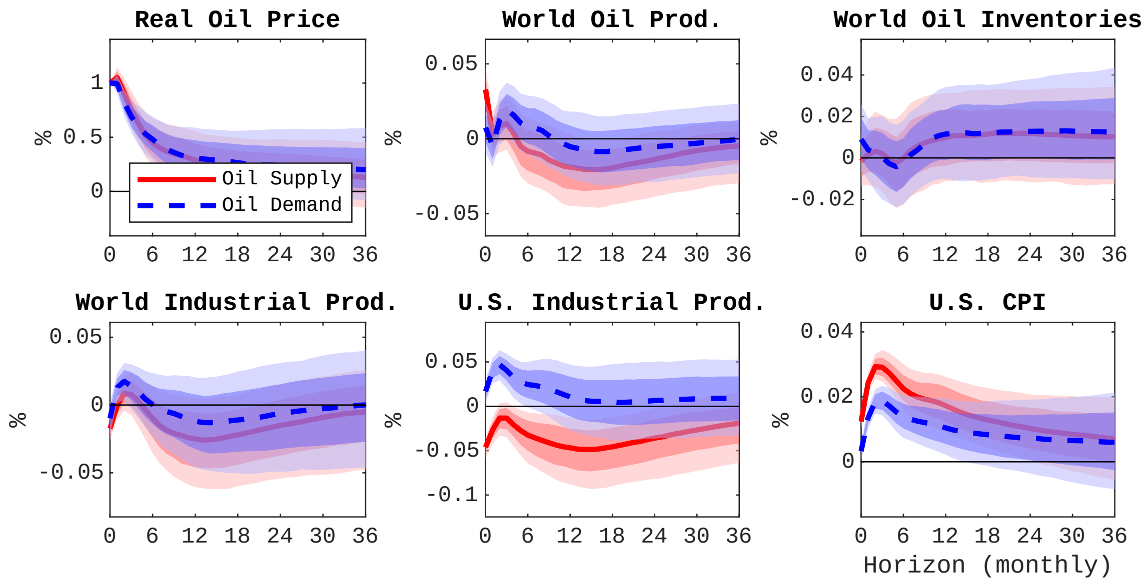

Note: Solid red: impulse responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies, constructed based on the comovement on OPEC 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: 1984:1–2020:9.

nomic activity available. Figure C.2 shows that the demand and oil supply 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 shock to oil supply expectations and expansionary for a positive shock to oil demand expectations. However, residual traces of puzzles in the response of world industrial production are still visible.

C.3 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 01/01/1984 and 05/01/2021 there are 143 announcement days. Of these, 19 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 19 observations that coincide with market holidays. The impulse responses obtained are in line with the baseline results.

Figure C.3: Removing market holidays

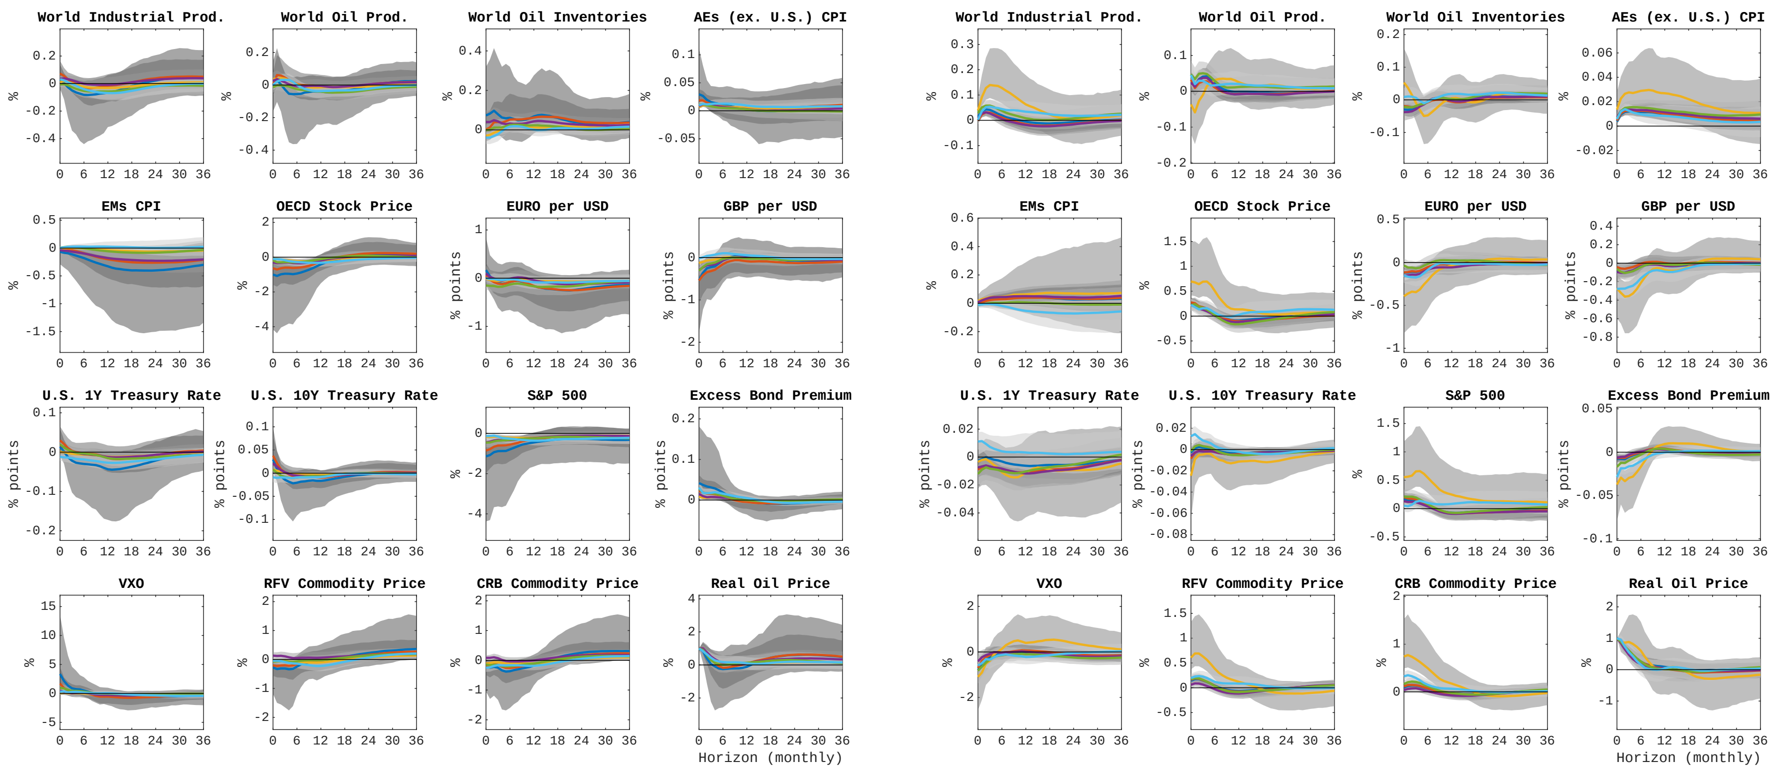

(a) Shock to oil supply expectations

(b) Shock to oil demand expectations

Note: Impulse responses obtained by dropping market holidays from the regression of the reduced-form VAR innovations on the instrument. Left panel: responses to a shock to oil supply expectations. Right panel: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% 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 11. **Blue**: S&P 500; **Red**: DS World; **Yellow**: DS Airlines; **Purple**: FTSE 100; **Green**: TOPIX; **Cyan**: KOSPI. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1980:3–2020:8. All proxies span the period 1984:1–2021:1.

C.4 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 represents 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.4 displays the responses to a negative shock to oil supply expectations (left) and to a positive shock to oil demand expectations (right) identified with the robust instruments based on the S&P 500 (solid red) and with the same instrument without the 6 MIOs. 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 are robust to the exclusion of the 6 MIO for both shocks, and while for the shock to oil supply expectations confidence regions become considerably larger, for the shock to oil demand expectations the precision of the responses does not change, and the magnitude of the expansion is actually even larger.

Figure C.4: Removing Most Influential Observations

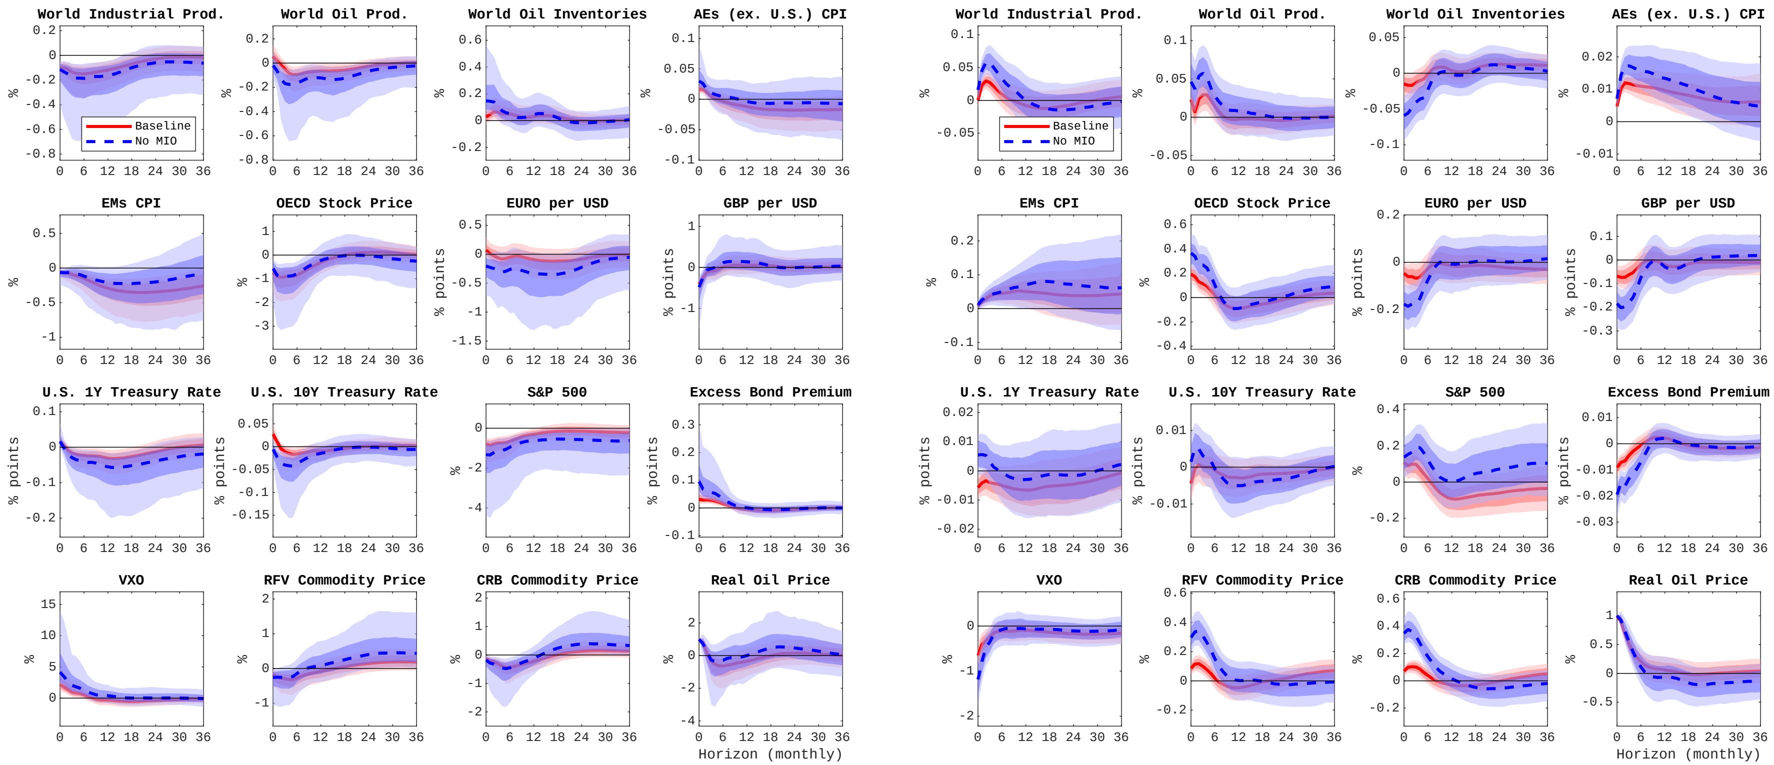

(a) Shock to oil supply expectations

(b) Shock to oil demand expectations

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 a shock to oil supply expectations. Right panel: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies. BVAR(12). Shaded areas represent 90% posterior coverage bands. Sample: 1984:1–2020:8.

Table C.1: Most Influential Observations

Oil demand expectations		Oil supply expectations	
Date	DFBETA	Date	DFBETA
2014:11	-0.451	2008:2	-0.315
1994:3	-0.307	2020:6	-0.233
2008:3	-0.220	2015:6	-0.178
1988:6	-0.212	2019:7	-0.170
2004:9	-0.197	2006:3	-0.164
1986:8	0.707	2016:11	0.595
2020:3	0.375	2017:5	0.426
1998:11	0.249	1996:12	0.330
2008:10	0.192	1986:12	0.195
1997:6	0.122	1987:12	0.163

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 oil demand shocks; right: for oil supply shocks) and a constant. The proxies are based on the S&P 500. The innovations are from the system represented in Figure C.4, estimated on the sample 1984:1–2020:8. DFBETA measures the difference in the parameter estimate with and without the influential observation.

C.5 Alternative measures of industrial production

Figure C.5 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.5: Alternative measures of industrial production

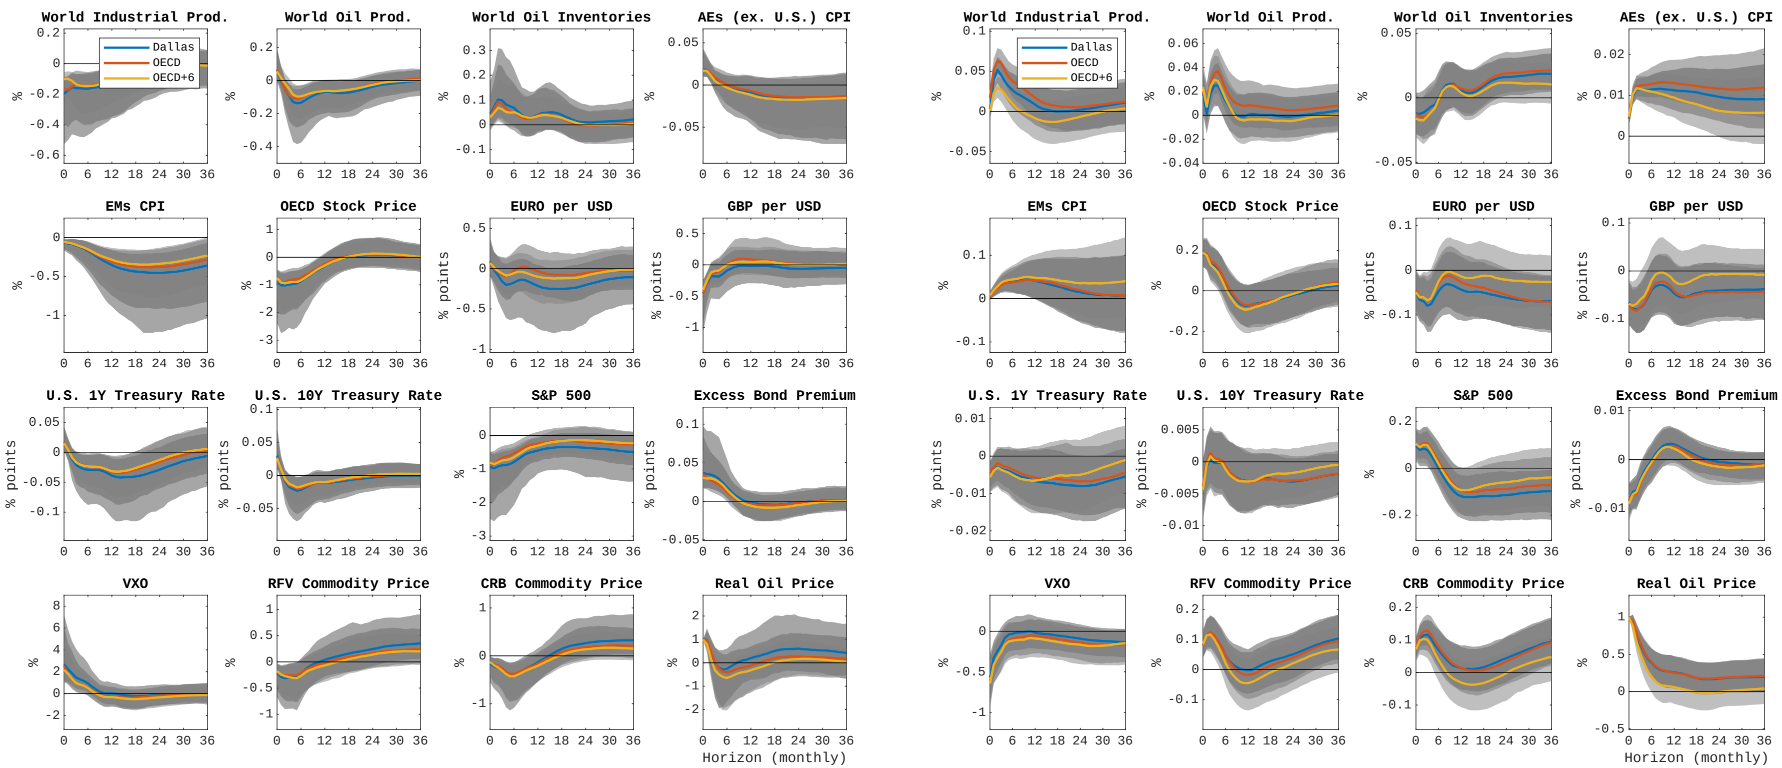

(a) Shock to oil supply expectations

(b) Shock to oil demand expectations

Note: Comparison among three alternative measures of World Industrial Production. Left panel: responses to a shock to oil supply expectations. Right panel: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% 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: 1984:1–2020:8.

Figure C.6: Transmission to the Global Economy excluding Covid-19

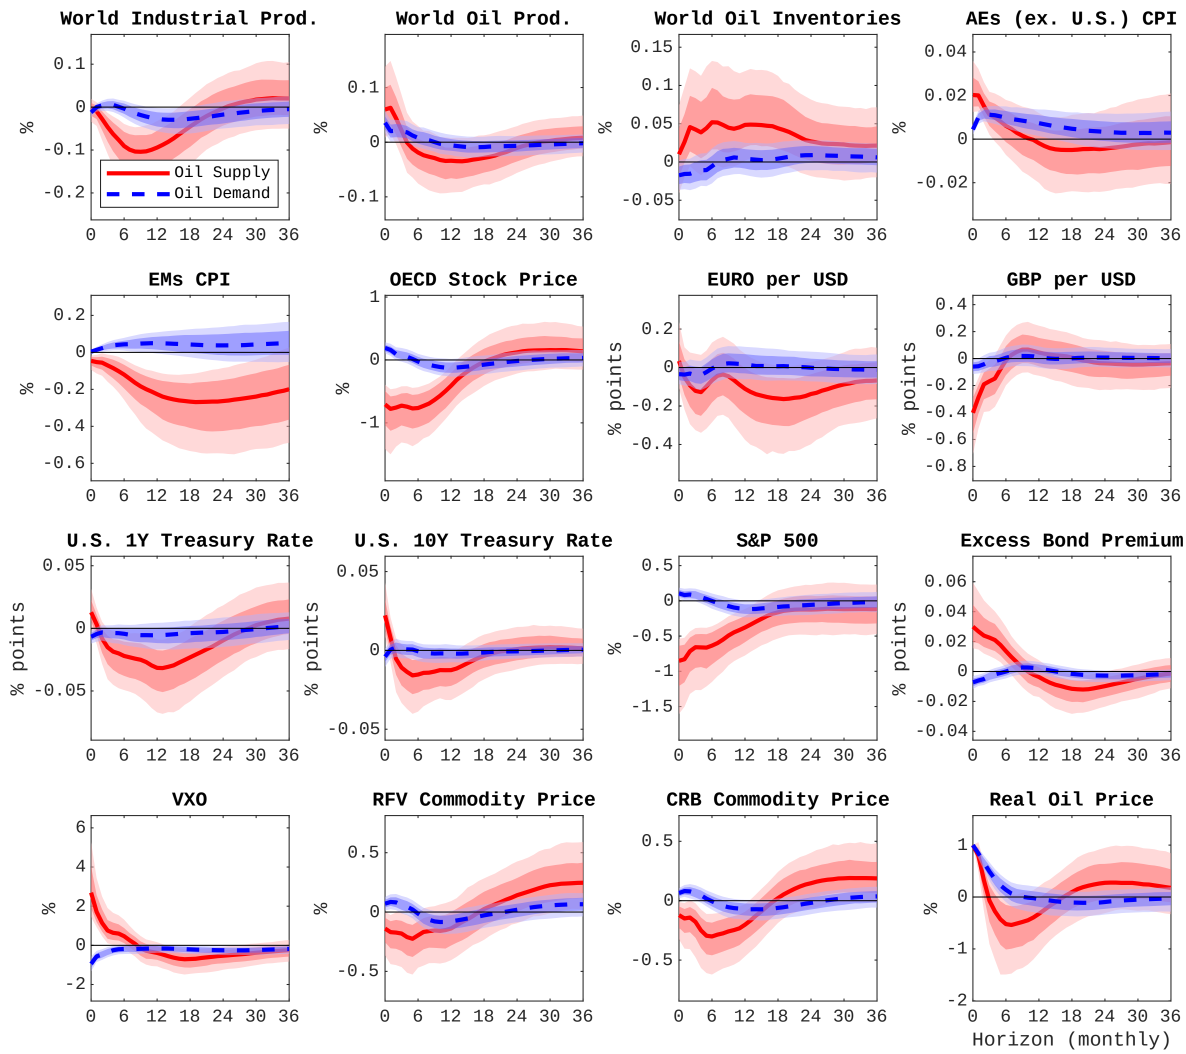

Note: Solid red: impulse responses to a shock to oil supply expectations. Dashed blue: responses to a shock to oil demand expectations. Both shocks are normalised to induce a 1% increase in real oil price. The shocks are identified using the robust proxies for shocks to oil supply and demand expectations. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: 1984:1–2019:12.

C.6 Excluding the Covid-19 period

Results are robust to cutting the sample at December 2019, before the Covid-19 pandemic. This is arguably an important check as Covid-19 produced deep discontinuities in the time series used in the analysis. Figure C.6 shows that the impulse responses coincide with the baseline results, although the expansion in world industrial production following a shock to oil demand expectations is now somewhat limited. Consistently with this result, 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.7: Country-specific responses of Industrial Production

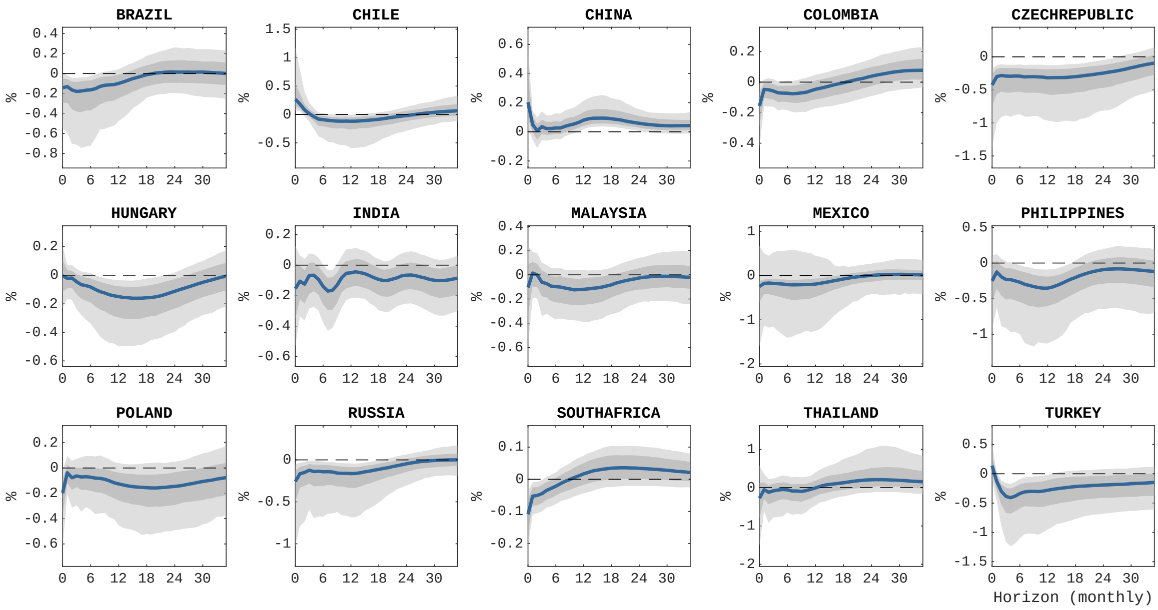

Note: Responses to a shock to oil supply expectations normalised to induce a 1% increase in real oil price. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1.

C.7 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.7 shows the response of industrial production for all 15 countries to a shock to oil supply expectations normalised to induce a 1% increase in real oil price. Figure C.8 displays the responses of the bilateral domestic/U.S. dollar exchange rate to the same shock. Finally, Figure C.9 shows the response of the stock price index to the shock.

Figure C.8: Country-specific responses of the Exchange Rate

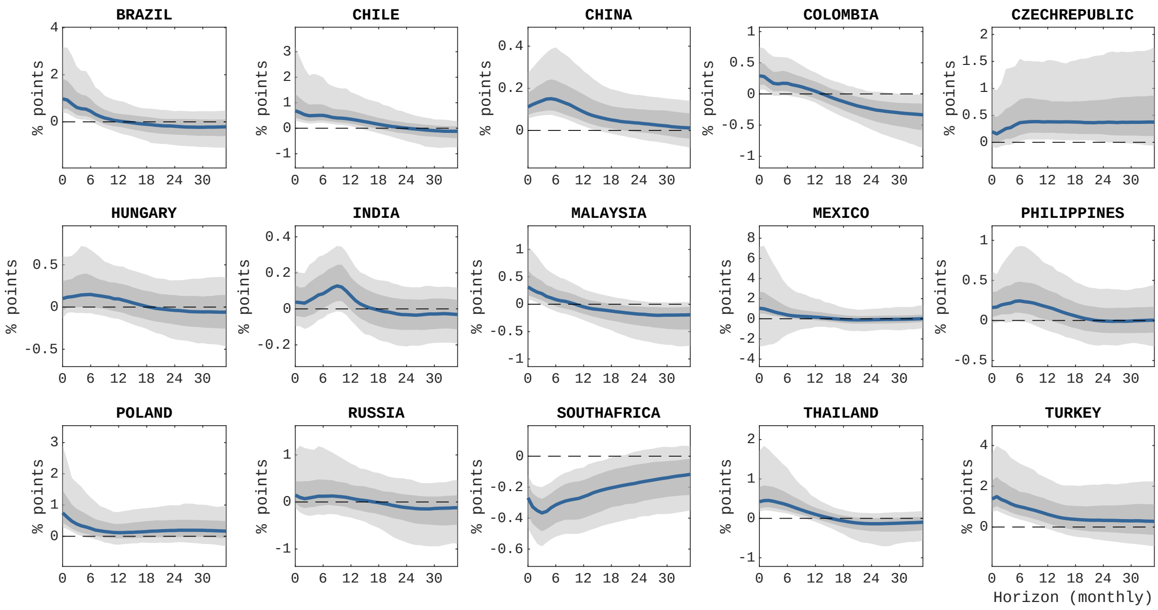

Note: Responses to a shock to oil supply expectations normalised to induce a 1% increase in real oil price. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1.

Figure C.9: Country-specific responses of Stock Prices

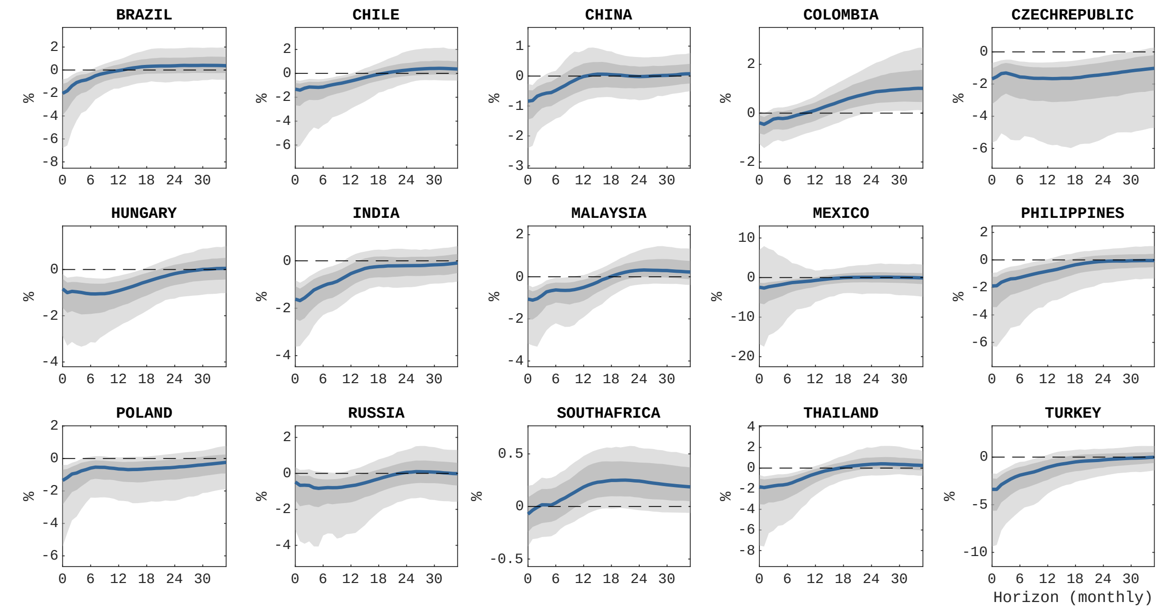

Note: Responses to a shock to oil supply expectations normalised to induce a 1% increase in real oil price. BVAR(12). Shaded areas represent 68% and 90% posterior coverage bands. Sample: see Table B.1.