Dauphine | PSL UNIVERSITÉ PARIS
LEDa UMR 8007-260

LEDa WORKING PAPERS

Reading the Future of Oil: A Noncausal Approach to Supply News Shocks

Stéphane Auray, Zakaria Moussa, Arthur Thomas

WP 2026-01
Janvier 2026

Reading the Future of Oil: A Noncausal Approach to Supply News Shocks*

Stéphane Auray† Zakaria Moussa‡ Arthur Thomas§

January 13, 2026

Abstract

This paper proposes a new strategy to identify oil supply news shocks by combining a Bayesian noncausal structural VAR with a Max-Share approach. The framework jointly resolves the problems of non-fundamentalness and recoverability that undermine standard (proxy) SVAR methods. Exploiting non-Gaussianity in a multivariate Student$-t$ specification, we recover structural shocks from a two sided moving average representation and isolate the expectation driven component of oil supply innovations without external instruments. Applied to global oil market data, the model supports a non fundamental representation and detects anticipatory price and inventory movements consistent with rational expectations storage behavior. The identified shocks explain a substantial fraction of real oil price fluctuations, notably in the late 1970s–1980s and during the 2014–15 collapse, while the COVID-19 episode is predominantly demand driven. Decomposing global supply shows that these shocks are primarily OPEC-driven and generate stagflationary responses in output and inflation, underscoring the central role of expectations in oil market dynamics.

JEL classification: C32, D84, E32, Q41, Q43

Keywords: Oil supply news shocks, Global oil market, Max-Share methodology, Non-fundamentalness, Structural Non-causal VAR

*We thank Christian Gouriéroux, Lutz Kilian, Jaakko Nelimarkka and Ricardo Degasperi for many useful suggestions and detailed comments, and many seminar and conference participants for their insights at the 7th Rimini Center for Economic Analysis (RCEA) Workshop, the 37th International Conference of the French Finance Association (AFFI) and the IAAE annual conference 2023.

†CREST-Ensai and Rennes School of Business, stephane.auray@ensai.fr

‡Nantes University, LEMNA, France, zakaria.moussa@univ-nantes.fr

§Université Paris-Dauphine, Université PSL, LEDa, CNRS, IRD, 75016 PARIS, FRANCE, arthur.thomas@dauphine.psl.eu

1 Introduction

The global oil market has faced recurrent episodes of pronounced volatility, driven by both demand and supply shocks that produce large price fluctuations. According to projections from the International Energy Agency (IEA), global supply is expected to broadly track demand through 2050, though with significant uncertainty reflecting technological, economic, and geopolitical constraints on resource development.[1] Intensifying geopolitical fragmentation over energy and climate policy has further heightened these uncertainties, underscoring the need to understand how expectations about future oil supply shape macroeconomic and financial outcomes.

A large body of macroeconomic research emphasizes the role of anticipated changes in future productivity, so called news shocks, as key drivers of business cycles (Beaudry and Portier, 2006; Barsky and Sims, 2011; Forni et al., 2014; Beaudry and Portier, 2014; Kurmann and Sims, 2021, among others). Building on this foundation, we define oil supply news shocks as exogenous shifts in the information set used by economic agents to form expectations about future global oil supply. Despite their conceptual relevance, few empirical studies have successfully identified or quantified these shocks.[2] Arezki et al. (2017), for instance, use the exogeneity of discovery timing in oil exploration as a quasi natural experiment to highlight the role of anticipated supply changes in shaping macroeconomic outcomes.

More recently, Känzig (2021) proposed a high frequency identification strategy combining OPEC announcement windows with the proxy SVAR framework to isolate oil supply news shocks. This approach provides compelling evidence that revisions in oil supply expectations matter for oil market and macroeconomic dynamics. Yet, it also raises the issue of whether high-frequency surprises can be treated as clean measures of

[1] For more details about different scenarios, see "Oil 2024 - Analysis and forecast to 2030".

[2] This news shock terminology is to be distinguished from the exogenous shock identified using the news component of either announcement of macroeconomic data releases (Andersen et al., 2003; Ramey, 2011; Kilian and Hicks, 2013, among others) or that used in the monetary policy literature exploiting central bank announcements to measure monetary policy shocks (Kuttner, 2001; Romer and Romer, 2004; Gürkaynak et al., 2005; Bernanke and Kuttner, 2005; Paul, 2020, among others), applied recently by Känzig (2021) on the OPEC announcements to identify an oil supply news shock.

low-frequency structural shocks, given potential aggregation and measurement problems highlighted by Kilian (2024). Moreover, Kilian and Zhou (2023) and Degasperi et al. (2025) argue that these OPEC-related surprises are better interpreted as shifts in broader oil price expectations rather than pure supply news. Of particular importance in the context of our paper, Plagborg-Møller and Wolf (2022) clearly demonstrated the non-recoverability of Kanzig's proxy-SVAR in the oil market, highlighting the need to address this issue. Recent advances by Chahrour and Jurado (2021), Plagborg-Møller and Wolf (2022), and Forni et al. (2025) propose alternative frameworks that test or bound recoverability in non invertible systems and embed forward looking dynamics explicitly, but the oil market literature has yet to implement a coherent solution that handles both issues simultaneously.

Our first contribution, therefore, is to provide such a unified solution. Using a stylized model of the oil market, we demonstrate that a noncausal structural VAR (NC-VAR) naturally solves the recoverability problem. The noncausal specification explicitly allows variables to depend on both past and future shocks. These models, introduced by Lanne and Saikkonen (2011) and extended to the multivariate setting by (Lanne and Saikkonen, 2013; Lanne and Luoto, 2016; Davis and Song, 2020; Velasco, 2023; Gourieroux and Jasiak, 2023), capture forward-looking behavior consistent with rational expectations and potentially infinite-variance disturbances (Gouriéroux et al., 2020), while maintaining the stationarity condition. These models have been successfully applied to fiscal foresight (Nelimarkka, 2017b), technology news (Nelimarkka, 2017a), and the New Keynesian Phillips curve (Lanne and Luoto, 2013).[3] In a closely related approach, Chahrour and Jurado (2021) provide a recovery-based framework that establishes conditions under which structural shocks can be uniquely identified from two-sided moving average

[3]Two alternative representations exist for multivariate mixed causal–noncausal processes. The VMAR(r,s) model of Lanne and Saikkonen (2013) adopts a multiplicative form $\Pi(L^{-1})\Phi(L)Y_t = \epsilon_t$, requiring the number of roots in each component to be a multiple of the number of variables. The VAR(n_1,n_2,p) representation of Davis and Song (2020); Gourieroux and Jasiak (2023) instead uses $Y_t = \sum_{j=1}^{p} \Theta_j Y_{t-j} + u_t$, where n_1 and n_2 indicate the roots outside and inside the unit circle, respectively, without multiplicity constraints. We adopt the VMAR(r,s) form because it separates causal and noncausal polynomial matrices, allowing economically meaningful structural restrictions on forward- and backward-looking components. A minor contribution of this paper is to show that the VMAR representation is necessary for identifying news shocks, as the structural interpretation requires explicit separation of forward-looking dynamics that is not available in the VAR(n_1,n_2,p) form.

(MA) representations in non-fundamental environments. However, their recoverability conditions impose the two-sided lead–lag structure as an assumption rather than testing whether such a structure is empirically identified in the data. While their framework guarantees unique identification given the assumed lead–lag structure, it does not recover the causal–noncausal decomposition itself, which requires exploiting higher-order moments that are absent under Gaussianity.[4] Following Gourieroux and Jasiak (2025), we exploit non-Gaussianity to break the observational equivalence between fundamental and non fundamental representations, enabling the identification of true impulse responses to anticipated shocks. However, we still need an identification strategy to isolate a news shock from the reduced-form error term of the NC-VAR.

Our second contribution is to provide such identification strategy: we adopt the Max-Share identification strategy *à la* Chahrour et al. (2023) in a noncausal VAR framework estimated for standard oil market variables-global oil production, global economic activity, real oil price, and global oil stocks. In this setting, the oil news shock is defined as the innovation that maximizes the finite-horizon forecast error variance share of global oil production. This adaptation requires no external proxies and remains robust to non-fundamentalness, since the structural shocks are, by construction, non-invertible. A minor contribution is to provide a framework for computing historical decompositions in non-causal VAR systems, which extends standard decomposition methods to environments with forward-looking dynamics.

Finally, exploiting institutional and behavioral differences between OPEC and non OPEC producers allows us to trace the origins of the identified news shocks. OPEC's coordinated production management contrasts with the market driven behavior of non OPEC producers, yielding distinct supply elasticities and adjustment speeds (Pierru et al., 2018, 2020; Baumeister and Hamilton, 2024). Disentangling these components reveals that OPEC-driven shocks dominate aggregate oil supply news and exhibit more pronounced macroeconomic effects.

[4]Higher moments or non-Gaussian likelihoods are required for that step, as causal and noncausal processes are second-moment equivalent under Gaussianity (Rosenblatt, 2000; Davis and Song, 2020; Gourieroux and Jasiak, 2017; Gouriéroux et al., 2020; Velasco, 2023; Gourieroux and Jasiak, 2023).

Empirically, the data provide strong support for a noncausal, fat tailed specification, confirming the relevance of a non-fundamental representation of the global oil market. The identified oil supply news shocks display clear anticipatory features: real oil prices and inventories increase in the leads, while production and activity fall with delay, consistent with rational expectations storage behavior. Historical decompositions show that these shocks explain a significant and time varying share of real oil price movements, especially in the late 1970s–1980s and around 2014, while the COVID-19 collapse appears demand-driven. At the global and U.S. level, negative oil supply news shocks generate stagflationary dynamics, raising inflation and uncertainty while weighing on output and employment. Yet, their effects on equity and financial markets remain limited, suggesting constrained but not destabilizing spillovers.

The remainder of the paper is organized as follows. Section 2 presents a stylized model of the oil market that illustrates how non-fundamentalness arises from a news shocks and demonstrates how the NC-VAR framework resolves the recoverability problem. Section 3 details the NC-VAR specification, impulse response function computation, historical decomposition, and the Bayesian estimation procedure. Section 4 describes the data, presents our Max-Share identification strategy. Section 5 reports the baseline results for the global oil market and extended macroeconomic variables at both global and U.S. levels. Section 6 disaggregates the analysis to trace the origins of oil supply news shocks, contrasting OPEC-driven versus non-OPEC production dynamics. Section 7 concludes.

2 Stylised model of the oil market with news shocks

In this section, we consider a stylized rational expectations model of the global oil market that explains oil price movements driven by expectations about future oil supply. We show how non-fundamentalness arises from the presence of lagged effects of oil supply shocks on oil prices, and how the use of a noncausal VAR approach implies the recoverability of news shocks in the sense of Chahrour and Jurado (2021). Consider a process for global oil production q_t that evolves as an AR(1) with both contemporaneous and anticipated

effects[5]:

$$q_t = \rho q_{t-1} + \chi \epsilon_t + \epsilon_{t-l} \tag{1}$$

where $|\rho| < 1$, ϵ_t is the oil supply news shock following a strong white noise process, and $l \geq 1$ denotes the anticipation horizon.[6] The parameter $\chi \geq 0$ captures the contemporaneous impact of the shock on oil production. When $\chi < 1$, the anticipated lag term ϵ_{t-l} dominates the dynamics of q_t, which is the key condition for non-fundamentalness to arise.

Suppose then that the oil price p_t (or any forward-looking variable, including inventories, except for capacity constraints) is determined by the following equilibrium condition:

$$p_t = \beta \mathbb{E}_t \left(p_{t+1} \right) + q_t + \nu_t \tag{2}$$

where $\beta < 1$ is a discount factor, $\mathbb{E}_t[\cdot]$ denotes the conditional expectation with respect to the information set containing the history of $\{q_t, p_t, \epsilon_t, \nu_t\}$, and ν_t is an exogenous disturbance (e.g., a demand or speculative shock) that affects the oil price contemporaneously but is orthogonal to the oil supply news shock: $\mathbb{E}_t[\nu_{t+h} \mid \epsilon_{t+h'}] = 0$ for all h, h'. The inclusion of ν_t allows the price to respond immediately to factors other than the news shock, which is empirically relevant as oil prices react to a variety of contemporaneous disturbances.

The forward-looking solution of equation (2), assuming $|\beta| < 1$ for convergence, is given by the present discounted value of expected future oil supply plus the exogenous price disturbance:

$$p_t = \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t \left(q_{t+j} \right) + \nu_t \tag{3}$$

To derive the closed-form expression for p_t, we compute $\mathbb{E}_t(q_{t+j})$ for $j \geq 0$. From

[5] As in our identification strategy, we allow the news shock to have an immediate impact on oil production; see Section 4.2

[6] $l \geq 1$ is necessary for non-fundamentalness to be treated within the NC-VAR framework. If l is equal to zero, solution exists using in the noncausal VARMA framework where noncausality is located in the MA part and not in the AR part (Gouriéroux et al., 2020).

equation (1), iterating forward yields:

$$\mathbb{E}_t(q_{t+j}) = \rho^j q_t + \sum_{i=0}^{j-1} \rho^{j-1-i} \mathbb{E}_t(\chi \epsilon_{t+1+i} + \epsilon_{t+1+i-l}) \tag{4}$$

Using $\mathbb{E}_t(\epsilon_{t+k}) = 0$ for $k > 0$ and $\mathbb{E}_t(\epsilon_{t+k}) = \epsilon_{t+k}$ for $k \leq 0$. As an illustration we present results for $l = 2$, the general case is treated in Appendix B:

$$\mathbb{E}_t(q_{t+1}) = \rho q_t + \epsilon_{t-1} \tag{5}$$

$$\mathbb{E}_t(q_{t+2}) = \rho^2 q_t + \rho \epsilon_{t-1} + \epsilon_t \tag{6}$$

$$\mathbb{E}_t(q_{t+j}) = \rho^j q_t + \rho^{j-1} \epsilon_{t-1} + \rho^{j-2} \epsilon_t, \quad j \geq 2 \tag{7}$$

Substituting these expressions into equation (3) and collecting terms:

$$\begin{aligned} p_t &= \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t(q_{t+j}) + \nu_t \\ &= q_t + \beta(\rho q_t + \epsilon_{t-1}) + \sum_{j=2}^{\infty} \beta^j \left(\rho^j q_t + \rho^{j-1} \epsilon_{t-1} + \rho^{j-2} \epsilon_t \right) + \nu_t \\ &= q_t \sum_{j=0}^{\infty} (\beta\rho)^j + \epsilon_{t-1} \sum_{j=1}^{\infty} \beta^j \rho^{j-1} + \epsilon_t \sum_{j=2}^{\infty} \beta^j \rho^{j-2} + \nu_t \end{aligned} \tag{8}$$

Evaluating the geometric series:

$$\sum_{j=0}^{\infty} (\beta\rho)^j = \frac{1}{1 - \rho\beta} \tag{9}$$

$$\sum_{j=1}^{\infty} \beta^j \rho^{j-1} = \frac{\beta}{1 - \rho\beta} \tag{10}$$

$$\sum_{j=2}^{\infty} \beta^j \rho^{j-2} = \frac{\beta^2}{1 - \rho\beta} \tag{11}$$

Thus, the equilibrium price is:

$$p_t = \frac{1}{1 - \rho\beta} q_t + \frac{\beta}{1 - \rho\beta} \epsilon_{t-1} + \frac{\beta^2}{1 - \rho\beta} \epsilon_t + \nu_t \tag{12}$$

Substituting the expression for q_t from equation (1):

$$p_t = \frac{\rho}{1-\rho\beta}q_{t-1} + \frac{1}{1-\rho\beta}(\chi+\beta^2)\epsilon_t + \frac{\beta}{1-\rho\beta}\epsilon_{t-1} + \frac{1}{1-\rho\beta}\epsilon_{t-2} + \nu_t \tag{13}$$

The structural moving average representation of $(q_t, p_t)'$ is:

$$\begin{bmatrix} q_t \\ p_t \end{bmatrix} = \begin{bmatrix} \rho & 0 \\ \frac{\rho}{1-\rho\beta} & 0 \end{bmatrix} \begin{bmatrix} q_{t-1} \\ p_{t-1} \end{bmatrix} + \underbrace{\begin{bmatrix} \chi + L^2 & 0 \\ \frac{\chi+\beta^2}{1-\rho\beta} + \frac{\beta}{1-\rho\beta}L + \frac{1}{1-\rho\beta}L^2 & 1 \end{bmatrix}}_{=B(L)} \begin{bmatrix} \epsilon_t \\ \nu_t \end{bmatrix} \tag{14}$$

The process $(q_t, p_t)'$ is fundamental if and only if $|B(z)| \neq 0$ for all $|z| \leq 1$. Computing the determinant:

$$|B(z)| = \chi + z^2 \tag{15}$$

This polynomial has roots at $z = \pm\sqrt{-\chi} = \pm i\sqrt{\chi}$. When $\chi < 1$, the modulus of these roots is $\sqrt{\chi} < 1$, implying that there exist roots inside the unit circle. Therefore, the process has a non-fundamental MA representation. In other words, once the anticipated component of the news shock dominates ($\chi < 1$), the observables suffer from non-fundamentalness and thus no causal VAR representation of $(q_t, p_t)'$ exists for the structural shocks (Gambetti and Moretti, 2017; Kilian and Lütkepohl, 2017; Gouriéroux et al., 2020; Nelimarkka, 2017a, among others). To address the non-fundamentalness issue, we rewrite the system in noncausal form. From equation (1), we have:

$$(1-\rho L)q_t = (\chi + L^2)\epsilon_t = (L^2 + \chi)\epsilon_t \tag{16}$$

Rewriting the right-hand side:

$$(L^2 + \chi)\epsilon_t = (1 + \chi L^{-2})\epsilon_{t-2} \tag{17}$$

where the polynomial $1 + \chi z^2$ has roots outside the unit circle when $\chi < 1$.[7] Thus:

$$(1 - \rho L)(1 + \chi L^{-2})^{-1} q_t = \epsilon_{t-2} \tag{18}$$

This shows that q_t has a noncausal representation with leads.

As shown by Nelimarkka (2017a) and in Appendix B, when a news shock is identified within the representation (14), the underlying noninvertibility of the MA part can be addressed by a noncausal VAR representation. More specifically, and thus demonstrates that a NC-VAR$(1, s)$ representation, with s appropriately chosen to capture the anticipation horizon, produces impulse response functions that perfectly match the theoretical IRFs from the structural model (1)–(2) for any value of $\chi \in [0, 1)$. To illustrate the noncausal representation explicitly, consider the special case $\chi = 0$, where the news shock has no contemporaneous effect on oil production and affects it only with two lags:

$$q_t = \rho q_{t-1} + \epsilon_{t-2} \tag{19}$$

In this case, equation (13) simplifies to:

$$p_t = \frac{\rho}{1 - \rho\beta} q_{t-1} + \frac{\beta^2}{1 - \rho\beta} \epsilon_t + \frac{\beta}{1 - \rho\beta} \epsilon_{t-1} + \frac{1}{1 - \rho\beta} \epsilon_{t-2} + \nu_t \tag{20}$$

To address the non-fundamentalness issue, we rewrite the system in noncausal form. From equation (19):

$$(1 - \rho L) q_t = \epsilon_{t-2} \tag{21}$$

Using the representation to substitute the shocks in equation (20), p_t can be written in noncausal form:

$$p_t = q_t + \beta q_{t+1} + \frac{\beta^2}{1 - \rho\beta} q_{t+2} + \nu_t \tag{22}$$

The dynamics of $(q_t, p_t)'$ can then be mapped into a noncausal VAR(1,2) using equa-

[7] Note that the roots are automatically complex in this case, generating cyclical dynamics.

tions (19) and (22):

$$\underbrace{\left(I_2 - \begin{bmatrix}\rho & 0\\ \rho & 0\end{bmatrix} L\right)}_{\Pi(L)} \underbrace{\left(I_2 - \begin{bmatrix}0 & 0\\ \beta & 0\end{bmatrix} L^{-1} - \begin{bmatrix}0 & 0\\ \frac{\beta^2}{1-\rho\beta} & 0\end{bmatrix} L^{-2}\right)}_{\Phi(L^{-1})} \begin{bmatrix}q_t\\ p_t\end{bmatrix} = \begin{bmatrix}1 & 0\\ 1 & 1\end{bmatrix} \begin{bmatrix}\epsilon_{t-2}\\ \nu_t\end{bmatrix} \tag{23}$$

where $\Pi(L)$ and $\Phi(L^{-1})$ correspond to the causal (with one lag) and the noncausal (with two leads) polynomials, respectively. Importantly, as these latter matrices are invertible, p_t has a two-sided MA representation. The existence of a two-sided MA representation has a direct implication for the recoverability of the oil supply news shock ϵ_t, in the sense of Chahrour and Jurado (2021) and Plagborg-Møller and Wolf (2022).

Proposition 1. *Consider the noncausal VAR(1,2) representation of $(q_t, p_t)'$ given by equation (23), Since $|\rho| < 1$ and $|\beta| < 1$, both $\Pi(L)$ and $\Phi(L^{-1})$ have all roots outside the unit circle. Consequently, $(q_t, p_t)'$ admits a two-sided moving average representation:*

$$\begin{bmatrix}q_t\\ p_t\end{bmatrix} = \sum_{j=-\infty}^{\infty} \Psi_j \begin{bmatrix}\epsilon_{t-2-j}\\ \nu_{t-j}\end{bmatrix}$$

where $\Psi(z) = \Phi(z^{-1})^{-1}\Pi(z)^{-1}$. The oil supply news shock ϵ_t is then recoverable, i.e., it satisfies:

$$\epsilon_t = \mathbb{E}\left(\epsilon_t \mid \{q_s, p_s\}_{s\in\mathbb{Z}}\right)$$

In other words, ϵ_t is spanned by all leads and lags of the observable variables $(q_t, p_t)'$.

Proof. The proof follows from the invertibility of the two-sided MA representation. Since $\det \Pi(z) = 1 - \rho z \neq 0$ for $|z| \leq 1$ (as $|\rho| < 1$) and $\det \Phi(z) = 1 \neq 0$ for all z, the polynomial $\Psi(z) = \Phi(z^{-1})^{-1}\Pi(z)^{-1}$ is well-defined and has a two-sided inverse $\Psi(z)^{-1} = \Pi(z)\Phi(z^{-1})$. Consequently, the structural shocks $(\epsilon_{t-2}, \nu_t)'$ can be expressed as a two-sided linear combination of $(q_s, p_s)'$ for $s \in \mathbb{Z}$. By Chahrour and Jurado (2021), Proposition 1, this implies that the structural shocks are recoverable. Formally, defining $y_t = (q_t, p_t)'$ and $u_t = (\epsilon_{t-2}, \nu_t)'$, we have:

$$u_t = \sum_{j=-\infty}^{\infty} \Psi_j^{-1} y_{t-j}$$

where $\Psi^{-1}(z) = \sum_{j=-\infty}^{\infty} \Psi_j^{-1} z^j$ is the two-sided inverse of $\Psi(z)$. This establishes that the shock ϵ_t (appearing as ϵ_{t-2} in u_t) is a linear combination of all leads and lags of the observables, which is the definition of recoverability. □

In this illustrative example, this proposition establishes that the NC-VAR framework provides a natural solution to the recoverability problem identified by Plagborg-Møller and Wolf (2022) in the context of oil market SVARs with external instruments. While standard causal VAR representations fail to recover the news shock from past observables alone (due to non-fundamentalness), the two-sided MA representation of the NC-VAR ensures that the shock can be recovered from the entire history and future of the observable process. We also provide a solution that differs from Chahrour and Jurado (2021) as: while they establish necessary and sufficient recoverability conditions through a rank condition on the spectral characteristic matrix—thereby determining whether shocks are identifiable given a known causal-noncausal structure—our NC-VAR approach simultaneously identifies the lead-lag structure of the data-generating process through non-Gaussian likelihood and ensures recoverability by construction through the multiplicative VMAR(r, s) of (23) representation. Furthermore, we show that this multiplicative VMAR(r, s) representation of Lanne and Saikkonen (2013), which cannot be recast into the VAR(n_1, n_2, p) framework with i.i.d. errors proposed by Davis and Song (2020) and Gourieroux and Jasiak (2017). This incompatibility arises from the rank deficiency of the lead coefficient matrix Φ_2, a structural feature stemming from the triangular transmission mechanism of news shocks; a formal proof following Chapter 4 of Giancaterini (2023) is provided in Appendix A.

In the general case where both the anticipation horizon l and the contemporaneous impact parameter χ are unknown, the infinite two-sided MA representation must be truncated for practical estimation. This motivates the use of a flexible NC-VAR(r, s) specification, where the lag and lead orders (r, s) are determined by the data, thereby accommodating richer dynamics in which the dependence of oil production q_t may extend beyond a single lag; the formal derivation for arbitrary $l \geq 1$ and $\chi \in [0, 1)$ is provided in Appendix B.

3 Econometric framework

This section introduces the econometric framework underlying our empirical analysis. We first present the reduced-form noncausal VAR model and its two-sided moving average representation and detail the computation of impulse response functions and historical decompositions.

3.1 The noncausal structural VAR

Consider the Bayesian NC-VAR(r,s) model developed by Lanne and Luoto (2016), where y_t is generated by:

$$\Pi(L)\,\Phi\left(L^{-1}\right) y_t = \epsilon_t \tag{24}$$

where the causal polynomial $\Pi(L) = I_n - \Pi_1 L - \cdots - \Pi_r L^r$, the noncausal polynomial $\Phi(L^{-1}) = I_n - \Phi_1 L^{-1} - \cdots - \Phi_s L^{-s}$, L is the backward shift operator, and ϵ_t is a sequence of independent, identically distributed random vectors with zero mean and a finite positive definite covariance matrix Σ. The stationarity of the process and the existence of a two-sided MA representation are guaranteed by the assumption that both matrix polynomials have their roots outside the unit disc, i.e., $\det \Pi(z) \neq 0$ and $\det \Phi(z) \neq 0$ for $|z| \leq 1$. The process admits a two-sided MA representation:

$$y_t = \sum_{j=-\infty}^{\infty} \Psi_j \epsilon_{t-j}, \quad \Psi(z) = \Phi\left(z^{-1}\right)^{-1} \Pi(z)^{-1} \tag{25}$$

which can be rewritten to highlight the future-dependent component:

$$y_t = \Phi_1 y_{t+1} + \cdots + \Phi_s y_{t+s} + \sum_{j=0}^{\infty} M_j \epsilon_{t-j} \tag{26}$$

where the non-zero lead coefficients Φ_i indicate that y_t depends on future values, reflecting the non-fundamental nature of the process. A crucial requirement is that ϵ_t must be non-Gaussian. Under Gaussianity, the NC-VAR(r,s) becomes observationally equivalent to a causal VAR($r+s$), making it impossible to separately identify the causal and non-

causal dynamics Lanne and Saikkonen (2013). Non-Gaussianity—through higher-order moments—provides the additional information needed to distinguish leads from lags in the data-generating process and to identify structural shocks (Gouriéroux et al., 2020; Gourieroux and Jasiak, 2017, 2023, 2025). We assume a multivariate t-distribution for ϵ_t:

$$\epsilon_t = \omega_t^{-1/2} \eta_t, \tag{27}$$

where $\eta_t \sim \mathrm{N}(0, \Sigma)$, $\lambda\omega_t \sim \chi^2_\lambda$, and λ denotes the degrees of freedom. This specification generates fat tails for small λ and approaches Gaussianity as $\lambda \to \infty$. The scalar volatility factor $\omega_t^{-1/2}$ induces the excess kurtosis required for identification. The Bayesian estimation follows Lanne and Luoto (2016) using Gibbs sampling with Minnesota-type priors that shrink lead coefficients more heavily toward zero than lag coefficients. This helps to attain the unimodalilty when posterior distributions are more likely to be multimodal when estimating the NC-VAR model (Lanne and Luoto, 2016). Details on the prior specification and the sampling algorithm are provided in Appendix D.

3.2 Impulse response function computation

The computation of impulse response functions (IRFs) in the NC-VAR framework exploits the two-sided MA representation in equation (25). Given the structural decomposition of the reduced-form error term in equation (33), where $\epsilon_t = \bar{B}\bar{u}_t$ with $\bar{B} = \tilde{A}W$, the structural two-sided MA representation takes the form:

$$y_t = \sum_{h=-\infty}^{\infty} \Psi_h \bar{B} \bar{u}_{t-h} = \sum_{h=-\infty}^{\infty} \Theta_h \bar{u}_{t-h} \tag{28}$$

where $\Theta_h = \Psi_h \bar{B}$ denotes the structural impulse response matrix at horizon h. The response of variable i to a unit shock in structural innovation $\bar{u}_{j,t}$ at horizon h is then given by:

$$\frac{\partial y_{i,t+h}}{\partial \bar{u}_{j,t}} = e_i' \Psi_h \bar{B} e_j = e_i' \Psi_h \tilde{A} w_j = \Theta_{i,j,h} \tag{29}$$

where e_i is a selection vector with one in position i and zeros elsewhere, w_j is the jth column of the orthogonal rotation matrix W, and $\gamma_j = \tilde{A}w_j$ represents the jth column of the structural impact matrix $\bar{B}$. A key distinction between the NC-VAR from standard causal VARs is that the horizon index h ranges over all integers, $h \in \mathbb{Z}$, rather than being restricted to non-negative values. The coefficients Ψ_h for $h < 0$ capture the *anticipation effects*—the responses of variables to shocks that have not yet materialised but are already anticipated by forward-looking agents. In practice, the MA coefficient matrices Ψ_h are computed by inverting the causal and noncausal polynomials. From equation (25), we have $\Psi(z) = \Phi(z^{-1})^{-1}\Pi(z)^{-1}$ for all z on the unit circle. The coefficients $\{\Psi_h\}_{h=-\infty}^{\infty}$ are obtained via a Laurent series expansion, computed numerically by truncating the infinite sums at sufficiently large horizons. For the causal part ($h \geq 0$), this involves the standard recursive computation from the inverted lag polynomial $\Pi(z)^{-1}$. For the noncausal part ($h < 0$), the lead polynomial $\Phi(z^{-1})^{-1}$ generates the anticipation coefficients. The cumulated IRFs reported in our empirical analysis are obtained by summing the point-wise responses: $\text{IRF}_{i,j}^{cum}(H) = \sum_{h=-s}^{H} \Theta_{i,j,h}$, where the lower bound $-s$ corresponds to the lead order of the NC-VAR. This cumulation is particularly relevant for variables expressed in growth rates (such as global oil production), as it recovers the response in levels. The credible sets displayed in the figures are constructed from the posterior distribution of the IRFs, obtained by computing the structural IRF for each draw from the Gibbs sampler and extracting the relevant quantiles across draws.

3.3 Historical decomposition

This section presents the construction of historical decomposition in the noncausal VAR framework, building directly on the structural two-sided MA representation established in Section 3.2. From equation (28), the contribution of structural shock j to the observed value of variable i at time t is defined as

$$HD_{i,t}^{(j)} = \sum_{h=-\infty}^{\infty} \Theta_{i,j,h} \cdot \bar{u}_{j,t-h} \tag{30}$$

where $\Theta_{i,j,h}$ denotes the (i,j)-th element of the structural impulse response matrix $\Theta_h = \Psi_h \bar{B}$ at horizon h, and $\bar{u}_{j,t}$ is the j-th structural shock. By construction, the sum of contributions from all structural shocks recovers the observed variable: $y_{i,t} = \sum_{j=1}^{k} HD_{i,t}^{(j)}$. This decomposition allows us to assess how much of the observed variation in each variable at each point in time can be attributed to cumulated oil supply news shocks versus other structural disturbances.

In practice, the infinite summation in equation (30) must be truncated to a finite horizon. Due to the stationarity conditions implied by equation (24), the structural MA coefficients Θ_h decay exponentially as $|h|$ increases. This property ensures that truncation at a sufficiently large horizon H_{irf} introduces negligible approximation error. The practical implementation thus computes

$$\hat{HD}_{i,t}^{(j)} = \sum_{h=-H_{irf}}^{H_{irf}} \Theta_{i,j,h} \cdot \bar{u}_{j,t-h} \tag{31}$$

where the structural shocks are recovered from the estimated reduced-form residuals via the relation $\bar{u}_t = \bar{B}^{-1}\varepsilon_t$.

A distinctive feature of noncausal VAR models is the possibility of separating the historical decomposition into anticipation and realization components:

$$HD_{i,t}^{(j)} = \underbrace{\sum_{h=1}^{H_{irf}} \Theta_{i,j,-h} \cdot \bar{u}_{j,t+h}}_{\text{Anticipation effects}} + \underbrace{\sum_{h=0}^{H_{irf}} \Theta_{i,j,h} \cdot \bar{u}_{j,t-h}}_{\text{Realization effects}} \tag{32}$$

The first term captures how future shocks, which are already anticipated by forward-looking agents, influence current observations. The second term represents the conventional backward-looking contribution from current and past shocks. The Bayesian estimation framework provides a natural approach for quantifying uncertainty in the historical decomposition. For each posterior draw $m = 1, \ldots, M$ from the Gibbs sampler, we extract the reduced-form residuals $\varepsilon_t^{(m)}$, the structural impact matrix $\bar{B}^{(m)}$, and the impulse response functions $\Theta_h^{(m)}$. We then compute the structural shocks as $\bar{u}_t^{(m)} = (\bar{B}^{(m)})^{-1}\varepsilon_t^{(m)}$ and evaluate the historical decomposition according to equation (31). The posterior dis-

tribution of $HD_{i,t}^{(j)}$ is characterized by computing pointwise quantiles across the M draws.

4 Empirical strategy

This section describes our empirical implementation of the noncausal VAR framework developed in Section 3 to identify oil supply news shocks in global oil markets. We first present the dataset and the baseline specification of our four-variable NC-VAR model, including the choice of lag and lead orders and the Bayesian estimation procedure. We then detail our identification strategy, which adapts the Max-Share methodology of Chahrour et al. (2023) to the noncausal VAR context to recover the rotation matrix corresponding to an oil supply news shock.

4.1 Data and estimation set up

Estimation data span the period between January 1974 and July 2025.[8] We estimate a monthly NC-VAR model using the above standard global oil market endogenous variables (our baseline model) :

$$y_t = [\Delta wop_t, \Delta wip_t, RRAC_t, \Delta stocks_t],$$

where Δwop_t denotes the percentage change in world oil production, obtained from the US Energy Information Administration's Monthly Energy Review, wip_t is the growth rate of the world industrial production proxied by the OECD+6 industrial production proposed by Baumeister and Hamilton (2019), [9] $RRAC_t$ denotes the real oil price measured by the refiner acquisition cost (RAC) for imported crude oil and deflated by the US consumer price index and $\Delta stocks_t$ is the proxy for the percentage change in global oil stocks as

[8] More details about the data and their sources can be found in Appendix C

[9] Our results are robust when we use the World Industrial Production, the Global Steel Production Factor, and the Real Commodity Price Factor as the measures of global economic activity discussed in Baumeister and Guérin (2021). The response of the Kilian index, however, is qualitatively different from that of the other variables. For brevity, we do not report these results, but they are available upon request.

constructed in Kilian and Murphy (2014).[10] It should be noted that this variable is very useful for our analysis not only because it solves the informational deficit of the VAR system as already explained in Kilian and Murphy (2014), but also because it helps judge the validity of our oil news shock. When agents anticipate, for example, a future oil shortfall, they increase their oil stock level, whereas an unanticipated oil supply shortage has an opposite effect on the stock as agents who did not anticipate this drop in oil supply will have to draw on their reserves.

The NC-VAR(r,s) models are estimated by employing a lag and lead length of 12 to capture the full dynamic of observables. This choice enables observables to fully capture oil supply news innovations when the NC-VAR(r,s) reduces to a standard autoregressive VAR in the absence of any non-fundamentalness issues. Moreover, the noncausal part will have a rich structure if a non-fundamentalness problem arises.[11] Results are based on 10,000 posterior draws obtained after a burn-in period of 50,000 draws. The choice of priors for our four baseline models is reported in Appendix D.

4.2 Oil supply news shock identification

To implement our identification strategy, we build on the extensive econometric literature on the identification of news shocks based on the Max-Share methodology (Beaudry and Portier, 2006; Barsky and Sims, 2011; Forni et al., 2014; Beaudry and Portier, 2014; Kurmann and Sims, 2021, among others). More precisely, we adapt the identification procedure of Chahrour et al. (2023) to the noncausal VAR context, to help us to recover the rotation matrix in the set of admissible representations that correspond to a news shock.

This approach identifies the oil news shock as the shock that accounts for the largest fraction of anticipated fluctuations in the variable targeted for identification at relatively

[10] Our baseline regression uses a similar specification as in Baumeister and Hamilton (2019) and Kilian and Murphy (2014), except for the global oil stocks variable which is expressed in logarithms before differentiating it. This is for ease of comparison with related empirical studies, particularly with Känzig (2021) and Degasperi et al. (2025).

[11] This is consistent with the literature, as the pseudo-causal VAR(p) uses $p = r + s = 24$ (see Kilian and Zhou, 2023).

long horizons in the future. This strategy allows us to isolate the anticipated component of the targeted variable from the unanticipated shocks. Standard identification restrictions on the reduced-form error terms from the NC-VAR implied by equation (27) are imposed:

$$\epsilon_t = \omega_t^{-\frac{1}{2}} \eta_t = \bar{B} \bar{u}_t \tag{33}$$

where the t-distributed structural shock vector, $\bar{u}_t = \omega_t^{-\frac{1}{2}} u_t^* = [\bar{u}_{1,t} \cdots \bar{u}_{k,t}]' \sim t_\lambda(I_k)$ is a product of two latent factors, a k-dimensional vector of Gaussian shocks $u_t^* \sim N(0, I_k)$ and the volatility term $\omega_t^{-\frac{1}{2}}$. Denote by W in equation (33) an orthogonal $(k \times k)$ matrix with w_i on its ith column, and $\tilde{A}$ a Cholesky factor such that $\Sigma = \tilde{A}\tilde{A}'$. The rotation matrix $\bar{B}$ is thus now given by $\bar{B} = \tilde{A}W$.

An additional identifying assumption is the imposition of zero restrictions on the lead coefficient matrices of the noncausal polynomial. Specifically, we constrain the noncausal polynomial $\Phi(L^{-1})$ to have the following block-triangular structure:

$$\Phi(L^{-1}) = I - \Phi_1 L^{-1} - \ldots - \Phi_s L^{-s}, \quad \Phi_i = \begin{bmatrix} 0 & \mathbf{0}_{1\times(n-1)} \\ \Phi_{21,i} & \Phi_{22,i} \end{bmatrix}, \quad \phi = R_\phi \phi_r, \quad i = 1, \ldots, s \tag{34}$$

where $\phi = \text{vec}(\Phi)$ denotes the vectorized lead coefficients, ϕ_r is the $((n^2 s - s^*) \times 1)$ vector of unrestricted parameters, and R_ϕ is the $(n^2 s \times (n^2 s - s^*))$ selection matrix that maps ϕ_r to ϕ while enforcing the $s^* = n \times s$ zero restrictions.

This restriction serves two purposes. First, by imposing zeros in the first row of each lead coefficient matrix Φ_i, we ensure that global oil production does not depend on future values of any variable in the system. This guarantees that oil production is not a forward-looking variable and that the structural shock identified in the first position of the system corresponds exclusively to an oil supply shock, rather than being contaminated by anticipation effects from other variables. Second, we extend this restriction to the second row by setting zeros in the position corresponding to world industrial production (i.e., we also restrict the $(2, 1)$ block to zero when relevant). This prevents the real oil price—which occupies the third position in our system—from being polluted by any anticipated

demand-side dynamics that could otherwise appear through the lead structure of the NC-VAR. In other words, these restrictions ensure that the forward-looking behavior captured by the noncausal component is channeled exclusively through the price and inventory variables, which are the natural forward-looking variables in the oil market, while production and economic activity remain backward-looking as economic theory would suggest. The technical implementation of these restrictions within the Bayesian estimation framework is detailed in Appendix D.[12]

Using the two-sided MA representation of the NC-VAR model from (25), the time-t forecast revision of expected global oil production at horizon τ is given by:

$$\mathbb{E}_t[wop_{t+\tau}] - \mathbb{E}_{t-1}[wop_{t+\tau}] = e_1' \Psi_\tau \tilde{A} W \bar{u}_t \tag{35}$$

with the vector $e_1 = [1\ 0\ 0\ 0]'$ selecting global oil production. Our objective is to identify the first shock as that which best explains the variance of forecast revisions at horizons between H_1 and H_2. Following Chahrour et al. (2023), we seek to find:

$$w_1^{news} = \underset{w_1}{\text{argmax}} \sum_{\tau=H_1}^{H_2} e_1' \Psi_\tau \tilde{A} w_1 w_1' \tilde{A}' \Psi_\tau' e_1 \quad \text{s.t.} \quad w_1' w_1 = 1 \tag{36}$$

This optimization problem can be solved as an eigenvalue problem. Specifically, the optimal w_1 is the eigenvector associated with the largest eigenvalue of the matrix:

$$S = \sum_{\tau=H_1}^{H_2} \left(\Psi_\tau \tilde{A}\right)' e_1 e_1' \left(\Psi_\tau \tilde{A}\right) \tag{37}$$

In contrast with the Barsky and Sims (2011) identification procedure, we do not impose orthogonality of the shock to current global oil production. This choice is motivated by two considerations, as discussed in Chahrour et al. (2023) and Kurmann and Sims (2021). First, due to the two-sided MA representation (25), identifying restrictions on the impact effects are generally difficult to implement in the noncausal VAR as the timing of the

[12] Figure 13, in Appendix E, shows the impulse response functions of the baseline model when the zero restrictions are relaxed (i.e., $R_\phi = I_{n^2 s}$). The results show that the main findings remain qualitatively unchanged.

shock is unknown a priori. Second, relaxing the impact effect restriction is potentially more robust to measurement errors and revisions in global oil production data as well as to small sample bias involved in long-run identification schemes. However, following Chahrour et al. (2023), we set $H_1 > 0$ to avoid designing an identification scheme that actively rewards nonzero impulses early on in the response, which would tend to mix short-term surprises and longer-term anticipated fluctuations in global oil production. We discuss in the next section our choice of the finite truncation horizons H_1 and H_2. This medium-run identification strategy is essentially the same as the *Max-Share* identification proposed by Uhlig (2004) and Francis et al. (2014), adapted here following Chahrour et al. (2023).

5 Baseline evidence on oil supply news shocks

In this section, we present the results of the impulse response function analysis of the oil supply news shock. Recall that in this particular case, the key advantage of the NC-VAR model lies in providing an attractive solution to the underlying problem of non-fundamentalness.

Our first step is to check whether a non-fundamental representation is supported by the data. As pointed out by Lanne and Saikkonen (2013), Gourieroux and Jasiak (2017) and Davis and Song (2020), the non-Gaussian assumption of the error term is a necessary and sufficient condition for uniqueness and is thus required in the NC-VAR estimation. Figure 1 plots the estimated marginal posterior density of the degrees of freedom (DOF) λ. The histogram indicates strong evidence in favor of fat tails as the posterior density is centered around 3.75 degrees of freedom; with the probability of λ being greater than 6 at almost nil.[13] This suggests that the normality assumption is inappropriate, confirming the choice of multivariate t-distribution for the error term, which makes it possible to identify a unique NC-VAR(r,s) specification. Therefore, noncausal representation is both statistically preferred and necessary for the unique identification of structural shocks.

[13] We set the prior mean of λ to 8. With a very low posterior mean of λ (less than 5), data dominate the assumed prior mean of λ.

Figure 1: Posterior density of DOF parameter λ

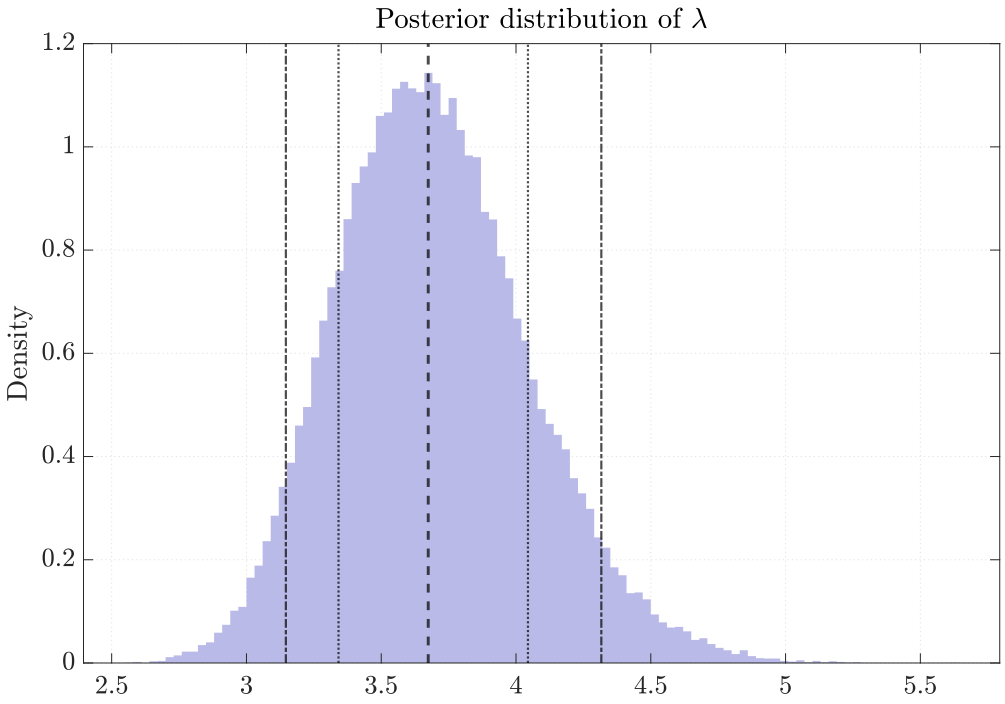

Notes: grey bars represent the frequency distribution of the DOF parameters from the 4-variable NC-VAR(12,12).

This result aligns with the theoretical arguments of Chahrour and Jurado (2021) and Gouriéroux et al. (2020) that non-fundamentalness is empirically relevant when expectations react before physical quantities.

5.1 How does the oil supply news shock diffuse to the oil market variables?

Figure 2 shows the cumulative impulse responses to an oil supply news shock which is normalized to correspond to a 10 percent change in the real oil price on impact. This normalization does not impose any sign restriction. The direction of the price response is entirely determined by the data through the Max-Share identification. Light and dark shaded bands represent 90 and 68% posterior credible sets, respectively. In the noncausal model context, the left side of the x-axis is added to represent responses related to the lead terms of the MA representation. The news shock is measured with a truncation horizon of $[H_1 = 24, H_2 = 24]$ but similar results, reported in Section 5.3, are obtained

with shorter and longer truncation horizons.

The choice of $H1 = 24$ and $H2 = 24$ for the Max-Share identification is economically justified by the institutional realities of global oil supply dynamics and the rational-expectations storage mechanism central to the paper. There are different reasons to do so. First, there is medium-run supply adjustment horizon. Oil production responds to news with 2-3 year lags, not instantaneously. Geopolitical shocks (Iran Revolution, Gulf War) or OPEC quota announcements affect extraction decisions with multi-quarter implementation delays due to drilling rig mobilization, field maintenance, and contractual rigidities. Moreover,empirical evidence shows the identified shock triggers a -4% production drop peaking at 2-3 months but stabilizing over 24 months—precisely the window where news about future supply should dominate forecast revisions. A 24-month horizon targets this implementation lag where agents update supply expectations before quantities fully adjust, avoiding short-run noise (H1=3 confounds surprises) or ultra-long-run contamination (H2=60 mixes permanent trends like shale).

Second, there is inventory precautionary motive. Global oil stocks act as the forward-looking buffer validating news shocks. Rational storage theory (Baumeister and Hamilton, 2019; Kilian and Murphy, 2014) predicts agents build inventories when expecting future scarcity within their planning horizon (1-3 years). H1=H2=24 captures the peak precautionary accumulation (+3% stocks on impact), which requires knowing supply will be tight over the next two years—not 5+ years (H2=60) where storage costs dominate. Shorter windows miss this dynamic; symmetric 24 perfectly aligns with observed stock build-up before production falls.

Third, oil prices reflect intertemporal scarcity pricing: a credible 2-year supply shortfall justifies +10% spot price jumps via convenience yield and risk premia, without requiring immediate quantity disruption. Futures curves embed expectations over 1-3 years; H1=24 isolates this news premium from high-frequency speculation (too short) or secular trends (too long).

Finally, all horizons yield similar price IRFs, but 24-month delivers tightest credible sets, confirming economic relevance over statistical convenience. H1=H2=24 targets

Figure 2: Impulse-response functions to the oil supply news shock

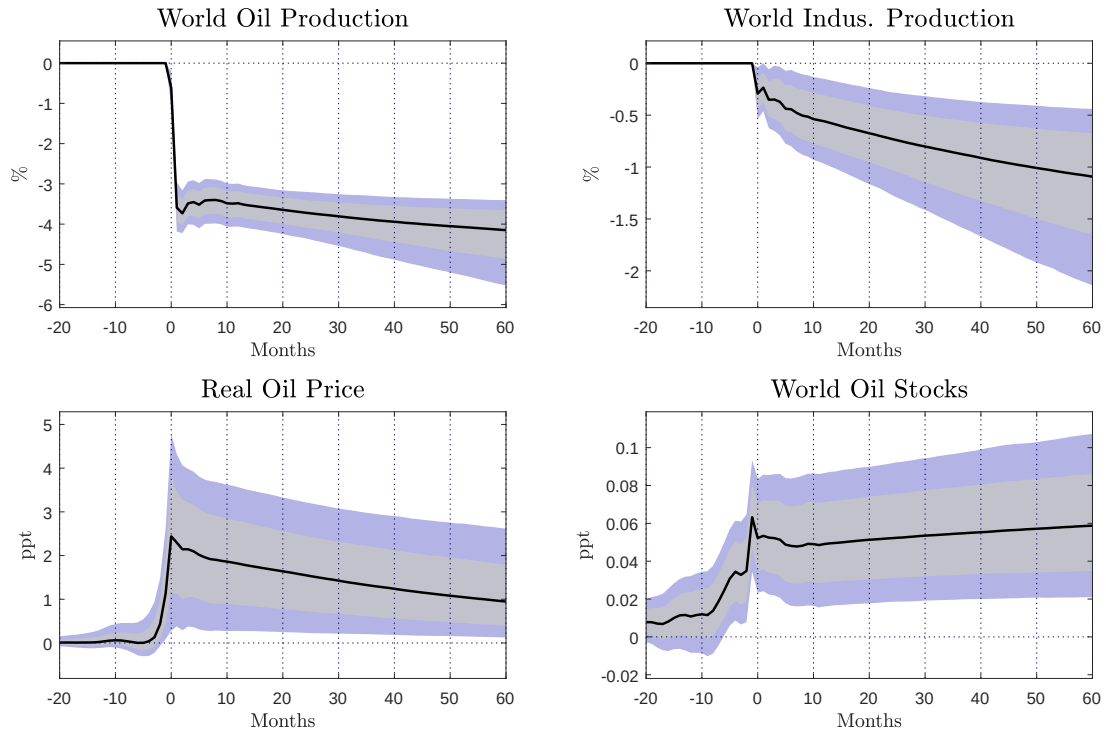

Notes: The black solid lines are the posterior median responses of the 4-variable baseline model from Baumeister and Hamilton (2019). The solid lines are the posterior median cumulated impulse responses of the NC-VAR(12,12). Shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation of NC-VAR.

OPEC-style news shocks—coordinated quota changes with 1-2 year horizons that dominate historical decompositions (1970s-1980s, 2014-2016). This matches the institutional fact that most economically-relevant oil supply news arrives via strategic announcements, not geological surprises, making the identification both data-driven and economically interpretable.

By construction, global oil production exhibits no response at leads under the Max-Share identification, while a contemporaneous response is allowed. The post-shock dynamics are highly informative: although an impact response is allowed, global oil production does not adjust instantaneously. Instead, it declines with a short delay, reaching approximately 4 percent within two to three months and then stabilizing at roughly that level. This pattern reflects the information captured by the Max-Share identification, which is entirely inferred from the observed production behavior. Given the speed of the response, the shock is likely driven by news that can influence current extraction decisions-such as

geopolitical events, strategic announcements (Känzig, 2021; Pinchetti, 2025) , or temporary supply disruptions-rather than long-term changes in production capacity arising from new discoveries (Arezki et al., 2017) or technological innovations. Global real activity also displays a delayed but long-lived contraction, reaching a trough near 1 percent only after several years, which is exactly what one expects from a news-driven supply disturbance that first works through precautionary pricing and intertemporal demand reallocation before propagating to the real side of the world economy.

But relevant to our analysis, there is clear evidence from the estimated impulse responses that a substantial proportion of the oil supply reduction triggered by the news shock can be anticipated before the drop materialises. The median response estimates from the noncausal model suggest large and significant responses of noncausal components for the most forward-looking variables, namely the real oil price and global oil stocks. As for the causal part, after rescaling the impulse responses, the real oil price increases by 10 percent on impact and remains about 5 percent above baseline over the medium run. Global oil stocks increase on impact by roughly 3 percent (0.06 ppt) and continue to trend up, confirming the interpretation of the identified disturbance as a genuine news shock: agents anticipate higher future scarcity and build precautionary stocks rather than drawing them down as they would after an unanticipated physical disruption.

Overall, Figure 2 shows that once noncausality is allowed, the impulse responses line up with a rational-expectations storage model with anticipated supply shortfalls: prices and inventories predictably move first, physical quantities follow with lags, and real activity reacts more sluggishly but persistently.

5.2 How much the oil supply news shock explains real oil price fluctuations in different historical episodes?

The question of the origin of oil price fluctuations remains central to the global oil market analysis (Kilian, 2009; Juvenal and Petrella, 2015; Baumeister and Kilian, 2016; Baumeister and Hamilton, 2019; Caldara et al., 2019; Känzig, 2021, among others). In this section

we rely on the historical decomposition exercise (see Section 3) to get a sense of the overall importance of oil supply expectation shocks in driving real oil price fluctuations over the last few decades, and particularly for specific historical episodes.

Figure 3, illustrates the cumulative effect of the oil supply news shocks together with the actual real oil price for the period 1975-2025. The figure reveals that oil supply news shocks account for a sizeable and time-varying share of real oil price movements over the last five decades. More precisely, examining specific historical events, the contribution of oil supply news shocks to the sharp rise in oil prices following the Iranian Revolution was strongly negative in the mid to late 1970s. This was followed by a substantial positive contribution of supply news around the Iran-Iraq War and until late 1985, even as the real price gradually declined, indicating that other shocks dominated the overall price movement. During the period 1980-1985, the non-OPEC production increased by about 15% (Gately, 1986). These increases were made by a large number of relatively small producers whose exploration and development activities had been stimulated by the price rises of the 1970s, and who had grown sufficiently in size and incurred sufficient fixed costs not to be discouraged by the price declines of the early 1980s. Thus, the contribution of expectations of an increase in the oil supply due to new discoveries is at least as important as the expectation of OPEC's collapse, when Saudi Arabia abandoned production constraints in late 1985 (Kilian and Murphy, 2014; Känzig, 2021). Around the time of the 1986 price collapse, the contribution of news became strongly negative, reflecting upward revisions to expected future supply arising from both non-OPEC expansions and the erosion of OPEC's ability to sustain high prices.

By contrast, during the Gulf War of 1990–91, the news played only a modest role in explaining the temporary spike. This is consistent with a narrative in which market participants anticipated compensating production from core OPEC, meaning they did not revise long-term supply expectations as dramatically as spot prices suggested. Oil news shocks seem to be better able to explain the trend reversal that took place in the late 1990s and ended with the onset of the Venezuelan crisis in late 2002. Despite a brief positive impact prior to the collapse of 2008, the influence of oil supply news remained

Figure 3: Contribution of oil supply news shocks to real oil price fluctuations

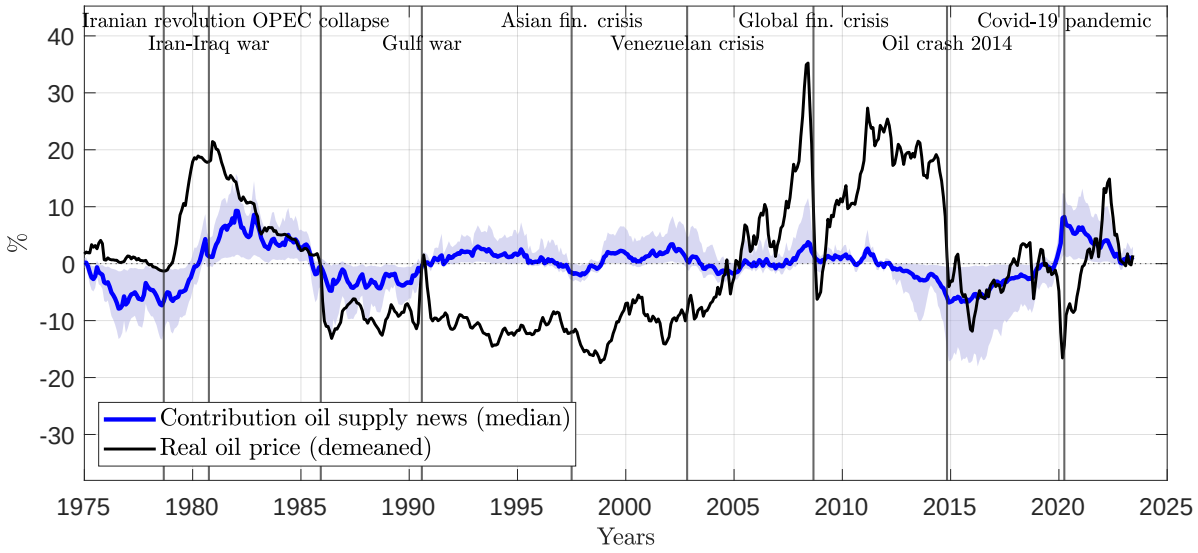

Notes: The solid blue line shows the average contribution of the oil supply news shock to the real crude oil price, estimated using the NC-VAR(12,12). The light blue shaded area represents the 90% credible intervals. The solid black line depicts the demeaned real oil price.

subdued from the mid-2000s until the end of 2013. After this period, its influence turned strongly negative around the time of the 2014 oil price crash.

Crucially, the events in the global oil market at the onset of the COVID-19 pandemic in March and April 2020 constitute a key milestone. The unprecedented collapse in oil prices was driven entirely by an unexpectedly severe contraction in global oil demand. This episode provides a natural validation check for our shock identification, since oil supply news shocks should be positive but should not account for the observed price collapse. Consistent with this interpretation, the historical decomposition shows that the early-2020 price collapse was entirely demand-driven. Although the contribution of oil supply news is positive, it does not explain the sharp decline in prices. This supports the identification strategy, as an unexpected global demand freeze should not be misclassified as news about future oil supply. Finally, oil supply news shocks play a meaningful role in explaining the subsequent price recovery following the pandemic, when oil-producing countries were expected to cut production in response to the demand contraction.

5.3 What do truncation horizons reveal about oil market information?

As noted above, our results are robust to alternative choices of the truncation window and to different values of H_1 and H_2 used to identify oil supply news shocks in the baseline model. Figure 4 evaluates this robustness by considering a range of Max-Share truncation windows, namely $(H_1, H_2) \in \{(6,24),(6,60),(12,24),(12,60),(24,24),(24,60),(60,60)\}$. Across these specifications, the responses of global oil production, real economic activity, and oil inventories remain remarkably stable in terms of sign, timing, and persistence. This stability indicates that the informational content of the identified shock-extracted from observed oil production-is robust and does not depend on a particular choice of truncation horizons.

Figure 4: News shock identification with different truncation windows (H_1; H_2)

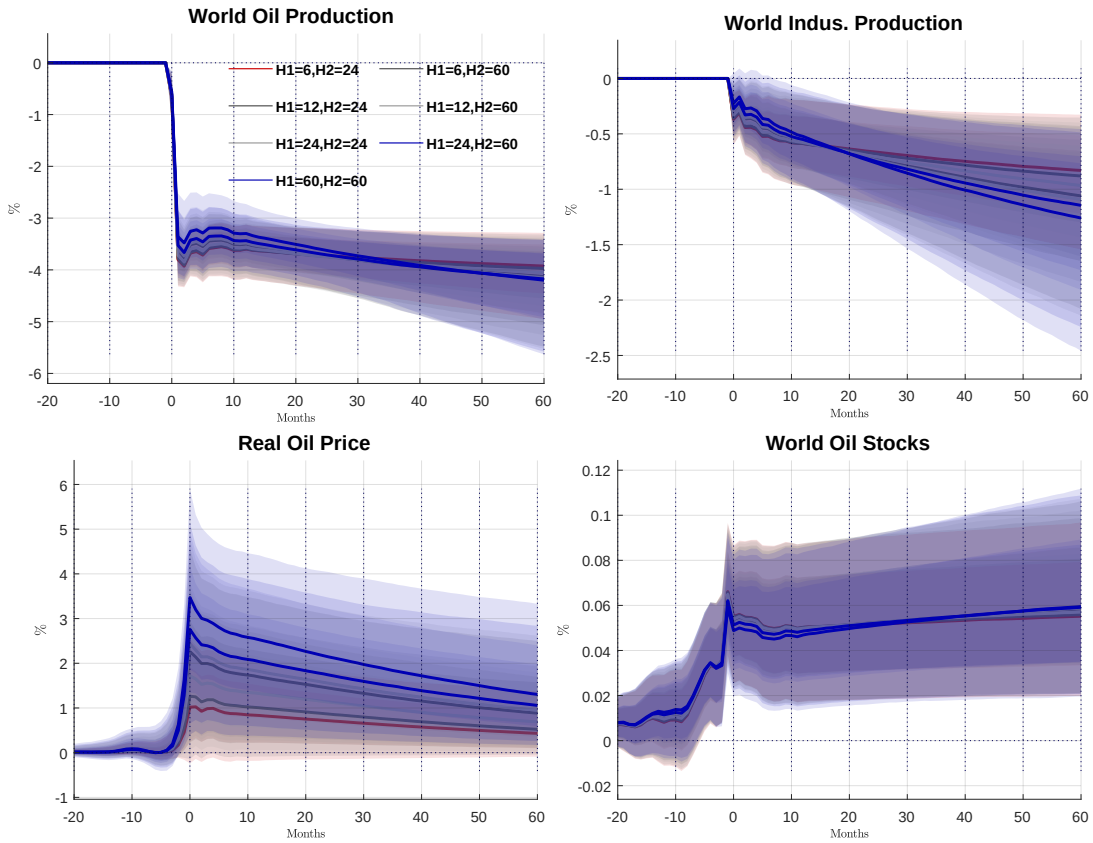

Notes: The black dashed lines are the posterior median responses of the 4-variable baseline model. The solid lines are the posterior median cumulated impulse responses of the NC-VAR(6,6). Light and dark grey shaded regions are the 90 % and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation of NC-VAR.

Although oil inventories are forward-looking, their responses remain largely invariant to alternative choices of the Max-Share window. This pattern is consistent with rational-expectations storage models, in which inventories adjust smoothly to anticipated scarcity, subject to physical, institutional, and strategic constraints. By contrast, the response of the real oil price is more sensitive to the truncation window. This sensitivity is economically intuitive and informative: The Max-Share approach is designed to identify shocks that are predictive of future oil production, rather than to target contemporaneous price movements. As a forward-looking valuation variable, the real oil price reflects how markets price information about future supply conditions at different horizons. Consequently, variations in (H1,H2) affect the timing and magnitude of the price response by shifting the horizon over which the informational content of the shock becomes most relevant, without altering the underlying real dynamics. The hump-shaped response peaks early and decays gradually, with little sensitivity to horizon choice—shorter windows (H1=6) yield slightly sharper peaks, while longer ones (H2=60) produce smoother but equally persistent effects. Across all specifications, the price responses are statistically significant at the 90% credible set (dark shaded areas) from leads -10 through +40 months, with 68% credible sets (light shaded) confirming pointwise significance over even longer horizons. This tight banding around the median trajectory underscores the robustness of the forward-looking price reaction to identification assumptions.

Sensitivity checks: Figure 14 in Appendix E show that re-estimating the model using a smaller sample size produces nearly identical benchmark results and continues to support the interpretation of the identified shock as an oil supply shock. In particular, impulse responses over the 1974:2–1989:12 sample are virtually indistinguishable from those of the benchmark specification. A negative shock to oil supply expectations induces a contraction in oil production, a gradual and persistent decline in global real activity, and a pronounced increase in both the real oil price and global oil inventories. However, when using the 1990:1–2025:7 sample, oil production and global real activity continue to respond negatively and significantly. Conversely, the positive responses of the real oil price and global oil stocks become statistically weaker. It should be noted that, this result is

unsurprising and can be explained by the declining contribution of oil supply news shocks to real oil price fluctuations in the second subsample, as illustrated by the historical decomposition in Figure 3 discussed the previous section. Moreover, this exercise shows that our results do not lead to puzzles, a problem pointed out by Degasperi et al. (2025) when estimating the effects of the oil supply news shock of Känzig (2021) for shorter samples. Finally, as we show in the next section, our results are robust to the inclusion of additional macroeconomic and financial variables in the system.

5.4 What impact does a news shock have on global and US macroeconomic variables?

The macroeconomic consequences of a negative exogenous oil supply shock have been extensively debated, both theoretically and empirically, particularly in the context of the debate on the role of oil shocks in generating stagflation (Hamilton, 1983, 2009; Gisser and Goodwin, 1986; Barsky and Kilian, 2002, 2004, among others). However, there are fewer studies on the effect of oil supply expectation shocks on macroeconomic variables (Arezki et al., 2017; Känzig, 2021; Degasperi et al., 2025). In this section, we first examine how this shock affects the global economy more broadly, before looking at the US variables.

Figure 5 extends the analysis to global macro-financial indicators, namely the Global Economic Conditions (GECON) indicator of Baumeister et al. (2022), the OECD passenger car registrations (PCR_OECD), the global supply chain pressure index (GSCPI), the OECD consumer confidence index (CCI_OECD), geopolitical risk index (GPR) and as a financial index, the world stock price index (MSCI).[14]

It is interesting to note that all global variables, except for the MSCI react significantly to the negative oil supply news shock. GECON gradually declines, reaching around -1 percent after roughly two years and remains depressed, echoing the slow but sizeable contraction in global industrial production. We observe roughly the same reaction from the passenger car registration, which is in line with the sensitivity of durable consumption and

[14]The additional variables are included in the system, one variable at a time in order to avoid estimating a large model.

Figure 5: Reactions of global macroeconomic variables to a news shock

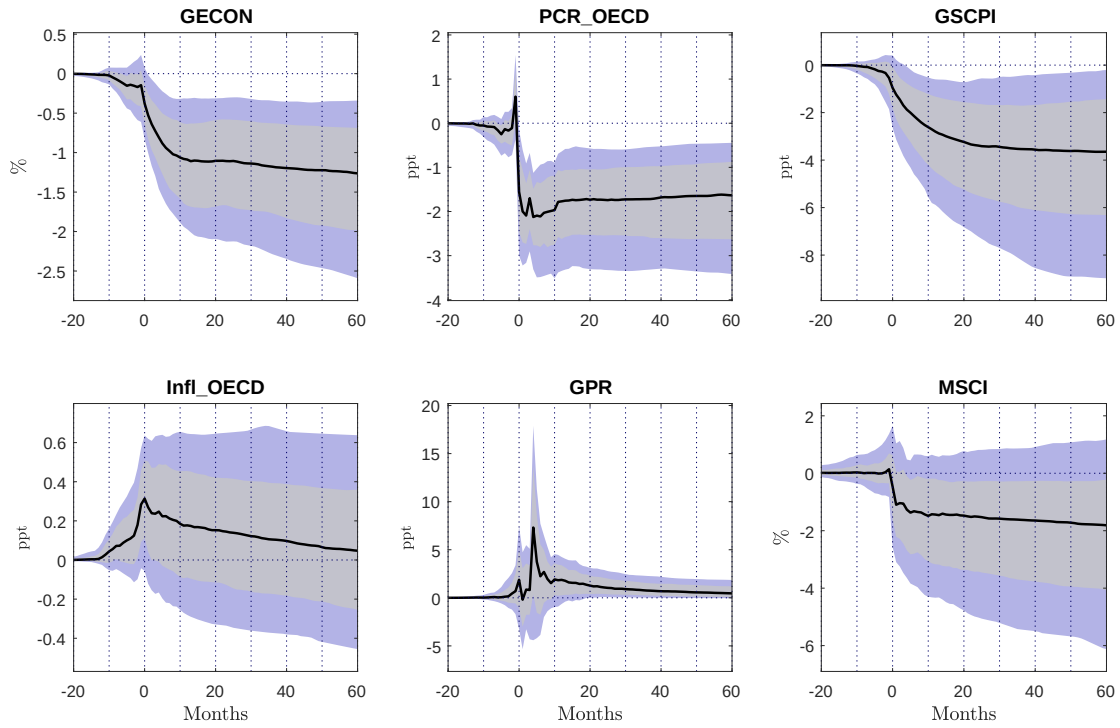

Notes: The solid black lines are the posterior median cumulated impulse responses of the NC-VAR(12,12) and shaded areas represent the 90 and 68 % credible sets. Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation of NC-VAR.

particularly car demand to expected future fuel costs and heightened macro-uncertainty. The behavior of the Global Supply Chain Pressure Index (GSCPI) is particularly informative, as its response to anticipated oil supply shocks differs fundamentally from its response to unanticipated supply disruptions. The GSCPI declines gradually and persistently, indicating that news about future supply shortfalls-unlike unexpected supply disruptions-may reduce global supply chain pressures by inducing forward-looking adjustments in production, sourcing, and trade volumes. These anticipatory adjustments help smooth logistics flows and alleviate congestion before the shock materializes.

The geopolitical risk index (GPR) spikes on impact by about 5 points and then gradually declines, remaining persistently around 1 percentage point. However, the response becomes statistically significant only after approximately two years. This pattern suggests a two-way interaction: geopolitical tensions often generate oil supply news, but once oil prices and inventories adjust, they in turn validate and reinforce perceptions of geopolitical risk. Consumer confidence drops immediately and significantly, with the ef-

fect starting already in the lead terms and stabilizing at a negative level. Finally, the global stock price index MSCI does not react on impact but subsequently falls sharply by about 1 percent, although the effect of the shock remains statistically insignificant. The effect of anticipated oil supply shocks on stock prices appears to be as negligible as that of unanticipated shocks (Kilian and Park, 2009).[15] This pattern may reflect the anticipatory nature of oil supply news shocks, the aggregation of offsetting effects across oil-importing and oil-exporting economies, and the fact that such shocks tend to tighten financial conditions without posing an immediate threat to global financial stability.

Figure 6: Reactions of US macroeconomic variables to a news shock

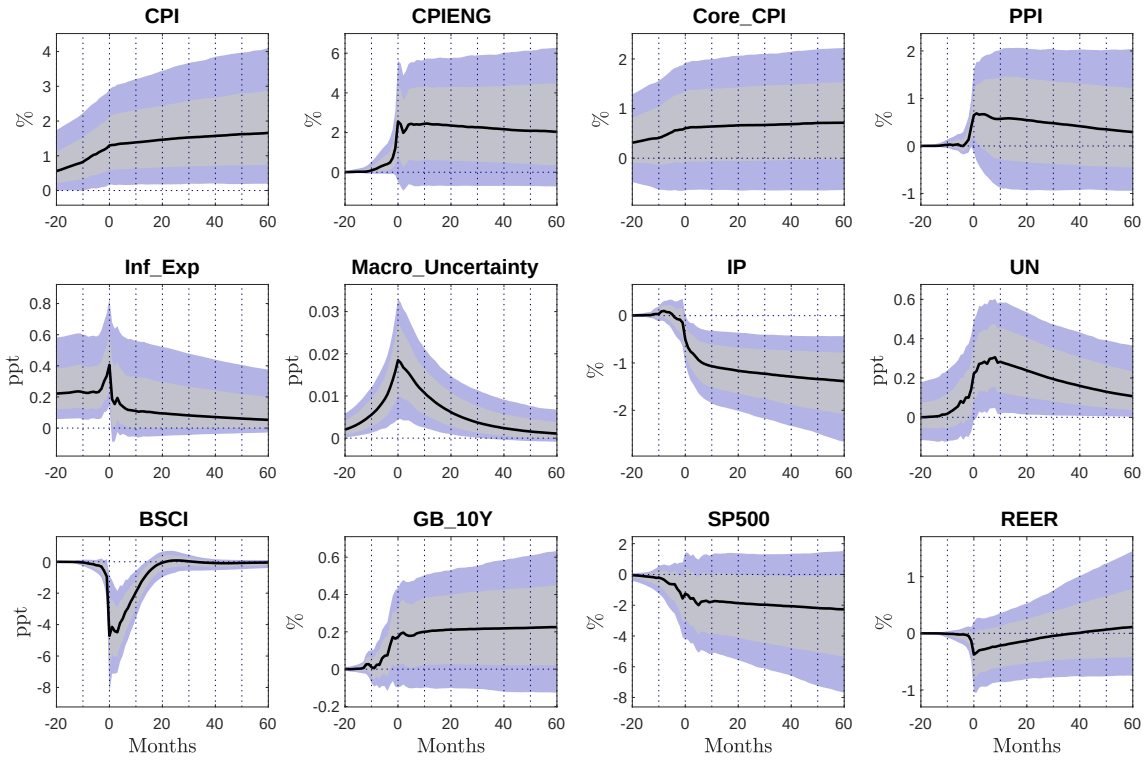

Notes: The solid lines are the posterior median cumulated impulse responses of US macroeconomics variables. Blue and grey shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation.

As for the effect of the oil supply news shock on the US economy, Figure 6 shows the reactions of US macroeconomic variables[16] belonging to different categories spanning

[15] It is worth noting that, according to Mumtaz et al. (2018), there is evidence of nonlinear stock price dynamics in response to unanticipated oil supply shocks. Mumtaz et al. (2018) show that the stock price reaction can be significant during the regime characterized by low oil inflation.

[16] All variables come from the FRED database, except for the excess bond premium which is an updated version of the measure of Gilchrist and Zakrajšek (2012) available from the Fed website: https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html

multiple dimensions of the US economy.

As shown in the top panel of Figure 6, all price variables except core CPI increase sharply on impact. Headline CPI increases already in the leads, peaks at around 1 percent on impact, and remains roughly at that higher level thereafter. The Energy CPI shows anticipatory movement, rising to 2% on impact and remaining at this level, albeit with a statistically insignificant response. The core CPI, however, remains muted from the anticipatory to the causal part. This pattern confirms that energy prices constitute the primary transmission channel, while second-round effects gradually spread to non-energy components over time. The producer price index (PPI) begins to respond in the leads, with an impact effect of around 0.5 percent, but quickly becomes statistically insignificant, signaling modest cost pressures in the production sector. As for inflation expectations (Inf_Exp), they exhibit a highly significant anticipatory response, rising on impact but remaining relatively short-lived. This pattern suggests that agents revise their expectations in the expected direction, while monetary policy credibility and the perceived transitory nature of the disturbance prevent a persistent de-anchoring of inflation expectations.

Macroeconomic uncertainty increases significantly already at its leads, reflecting the anticipatory component of the shock, and peaks on impact at about 0.02 percentage points before gradually declining thereafter. US industrial production (IP) declines after a short delay, mirroring the global real activity response and indicating that the real side of the economy adjusts sluggishly to the cumulative drag from higher energy costs and weaker demand. The unemployment rate (UN) increases by about 0.3 percentage points on impact and then gradually declines. The responses of price and activity variables suggest that anticipated oil supply shortfalls can simultaneously trigger inflationary and recessionary effects, producing a negative co-movement between CPI and real activity, consistent with stagflation. Consequently, this shock presents significant challenges for monetary policymakers, given the negative trade-off between inflation and economic activity.

Business confidence (BSCI) declines significantly by about –1 percent on impact, beginning to recover and returning to its initial level after roughly twenty months. The 10-year government bond yield (GB_10Y) rises persistently, though not significantly, in-

dicating that term premia and expected future short rates respond to the inflationary component of the shock despite the anticipated drag on economic activity. Consistent with the MSCI index, the S&P 500 exhibits a similarly negligible response to anticipated oil supply shocks, reinforcing the finding that equity markets remain unaffected by such supply news. Finally, the real effective exchange rate (REER) shows no statistically significant dollar appreciation, whereas, as a net oil importer, the United States would be expected to experience a currency depreciation following a supply shock that raises oil prices. This effect may be muted by the fact that oil is predominantly traded in U.S. dollars, as well as by the United States' transition between net oil exporter and net oil importer during the sample period.

In summary, these results suggest that news about future oil supply shortfalls can simultaneously generate inflationary and recessionary effects, while having only a limited impact on financial markets.

6 Two worlds of oil: Anticipation across OPEC and non-OPEC productions

While the global oil production delivers a clear supply-news shock, it remains silent on the underlying sources of this information. To shed light on the origin of the identified news, we reapply the same identification strategy to disaggregated production measures.

The distinction between OPEC and non-OPEC oil production is well grounded both institutionally and empirically. OPEC producers operate under a coordinated framework that explicitly targets supply management through production quotas and strategic adjustments, whereas non-OPEC producers are largely driven by market-based incentives and technological constraints, leading to fundamentally different supply responses.[17]

A growing body of empirical work documents these structural differences and their implications for oil market dynamics. Extending the framework of Pierru et al. (2018, 2020),

[17] see, e.g., the discussion provided by the U.S. Energy Information Administration: https://www.eia.gov/finance/markets/crudeoil/supply-opec.php

Almutairi et al. (2023) show that coordinated supply management by OPEC+ reduced oil price volatility by up to 50 percent, both before and during the COVID-19 pandemic. Importantly, they attribute most of this stabilization effect to OPEC's own actions, while the role of the Allies was primarily to support the price level rather than to dampen volatility. More recently, Baumeister and Hamilton (2024) emphasize the importance of accounting for heterogeneous supply and demand behavior when identifying oil market shocks. They show that fluctuations in Saudi Arabian oil production, together with endogenous inventory adjustments, have historically played a central role in stabilizing the global price of oil.

Taken together, these findings suggest that changes in OPEC production embed a substantial anticipatory component, reflecting coordinated and credible commitments to future supply management. Exploiting the differential behavior between OPEC and non-OPEC production therefore provides a natural and informative strategy for identifying oil supply news shocks, as anticipated supply developments are more likely to be transmitted through OPEC production decisions than through the largely market-driven responses of non-OPEC producers.

Figure 7 juxtaposes the responses from NC-VARs where the news shock is identified either from OPEC production or from non-OPEC production, and the contrast is highly revealing about the nature of each disturbance. Under OPEC-based identification, global oil production and real activity both fall gradually and persistently, while real prices and inventories rise, closely mirroring the baseline news shock and matching the narrative that coordinated OPEC decisions about future quotas drive expectations of lower future supply. Under the non-OPEC-based identification, oil production and global economic activity also decline, but the magnitude of these responses is substantially smaller. The key difference lies in the behavior of the real oil price and global oil inventories: the real oil price responds significantly, but in the opposite direction compared to OPEC shocks, falling rather than rising and global oil inventories respond positively, but the effect is statistically insignificant. This suggests that the market does not view non-OPEC supply news as a primary driver of long-run scarcity; instead, non-OPEC expansions

Figure 7: Impulse-response functions to the OPEC and Non-OPEC oil news shock

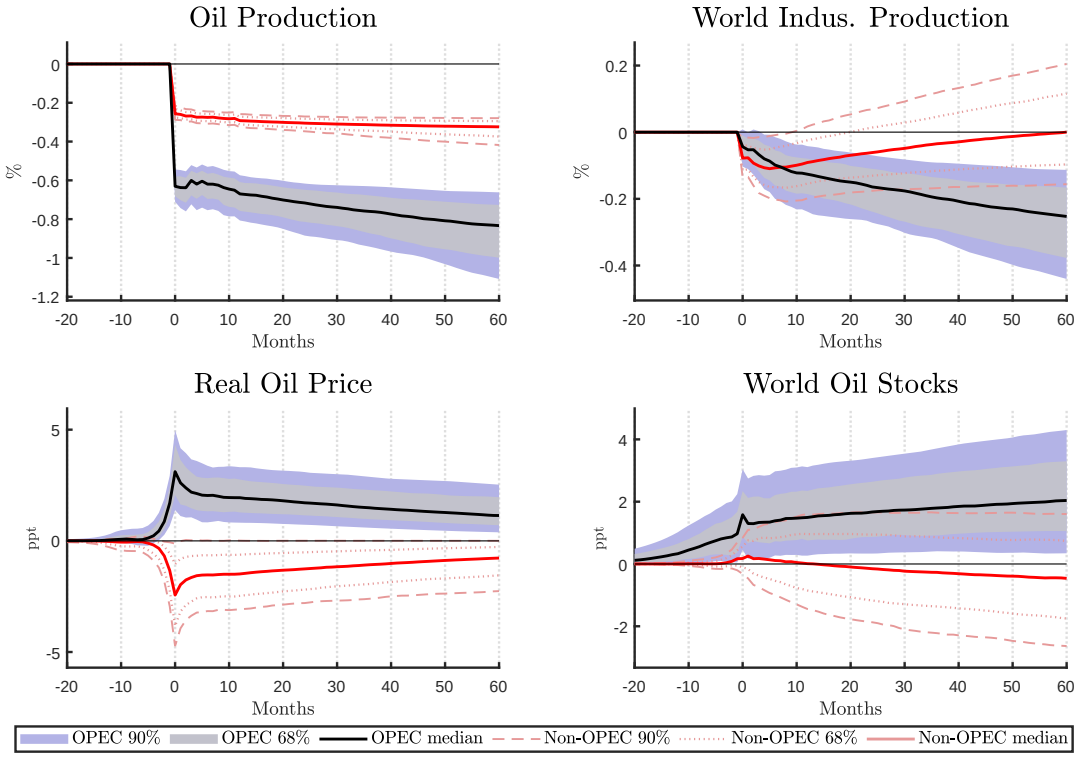

Notes: Black lines show the posterior median cumulated impulse responses of a 4-variable NC-VAR(12,12) identifying the oil supply news shock via OPEC production, while red lines show a separate NC-VAR(12,12) using Non-OPEC production. Shaded areas and dashed/dotted lines indicate 90% and 68% credible sets. Impulse responses are shown over negative and positive horizons due to the noncausal NC-VAR, with negative values reflecting lead terms of the MA representation.

often offset OPEC restrictions or respond to demand booms. Consequently, these shocks appear less consistent with a pure supply-side news shock and instead resemble a mixed, or potentially demand-dominated, disturbance. This pattern suggests that non-OPEC production responds largely endogenously to broader macroeconomic conditions rather than reflecting an exogenous supply-driven news shock.

In summary, the two specifications suggest that the Max-Share procedure captures different underlying fundamentals. For OPEC production, it isolates expectation shocks related to cartel behavior. For non-OPEC production, however, it tends to combine demand-related movements with capacity adjustments. Consequently, the global oil supply news shock is primarily an OPEC-driven phenomenon, whereas non-OPEC production conveys more information about cyclical demand than about exogenous supply shocks.

Figure 8 compares the historical contributions of OPEC- and non-OPEC-based news

shocks to real oil price fluctuations. A salient feature of the figure is that the contribution of non-OPEC news shocks is most often opposite in sign to that of OPEC news shocks. This divergence underscores their distinct economic meanings. OPEC-based news shocks account for large and persistent movements in real oil prices, particularly during episodes associated with changes in cartel coordination or anticipated supply restrictions, consistent with forward-looking supply news. In contrast, on-OPEC-based news shocks tend to offset these movements, exhibit smaller and less persistent contributions, and are more closely aligned with cyclical downturns and price declines, indicating demand-like effects rather than exogenous supply news.

Moreover, the episode surrounding the COVID-19 pandemic provides additional support for this interpretation. During this period, the collapse in oil prices is almost entirely driven by demand conditions, and the contribution of non-OPEC-based news shocks is both negative and significant. This finding offers further evidence that non-OPEC news shocks have demand-like effects and are consistent with demand-driven interpretations, rather than reflecting exogenous supply news.

Figure 8: Contribution of OPEC an Non-OPEC oil news shocks to oil price fluctuations

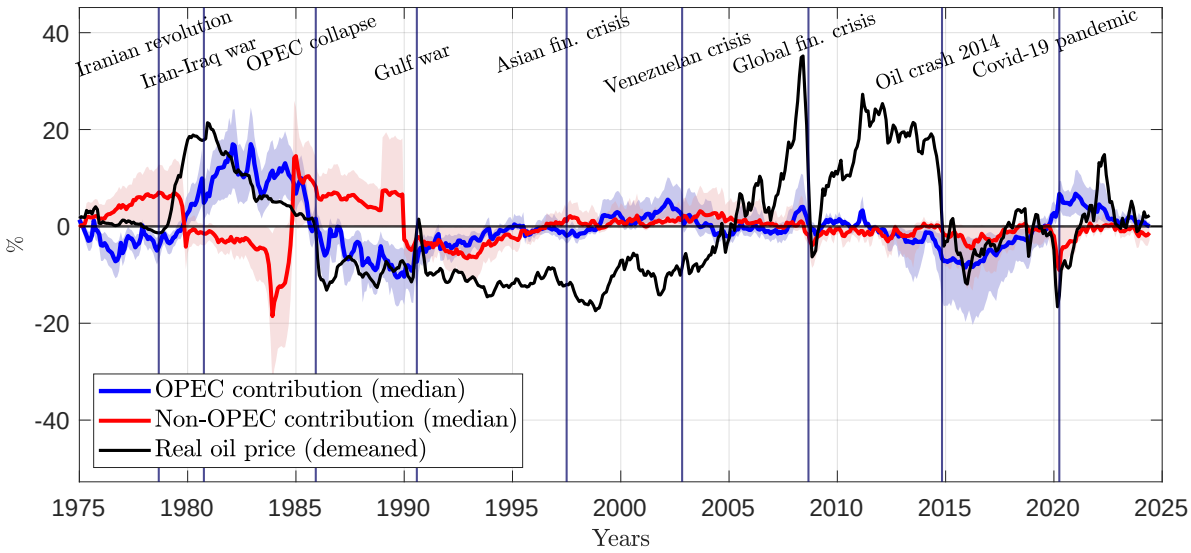

Notes: The solid blue line shows the median contribution of the oil supply news shock identified from OPEC oil production, while the solid red line shows the contribution identified from Non-OPEC oil production. Light blue and light red shaded areas indicate the 90% credible sets for OPEC and Non-OPEC shocks, respectively.

Figure 9 revisits the global macroeconomic indicators for an OPEC-based news shock and shows that the global macro-financial responses are very similar to those obtained for

the baseline shock. This finding is consistent with OPEC being the dominant source of global oil supply news. GECON and OECD passenger car registrations both respond with sizeable and persistent declines, mirroring the baseline, which shows that OPEC-related news is sufficient to generate global demand weakness in energy-intensive sectors. The global supply chain pressures (GSCPI) also decline in line with the benchmark model. OECD inflation rises more significantly than in the benchmark model, indicating that once the shock is filtered from the non-OPEC component, the effect on prices becomes more pronounced. Geopolitical risk, jumping immediately, reveals that markets view OPEC supply news as both a source and a symptom of heightened geopolitical fragility. Finally, the global stock index MSCI declines but not significantly, suggesting that although OPEC shocks constitute important global risk events, their financial impact is not persistent when viewed from a broad global portfolio perspective.

Figure 9: Reactions of global macroeconomic variables to a OPEC news shock

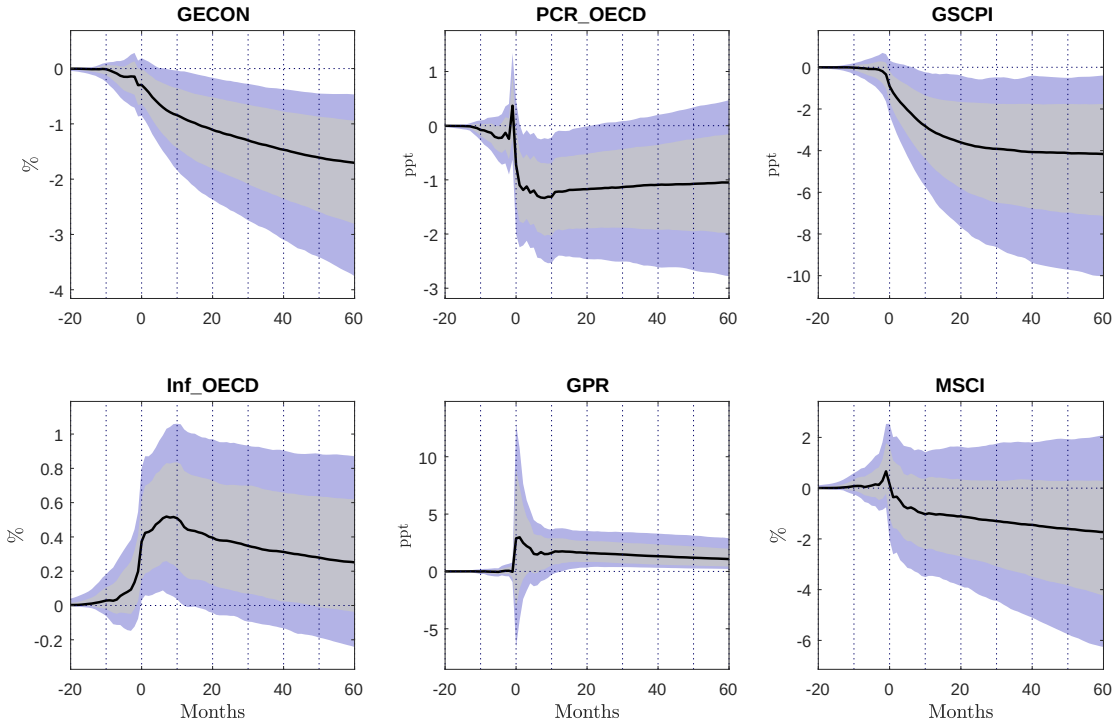

Notes: The solid lines are the posterior median cumulated impulse responses of global macroeconomics variables. Blue and grey shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation.

The close similarity between Figures 5 and 9 strengthens the argument that the base-

line news shock identified from aggregate production is effectively an OPEC-driven expectations shock at the global level.

In contrast, Figure 10 highlights the markedly different dynamic responses associated with news shocks identified using non-OPEC production. In this case, global macroeconomic reactions tend to be weaker and, in some instances, exhibit sign reversals. GECON and car registrations show smaller and sometimes even positive initial responses, suggesting that the identified shock may be picking up episodes where non-OPEC expansions accompany strong global demand rather than exogenous supply tightening. Geopolitical risk reacts only modestly, if at all, underscoring that markets do not interpret non-OPEC production surprises as major geopolitical events. Finally the global stock index response is mild and not significant, which is again consistent with a disturbance that is less tightly linked to perceived long-run scarcity.

Figure 10: Reactions of global macroeconomic variables to a NON-OPEC news shock

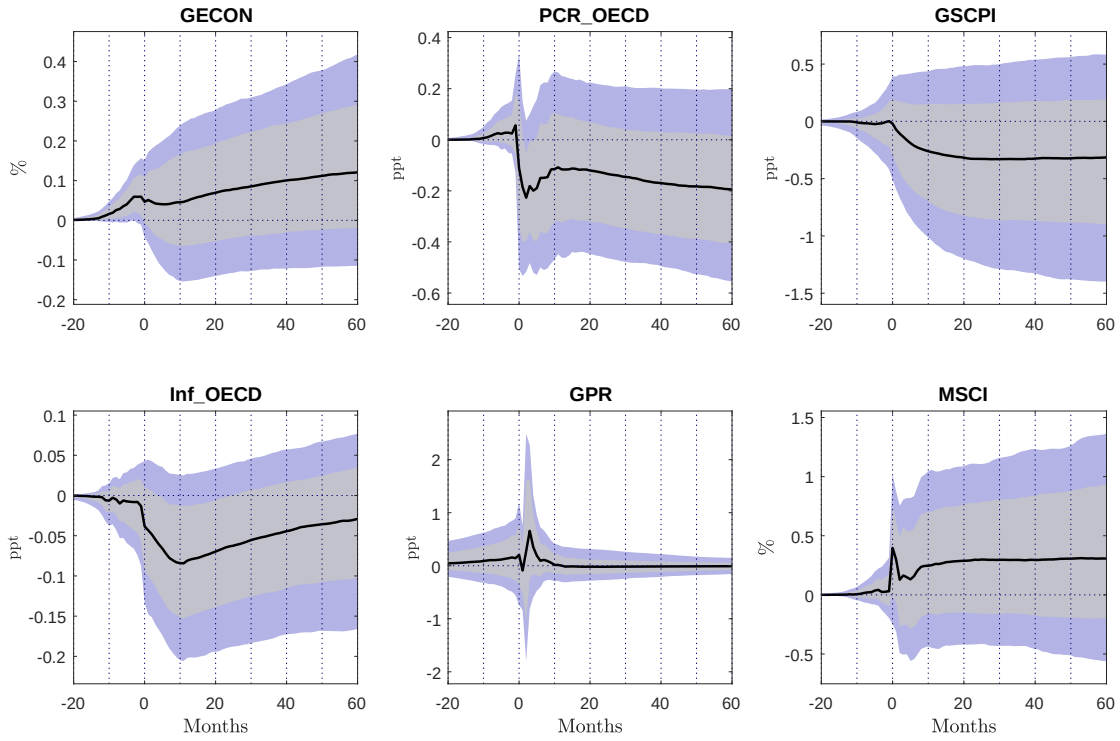

Notes: The solid lines are the posterior median cumulated impulse responses of global macroeconomics variables. Blue and grey shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation.

Figure 11 reproduces the U.S. macro-financial responses to an OPEC-based news shock

and confirms the stagflationary pattern documented in Figure 6, and clarifies that the US stagflation response is especially pronounced when the news shock is explicitly tied to OPEC.

Figure 11: Reactions of US macroeconomic variables to a OPEC supply news shock

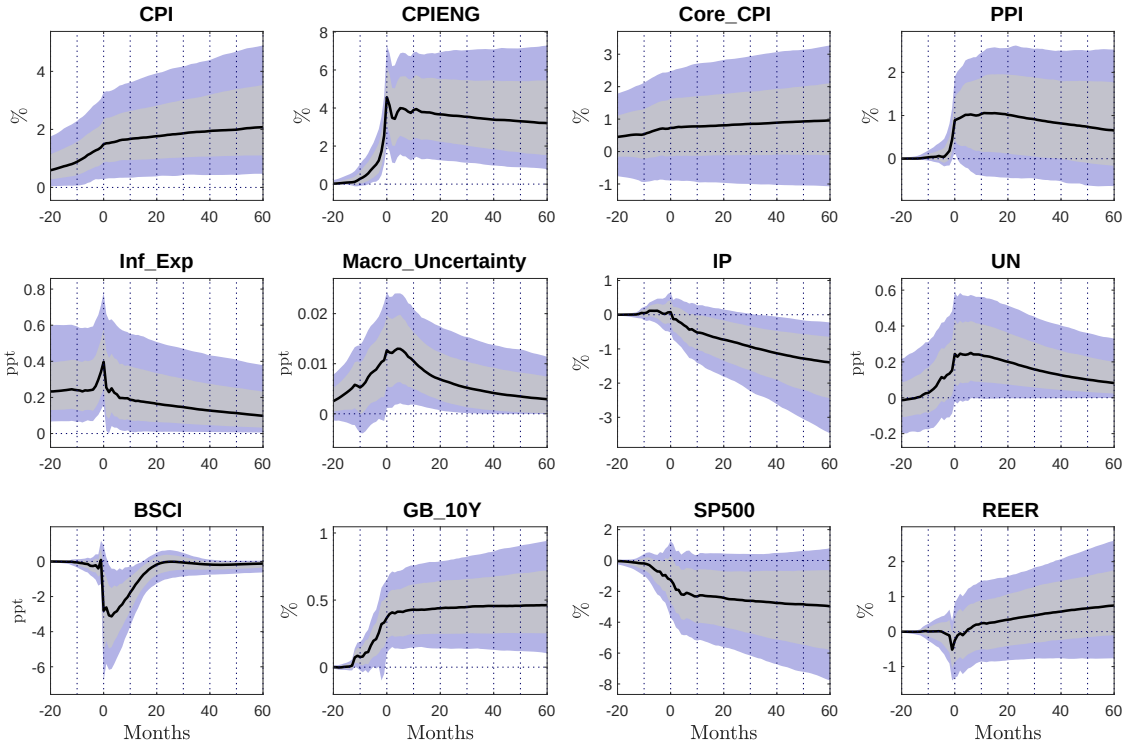

Notes: The solid lines are the posterior median cumulated impulse responses of US macroeconomics variables. Blue and grey shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation.

Headline and energy CPI, producer prices, and inflation expectations all rise following the shock, with headline and energy prices exhibiting particularly strong responses. This pattern is consistent with an environment in which OPEC-based news is rapidly transmitted to energy-intensive sectors, feeding directly into cost pressures and inflation expectations. In line with the results of the benchmark model, macroeconomic uncertainty increases and real economic activity deteriorates. Industrial production falls and unemployment rises, with lags and magnitudes that are comparable to those observed in the baseline specification. At the same time, business confidence index falls sharply on impact, underscoring the central role of expectations and sentiment as key amplification channels for OPEC-related news shocks. Financial markets respond in a manner consistent with heightened inflationary concerns. The 10-year government bond yield rises

more strongly than in the baseline, while equity prices and the real effective exchange rate exhibit more muted reactions.

Finally, Figure 12 shows that US responses to non-OPEC-based news shocks are weaker, sometimes opposite, and generally less coherent with a stagflation narrative. This supports the view that such shocks are closer to demand shocks.

Figure 12: Reactions of US macroeconomic variables to a Non-OPEC news shock

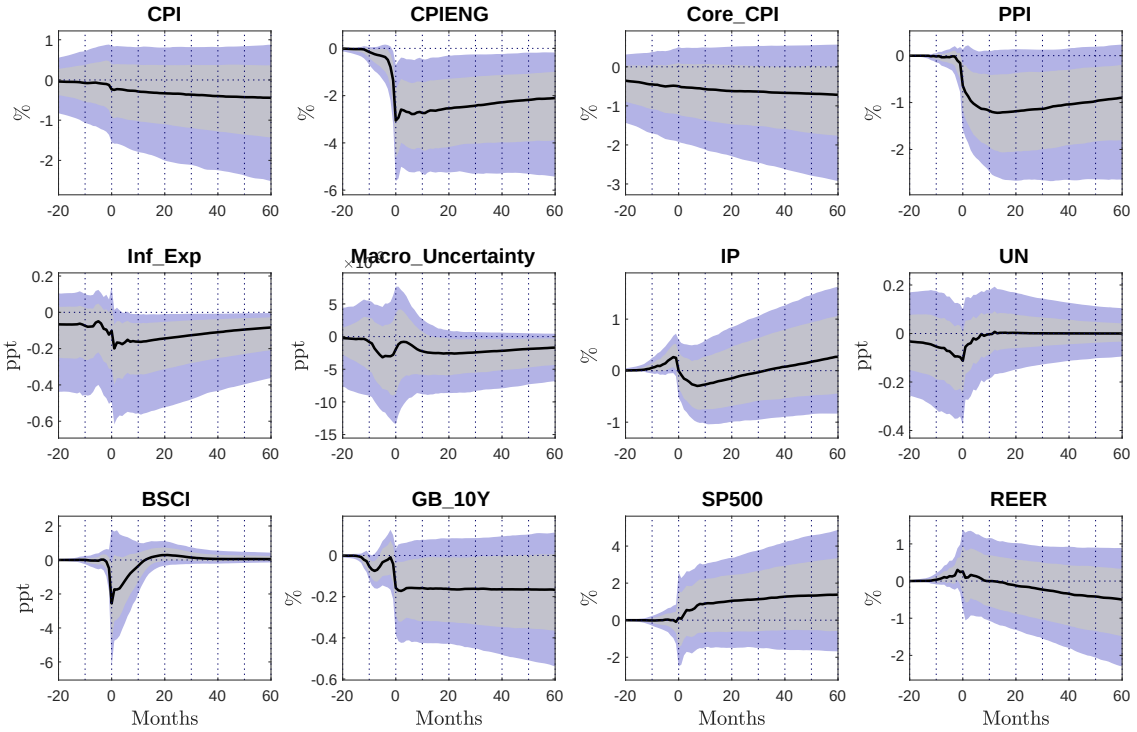

Notes: The solid lines are the posterior median cumulated impulse responses of US macroeconomics variables. Blue and grey shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA representation.

Consistent with the decline in the real oil price, price measures fall following the shock. Energy prices, in particular, contract sharply, suggesting that market participants do not interpret the disturbance as signaling a significant or persistent future supply constraint.The modest increase in macroeconomic uncertainty, alongside the slight declines in industrial production and limited rise in unemployment, further suggests that the macroeconomic impact of the shock is relatively contained. Confidence indices show only modest declines or even temporary improvements, while equity prices, long-term yields, and the real effective exchange rate exhibit limited movements, pointing to a milder financial transmission.

Overall, these results confirm that non-OPEC-based "news" shocks do not present the same macro-financial policy challenges as OPEC-driven news: they tend to be deflationary, less contractionary, and considerably less disruptive for the US economy.

7 Conclusion

This paper develops a unified framework to recover oil supply news shocks in an environment where expectations, not just fundamentals, drive market dynamics. By pairing a noncausal structural VAR with a Max-Share identification, we resolve the tension between non-fundamentalness and recoverability that constrains standard proxy SVAR approaches. The evidence reveals a global oil market that is forward looking and expectation driven: price and inventory movements anticipate production adjustments, while the real economy absorbs their delayed but powerful effects.

Beyond the oil market, our findings speak to a broader theme in macroeconomics, the centrality of beliefs in shaping real outcomes. Anticipation can be as powerful as realization, and when expectations are shared collectively, the future becomes a fundamental element of its own. Moving forward, extending this noncausal identification approach to other domains-such as fiscal foresight, climate policy, or geopolitical risk-offers a promising avenue to better understand how information, expectations, and time interact in driving global cycles.

References

Almutairi, H., Pierru, A., Smith, J.L., 2023. Oil market stabilization: The performance of OPEC and its allies. The Energy Journal 44, 1–22.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Vega, C., 2003. Micro effects of macro announcements: Real-time price discovery in foreign exchange. American Economic Review 93, 38–62.

Arezki, R., Ramey, V.A., Sheng, L., 2017. News Shocks in Open Economies: Evidence from Giant Oil Discoveries. The Quarterly Journal of Economics 132, 103–155.

Barsky, R., Kilian, L., 2002. Do we really know that oil caused the great stagflation? a monetary alternative, in: NBER Macroeconomics Annual 2001, Volume 16. National Bureau of Economic Research, Inc, pp. 137–198.

Barsky, R.B., Kilian, L., 2004. Oil and the macroeconomy since the 1970s. Journal of Economic Perspectives 18, 115–134.

Barsky, R.B., Sims, E.R., 2011. News shocks and business cycles. Journal of Monetary Economics 58, 273–289.

Baumeister, C., Guérin, P., 2021. A comparison of monthly global indicators for forecasting growth. International Journal of Forecasting 37, 1276–1295.

Baumeister, C., Hamilton, J.D., 2019. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. American Economic Review 109, 1873–1910.

Baumeister, C., Hamilton, J.D., 2024. Uncovering Disaggregated Oil Market Dynamics: A Full-Information Approach to Granular Instrumental Variables. Technical Report.

Baumeister, C., Kilian, L., 2016. Understanding the Decline in the Price of Oil since June 2014. Journal of the Association of Environmental and Resource Economists 3, 131–158.

Baumeister, C., Korobilis, D., Lee, T.K., 2022. Energy Markets and Global Economic Conditions. The Review of Economics and Statistics 104, 828–844.

Beaudry, P., Portier, F., 2006. Stock prices, news, and economic fluctuations. American Economic Review 96, 1293–1307.

Beaudry, P., Portier, F., 2014. News-driven business cycles: Insights and challenges. Journal of Economic Literature 52, 993–1074.

Bernanke, B.S., Kuttner, K.N., 2005. What explains the stock market's reaction to federal reserve policy? The Journal of Finance 60, 1221–1257.

Caldara, D., Cavallo, M., Iacoviello, M., 2019. Oil price elasticities and oil price fluctuations. Journal of Monetary Economics 103, 1–20.

Chahrour, R., Chugh, S.K., Potter, T., 2023. Anticipated Productivity and the Labor Market. Quantitative Economics 14, 897–934.

Chahrour, R., Jurado, K., 2021. Recoverability and expectations-driven fluctuations. The Review of Economic Studies 89, 214–239.

Davis, R.A., Song, L., 2020. Noncausal vector ar processes with application to economic time series. Journal of Econometrics 216, 246–267.

Degasperi, R., Hong, S.S., Ricco, G., 2025. Identification of Expectational Shocks in the Oil Market using OPEC Announcements. Working Paper. University of Warwick, Department of Economics.

Forni, M., Franconi, A., Gambetti, L., Sala, L., 2025. Asymmetric transmission of oil supply news. Quantitative Economics 16, 947–979.

Forni, M., Gambetti, L., Sala, L., 2014. No news in business cycles. The Economic Journal 124, 1168–1191.

Francis, N., Owyang, M.T., Roush, J.E., DiCecio, R., 2014. A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. The Review of Economics and Statistics 96, 638–647.

Gambetti, L., Moretti, L., 2017. News, Noise and Oil Price Swings. Research Technical Papers 12/RT/17. Central Bank of Ireland.

Gately, D., 1986. Lessons from the 1986 Oil Price Collapse. Brookings Papers on Economic Activity 17, 237–284.

Giancaterini, F., 2023. Essays on Univariate and Multivariate Noncausal Processes. Doctoral dissertation. Maastricht University.

Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. American Economic Review 102, 1692–1720.

Gisser, M., Goodwin, T.H., 1986. Crude oil and the macroeconomy: Tests of some popular notions: Note. Journal of Money, Credit and Banking 18, 95–103.

Gourieroux, C., Jasiak, J., 2017. Noncausal vector autoregressive process: Representation, identification and semi-parametric estimation. Journal of Econometrics 200, 118–134.

Gourieroux, C., Jasiak, J., 2023. Generalized covariance estimator. Journal of Business & Economic Statistics 41, 1315–1327.

Gourieroux, C., Jasiak, J., 2025. Nonlinear fore(back)casting and innovation filtering for causal-noncausal var models. URL: `https://arxiv.org/abs/2205.09922`, `arXiv:2205.09922`.

Gouriéroux, C., Monfort, A., Renne, J.P., 2020. Identification and Estimation in Non-Fundamental Structural VARMA Models. Review of Economic Studies 87, 1915–1953.

Gürkaynak, R.S., Sack, B., Swanson, E., 2005. Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. International Journal of Central Banking 1.

Hamilton, J.D., 1983. Oil and the macroeconomy since world war ii. Journal of Political Economy 91, 228–248.

Hamilton, J.D., 2009. Understanding crude oil prices. The Energy Journal 30, 179–206.

Jurado, K., Ludvigson, S.C., Ng, S., 2015. Measuring uncertainty. American Economic Review 105, 1177–1216.

Juvenal, L., Petrella, I., 2015. Speculation in the oil market. Journal of Applied Econometrics 30, 621–649.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. The American Economic Review 99, 1053–1069.

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

Kilian, L., Hicks, B., 2013. Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? Journal of Forecasting 32, 385–394.

Kilian, L., Lütkepohl, H., 2017. Structural Vector Autoregressive Analysis. Themes in Modern Econometrics, Cambridge University Press.

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

Kilian, L., Park, C., 2009. The impact of oil price shocks on the u.s. stock market. International Economic Review 50, 1267–1287.

Kilian, L., Zhou, X., 2023. The econometrics of oil market VAR models, in: Essays in Honor of Joon Y. Park: Econometric Methodology in Empirical Applications. Emerald Publishing Limited.

Kurmann, A., Sims, E., 2021. Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks. The Review of Economics and Statistics 103, 216–235.

Kuttner, K.N., 2001. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. Journal of Monetary Economics 47, 523–544.

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

Lanne, M., Luoto, J., 2013. Autoregression-based estimation of the new Keynesian Phillips curve. Journal of Economic Dynamics and Control 37, 561–570.

Lanne, M., Luoto, J., 2016. Noncausal bayesian vector autoregression. Journal of Applied Econometrics 31, 1392–1406.

Lanne, M., Saikkonen, P., 2011. Noncausal autogressions for economic time series. Journal of Time Series Econometrics 3.

Lanne, M., Saikkonen, P., 2013. Noncausal Vector Autoregression. Econometric Theory 29, 447–481.

Mumtaz, H., Pirzada, A., Theodoridis, K., 2018. Non-linear effects of oil shocks on stock prices. Working Papers 865. Queen Mary University of London, School of Economics and Finance.

Nelimarkka, J., 2017a. Evidence on News Shocks under Information Deficiency. MPRA Paper 80850. University Library of Munich, Germany.

Nelimarkka, J., 2017b. The effects of government spending under anticipation: the noncausal VAR approach. MPRA Paper 81303. University Library of Munich, Germany.

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

Pierru, A., Smith, JamesL., Hossa, A., 2020. OPEC's pursuit of market stability. Economics of Energy Environmental Policy 9.

Pierru, A., Smith, JamesL., Zamrik, T., 2018. OPEC's impact on oil price volatility: The role of spare capacity. The Energy Journal 39, 173–196.

Pinchetti, M., 2025. Geopolitical risk and inflation: the role of energy markets. Eco Notepad 420. Banque de France.

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

Ramey, V.A., 2011. Identifying Government Spending Shocks: It's all in the Timing*. The Quarterly Journal of Economics 126, 1–50.

Romer, C.D., Romer, D.H., 2004. A new measure of monetary shocks: Derivation and implications. American Economic Review 94, 1055–1084.

Rosenblatt, M., 2000. Gaussian and Non-Gaussian Linear Time Series and Random Fields. Springer. doi:10.1007/978-1-4612-1262-1.

Uhlig, H., 2004. Do technology shocks lead to a fall in total hours worked? Journal of the European Economic Association 2, 361–371.

Velasco, C., 2023. Identification and estimation of structural varma models using higher order dynamics. Journal of Business & Economic Statistics 41, 819–832.

A Equivalence with alternative NC-VAR specifications

The multiplicative VMAR(r,s) representation in (23), following Lanne and Saikkonen (2013), admits an alternative specification as a VAR(n_1,n_2,p) process in the sense of Davis and Song (2020) and Gourieroux and Jasiak (2017). The latter representation writes the n-dimensional process y_t as:

$$y_t = \Theta_1 y_{t-1} + \Theta_2 y_{t-2} + \cdots + \Theta_p y_{t-p} + u_t$$

where the characteristic polynomial $\det(I_n - \Theta_1 z - \cdots - \Theta_p z^p) = 0$ has n_1 roots outside and n_2 roots inside the unit circle, with $n_1 + n_2 = n \times p$. Following Giancaterini (2023), Chapter 4, the equivalence between these two specifications requires that the highest-order lead coefficient matrix in the noncausal polynomial has full rank. We now verify that this condition fails in our VMAR(1,2) specification. The noncausal polynomial $\Phi(L^{-1}) = I_2 - \Phi_1 L^{-1} - \Phi_2 L^{-2}$ has coefficient matrices:

$$\Phi_1 = \begin{bmatrix} 0 & 0 \\ \beta & 0 \end{bmatrix}, \quad \Phi_2 = \begin{bmatrix} 0 & 0 \\ \frac{\beta^2}{1-\rho\beta} & 0 \end{bmatrix}$$

Computing the determinant of the highest-order coefficient yields:

$$\det(\Phi_2) = \det \begin{bmatrix} 0 & 0 \\ \frac{\beta^2}{1-\rho\beta} & 0 \end{bmatrix} = 0 \times 0 - 0 \times \frac{\beta^2}{1-\rho\beta} = 0 \tag{38}$$

Hence $\text{rank}(\Phi_2) = 1 < 2 = n$, confirming that the full-rank condition is violated. This rank deficiency implies that the process $(q_t, p_t)'$ does *not* admit a VAR(n_1,n_2,p) representation with i.i.d. errors in the sense of Davis and Song (2020) and Gourieroux and Jasiak (2017). To see this explicitly, consider the companion form of the noncausal component. Defining $Z_t = (y_t', y_{t+1}')'$ and rewriting the noncausal part as $Z_t = \Upsilon Z_{t+1} + v_t$,

the companion matrix takes the form:

$$\Upsilon = \begin{bmatrix} \Phi_1 & \Phi_2 \\ I_2 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ \beta & 0 & \frac{\beta^2}{1-\rho\beta} & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \tag{39}$$

The matrix Υ is singular precisely because Φ_2 is rank-deficient. Consequently, the eigenvalue decomposition required for the Gourieroux and Jasiak (2017) representation theorem cannot be directly applied to obtain a standard VAR(n_1, n_2, p) with i.i.d. innovations. This rank deficiency does not compromise the validity of our approach. The VMAR(r, s) multiplicative structure of Lanne and Saikkonen (2013) remains well-defined and estimable, as both $\Pi(L)$ and $\Phi(L^{-1})$ satisfy the stability conditions $\det \Pi(z) \neq 0$ and $\det \Phi(z) \neq 0$ for $|z| \leq 1$.

B Noncausal representation and recoverability in the general news model

This appendix extends the stylized model of Section 2 to a general specification of the oil production process, following Nelimarkka (2017a). We derive the NC-VAR$(1, l)$ representation and prove recoverability of the oil supply news shock in the sense of Chahrour and Jurado (2021).

B.1 General oil production process with news

Consider the general oil production process:

$$q_t = \rho q_{t-1} + \chi \epsilon_t + \epsilon_{t-l} \tag{40}$$

where $|\rho| < 1$, $l \geq 1$ denotes the anticipation horizon, and $\chi \in [0, 1)$ captures the contemporaneous impact of the news shock on oil production. The case $\chi = 0$ corresponds to a pure news shock with no immediate effect on production. The oil price p_t (or any forward-looking variable, including inventories, except for capacity constraints) is determined by the forward-looking equilibrium condition:

$$p_t = \beta \mathbb{E}_t[p_{t+1}] + q_t + \nu_t \tag{41}$$

where $|\beta| < 1$ and ν_t is an exogenous price disturbance orthogonal to the news shock. The forward-looking solution of equation (41) is:

$$p_t = \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t[q_{t+j}] + \nu_t \tag{42}$$

From equation (40), iterating forward yields for $j \geq 1$:

$$q_{t+j} = \rho^j q_t + \sum_{i=0}^{j-1} \rho^{j-1-i} \left(\chi \epsilon_{t+1+i} + \epsilon_{t+1+i-l} \right) \tag{43}$$

Taking conditional expectations with $\mathbb{E}_t[\epsilon_{t+k}] = \epsilon_{t+k}$ for $k \leq 0$ and $\mathbb{E}_t[\epsilon_{t+k}] = 0$ for $k > 0$:

$$\mathbb{E}_t[q_{t+j}] = \rho^j q_t + \sum_{i=0}^{\min(j-1, l-1)} \rho^{j-1-i} \epsilon_{t+1+i-l}, \quad j \geq 1 \tag{44}$$

Substituting $k = l - 1 - i$ (so that $\epsilon_{t+1+i-l} = \epsilon_{t-k}$):

$$\mathbb{E}_t[q_{t+j}] = \rho^j q_t + \sum_{k=\max(l-j, 0)}^{l-1} \rho^{j-l+k} \epsilon_{t-k}, \quad j \geq 1 \tag{45}$$

Proposition 2. *Under $|\rho| < 1$, $|\beta| < 1$, and $\chi \in [0, 1)$, the equilibrium price is:*

$$p_t = \theta \rho q_{t-1} + \sum_{k=0}^{l} \gamma_k \epsilon_{t-k} + \nu_t \tag{46}$$

where $\theta = (1 - \rho\beta)^{-1}$ *and the coefficients are:*

$$\gamma_0 = \theta(\chi + \beta^l) \tag{47}$$

$$\gamma_k = \theta\beta^{l-k}, \quad 1 \leq k \leq l-1 \tag{48}$$

$$\gamma_l = \theta \tag{49}$$

Proof. Substituting (45) into (42) and collecting terms in q_t:

$$1 + \sum_{j=1}^{\infty} (\rho\beta)^j = \frac{1}{1 - \rho\beta} = \theta \tag{50}$$

so that $\theta q_t = \theta\rho q_{t-1} + \theta\chi\epsilon_t + \theta\epsilon_{t-l}$. For the anticipation terms, each ϵ_{t-k} with $k \in \{0, \ldots, l-1\}$ receives contributions from horizons $j \geq l - k$:

$$\sum_{j=l-k}^{\infty} \beta^j \rho^{j-l+k} = \frac{(\rho\beta)^{l-k}}{1 - \rho\beta} \cdot \rho^{k-l} = \theta\beta^{l-k} \tag{51}$$

Combining all contributions: ϵ_t receives $\theta\chi$ from fundamentals and $\theta\beta^l$ from anticipation, yielding $\gamma_0 = \theta(\chi + \beta^l)$; intermediate shocks ϵ_{t-k} for $1 \leq k \leq l-1$ receive only anticipation terms $\gamma_k = \theta\beta^{l-k}$; and ϵ_{t-l}, being already realized, receives only $\gamma_l = \theta$ from the fundamental component. □

Specializing to a two-period anticipation horizon ($l = 2$), we obtain $\gamma_0 = \theta(\chi + \beta^2)$, $\gamma_1 = \theta\beta$, and $\gamma_2 = \theta$. The equilibrium price becomes:

$$p_t = \theta\rho q_{t-1} + \theta(\chi + \beta^2)\epsilon_t + \theta\beta\epsilon_{t-1} + \theta\epsilon_{t-2} + \nu_t \tag{52}$$

We now examine the invertibility properties of this reduced-form representation. The structural moving average representation of $y_t = (q_t, p_t)'$ is:

$$y_t = \begin{bmatrix} \rho & 0 \\ \theta\rho & 0 \end{bmatrix} y_{t-1} + B(L) \begin{bmatrix} \epsilon_t \\ \nu_t \end{bmatrix} \tag{53}$$

where the MA polynomial $B(L)$ has the triangular form:

$$B(L) = \begin{bmatrix} \chi + L^l & 0 \\ b_{21}(L) & 1 \end{bmatrix} \tag{54}$$

with $b_{21}(L) = \gamma_0 + \sum_{k=1}^{l-1} \gamma_k L^k + \gamma_l L^l$.

Lemma 1. *The MA representation (53) is non-fundamental if and only if $\chi < 1$. The determinant polynomial is:*

$$\det B(z) = \chi + z^l \tag{55}$$

which has l roots:

$$z_k = \chi^{1/l} e^{i\pi(2k+1)/l}, \quad k = 0, 1, \ldots, l-1 \tag{56}$$

all with modulus $|z_k| = \chi^{1/l} < 1$ when $\chi < 1$.

Proof. From the triangular structure of $B(z)$:

$$\det B(z) = (\chi + z^l) \cdot 1 = \chi + z^l \tag{57}$$

Setting $\det B(z) = 0$ yields $z^l = -\chi = \chi e^{i\pi}$. The l complex roots are:

$$z_k = \chi^{1/l} e^{i\pi(2k+1)/l}, \quad k = 0, 1, \ldots, l-1 \tag{58}$$

Each root has modulus $|z_k| = \chi^{1/l}$. For $\chi < 1$, we have $|z_k| < 1$, so all l roots lie strictly inside the unit circle, rendering the MA polynomial non-invertible. □

The l roots are uniformly distributed on a circle of radius $\chi^{1/l}$ centered at the origin, with angular separation $2\pi/l$. For $l = 2$: $z_{0,1} = \pm i\sqrt{\chi}$. For $l = 3$: three roots at angles $\pi/3$, π, and $5\pi/3$.

To address this non-fundamentalness, we reformulate the model in its noncausal VAR representation. From equation (40):

$$(1 - \rho L) q_t = (\chi + L^l) \epsilon_t \tag{59}$$

Rewriting the right-hand side using lead operators:

$$(\chi + L^l)\epsilon_t = L^l(\chi L^{-l} + 1)\epsilon_t = (1 + \chi L^{-l})\epsilon_{t-l} \tag{60}$$

Lemma 2. *For $\chi < 1$, the polynomial $1 + \chi z^l$ has all roots outside the unit circle, hence $(1 + \chi L^{-l})$ is invertible on the unit circle.*

Proof. The roots of $1 + \chi z^l = 0$ are $z_k = \chi^{-1/l} e^{i\pi(2k+1)/l}$ with modulus $|z_k| = \chi^{-1/l} > 1$ for $\chi < 1$. □

Therefore:

$$(1 - \rho L)(1 + \chi L^{-l})^{-1} q_t = \epsilon_{t-l} \tag{61}$$

Expanding $(1 + \chi L^{-l})^{-1} = \sum_{j=0}^{\infty} (-\chi)^j L^{-lj}$, the price equation can be rewritten in terms of future values of q_t.

Proposition 3. *Under $|\rho| < 1$, $|\beta| < 1$, and $0 \leq \chi < 1$, the system (40)–(41) admits a noncausal VAR$(1, l)$ representation:*

$$\Pi(L)\Phi(L^{-1})y_t = B_0 u_t \tag{62}$$

where $y_t = (q_t, p_t)'$, $u_t = (\epsilon_{t-l}, \nu_t)'$, and

$$\Pi(L) = I_2 - \Pi_1 L, \quad \Pi_1 = \begin{bmatrix} \rho & 0 \\ \theta\rho & 0 \end{bmatrix} \tag{63}$$

$$\Phi(L^{-1}) = I_2 - \sum_{j=1}^{s} \Phi_j L^{-j}, \quad \Phi_j = \begin{bmatrix} \phi_{j,11} & 0 \\ \phi_{j,21} & 0 \end{bmatrix} \tag{64}$$

$$B_0 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \tag{65}$$

The lead polynomial coefficients depend on χ as follows:

(*i*) *If $\chi = 0$ (pure news shock), then $\phi_{j,11} = 0$ for all $j \geq 1$, and the truncation order is*

$s = l$.

(*ii*) *If* $\chi \in (0,1)$ *(mixed shock), the expansion* $(1+\chi L^{-l})^{-1} = \sum_{k=0}^{\infty}(-\chi)^k L^{-kl}$ *implies that non-zero entries in* $\phi_{j,11}$ *occur only at multiples of* l*:*

$$\phi_{kl,11} = -(-\chi)^k, \quad k = 1, 2, 3, \ldots \tag{66}$$

In practice, truncating at order $s \gg l$ *provides an accurate approximation.*

(*iii*) *For* $l = 2$ *with* $\chi \in (0,1)$*, the coefficients satisfy:*

$$\phi_{j,11} = \begin{cases} 0 & j \text{ odd} \\ -(-\chi)^{j/2} & j \text{ even} \end{cases} \tag{67}$$

$$\phi_{j,21} = \begin{cases} \beta(-\chi)^{(j-1)/2} & j \text{ odd} \\ \theta(\chi+\beta^2)(-\chi)^{j/2-1} & j \text{ even} \end{cases} \tag{68}$$

The impulse response matrices Ψ_j *solve the recursion:*

$$\Psi_j = \mathbf{1}_{j\geq 0}\Pi_1^j + \sum_{k=1}^{s} \Phi_k \Psi_{j+k} \tag{69}$$

with boundary condition $\Psi_j = 0$ *for* $j < -s$.

We now assess whether the news shock can be recovered from observables using the recoverability criterion of Chahrour and Jurado (2021).

B.2 Recoverability in the sense of Chahrour and Jurado (2021)

We now establish the main theoretical result: the oil supply news shock is recoverable from the observables.

Definition 1. *From Chahrour and Jurado (2021): let* $\{y_t\}$ *be obtained from structural*

shocks $\{u_t\}$ *by a linear transformation with spectral characteristic* $\varphi(\lambda)$*:*

$$y_t = \int_{-\pi}^{\pi} e^{i\lambda t}\varphi(\lambda)\Phi_u(d\lambda) \tag{70}$$

The shocks are ***recoverable*** *from* $\{y_t\}$ *if and only if:*

$$\text{rank}[\varphi(\lambda)] = n_u \quad \textit{for almost all } \lambda \in [-\pi, \pi] \tag{71}$$

Proposition 4. *Consider the NC-VAR*$(1, l)$ *representation* (62) *under:*

(A1) $|\rho| < 1$ *(stationarity of oil production)*

(A2) $|\beta| < 1$ *(convergence of forward-looking solution)*

(A3) $0 \leq \chi < 1$ *(news shock dominates contemporaneous effect)*

(A4) $l \geq 1$ *(anticipation horizon)*

Then the structural shocks $(\epsilon_t, \nu_t)'$ *are recoverable from the observables* $(q_t, p_t)'$ *in the sense of Chahrour and Jurado (2021).*

Proof. The NC-VAR$(1, l)$ representation (62) implies the two-sided MA form

$$y_t = \Phi(L^{-1})^{-1}\Pi(L)^{-1}B_0 u_t$$

, with spectral characteristic $\varphi(\lambda) = \Phi(e^{i\lambda})^{-1}\Pi(e^{-i\lambda})^{-1}B_0$. We establish that $\det[\varphi(\lambda)] \neq 0$ for all $\lambda \in [-\pi, \pi]$ by examining each factor. The causal polynomial $\Pi(z) = I_2 - \Pi_1 z$ has determinant $\det[\Pi(z)] = 1 - \rho z$, with unique root $z = 1/\rho$. Under Assumption (A1), $|\rho| < 1$ implies $|1/\rho| > 1$, hence $\det[\Pi(e^{-i\lambda})] = 1 - \rho e^{-i\lambda} \neq 0$ for all λ. The noncausal polynomial $\Phi(z) = I_2 - \sum_{j=1}^{s} \Phi_j z^j$ inherits its lower-triangular structure from the Φ_j matrices, so $\det[\Phi(z)] = (1 - \sum_j \phi_{j,11} z^j) \cdot 1$. When $\chi = 0$, all $\phi_{j,11} = 0$ and $\det[\Phi(z)] = 1$. When $\chi \in (0, 1)$, the convergence of the geometric series in χ ensures $\det[\Phi(e^{i\lambda})] \neq 0$ for

all λ. Since $\det[B_0] = 1$, we conclude that

$$\det[\varphi(\lambda)] = \frac{\det[B_0]}{\det[\Phi(e^{i\lambda})]\det[\Pi(e^{-i\lambda})]} \neq 0 \quad \forall \lambda \in [-\pi, \pi] \tag{72}$$

Thus $\text{rank}[\varphi(\lambda)] = 2 = n_u$ for all λ, which satisfies the Chahrour and Jurado (2021) recoverability condition. □

Corollary 1. *Under the conditions of Proposition 4, the structural shocks can be explicitly recovered as:*

$$u_t = B_0^{-1}\Pi(L)\Phi(L^{-1})y_t \tag{73}$$

which is a linear combination of a finite number of leads and lags of the observables:

$$\begin{bmatrix} \epsilon_{t-l} \\ \nu_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \left(y_t - \Pi_1 y_{t-1} - \sum_{j=1}^{l} \Phi_j y_{t+j} \right) \tag{74}$$

C Data

Table 1: Variables used, sources and transformations

Variables	Code	Period	Data source	Transformation
Global Oil Market Variables				
World oil production	WOP	1974:01-2025:07	US EIA's Monthly Energy Review	First log difference
World industrial production	WIP	1974:01-2025:07	Christiane Baumeister website	First log difference
Refiner acquisition cost	RAC	1974:01-2025:07	US EIA's Monthly Energy Review	Level
World oil stocks[a]	Stocks	1974:01-2025:07	US EIA's Monthly Energy Review	First difference
Global economic variable				
Global Economic Conditions indicators	GECON	1974:01-2025:07	Christiane Baumeister website	Level
Geopolitical Risk Index	GPR	1974:01-2025:07	Caldara et al. (2019)	Level
OECD Inflation	Infl_OECD	1974:01-2025:02	OECD database	Level
OECD passenger car registrations	PCR_OECD	1974:01-2019:1	OECD database	Level
World stock price index	MSCI	1974:01-2025:07	Datastream	First log difference
US macroeconomic variables				
Consumer price index for all urban consumers and all items	CPI	1974:01-2025:07	FRED Database	First log difference
Consumer price index for all urban consumers: Energy in U.S. City Average	CPIENG	1974:01-2025:07	FRED Database	First log difference
Consumer price index excluding food and energy	Core CPI	1974:01-2025:07	FRED Database	First log difference
Producer Price Index	PPIACO	1974:01-2025:07	FRED Database	First log difference
US Industrial production	INDPROD	1974:01-2025:07	FRED Database	First log difference
US Unemployment rate	UNRATE	1974:01-2025:07	FRED Database	Level
University of Michigan: Inflation Expectation	INF_EXP	1978:01-2025:07	Michigan Survey	Level
Business Tendency Surveys for Manufacturing	BSCI	1974:01-2025:07	FRED Database	Level
1-Year Ahead Macroeconomic Uncertainty[b]	Macro_uncertainty	1974:01-2025:07	FRED Database	Level
Long-term interest rate	10Y GB	1974:01-2025:07	FRED Database	Level
Stock price index	SP500	1974:01-2025:07	Bloomberg	First log difference
Real effective exchange rate	REER	1974:01-2025:07	FRED Database	First log difference

[a] Following Kilian and Murphy (2014)'s method, the world oil stock variable is constructed by rescaling U.S. crude oil inventories with the OECD-to-U.S. petroleum stocks ratio. All variables are from the EIA database and the series is seasonally adjusted.

[b] Macroeconomic uncertainty is proxied by 12-month ahead macro uncertainty measure of Jurado et al. (2015), drawn from large number of U.S. macroeconomic time series.

D Bayesian estimation of the noncausal VAR

This appendix provides a detailed description of the Bayesian estimation procedure for the NC-VAR(r,s) model, following Lanne and Luoto (2016) and Nelimarkka (2017a). We derive the Gibbs sampler algorithm that exploits the conditional normality of the likelihood function under the multivariate t-distribution assumption.

D.1 Model setup and notation

Consider the NC-VAR(r,s) model from equation (24):

$$\Pi(L)\Phi(L^{-1})y_t = \epsilon_t$$

where $\Pi(L) = I_n - \Pi_1 L - \cdots - \Pi_r L^r$ is the causal polynomial, $\Phi(L^{-1}) = I_n - \Phi_1 L^{-1} - \cdots - \Phi_s L^{-s}$ is the noncausal polynomial, and ϵ_t follows a multivariate t-distribution as specified in equation (27):

$$\epsilon_t = \omega_t^{-1/2}\eta_t, \quad \eta_t \sim \mathcal{N}(0, \Sigma), \quad \lambda\omega_t \sim \chi^2_\lambda$$

where $\omega_t^{-1/2}$ is a scalar volatility factor and λ denotes the degrees-of-freedom parameter.

Let Π and Φ be matrices stacking Π_i' for $i = 1, \ldots, r$ and Φ_i' for $i = 1, \ldots, s$, respectively. Define $\pi = \text{vec}(\Pi)$, $\phi = \text{vec}(\Phi)$, and denote the full parameter vector as $\theta = (\pi', \phi', \text{vech}(\Sigma)', \lambda)'$. For convenience, we also define $\vartheta = (\pi', \phi')'$. To impose s^* zero restrictions on matrix Φ, we introduce an $((n^2 s - s^*) \times 1)$ vector ϕ_r containing the unrestricted parameters of Φ and an $(n^2 s \times (n^2 s - s^*))$ deterministic selection matrix R_ϕ which maps the unrestricted parameters to the full vector as $\phi = R_\phi \phi_r$. When no restrictions are imposed, $R_\phi = I_{n^2 s}$ and $\phi_r = \phi$.

D.2 Conditional likelihood

Define the auxiliary variables:

$$v_t(\phi) = y_t - \Phi_1 y_{t+1} - \cdots - \Phi_s y_{t+s}$$

and the transformed error term:

$$\epsilon_t(\vartheta) = v_t(\phi) - \sum_{j=1}^{r} \Pi_j v_{t-j}(\phi)$$

The approximate conditional joint density of $y = (y_1, \ldots, y_T)$ given $\omega = (\omega_{r+1}, \ldots, \omega_{T-s})$ is:

$$p(y|\omega, \theta) \approx \prod_{t=r+1}^{T-s} p(\epsilon_t(\vartheta)|\omega_t, \Sigma)$$

where

$$p(\epsilon_t|\omega_t, \Sigma) = \frac{\omega_t^{n/2}}{(2\pi)^{n/2}|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}\omega_t \epsilon_t(\vartheta)'\Sigma^{-1}\epsilon_t(\vartheta)\right)$$

D.3 Prior distributions

We specify the following prior distributions:

$$\begin{aligned}
\pi &\sim \mathcal{N}(\underline{\pi}, V_\pi) \cdot \mathbf{1}(\pi) \\
\phi_r &\sim \mathcal{N}(\underline{\phi}_r, V_{\phi_r}) \cdot \mathbf{1}(\phi), \quad \phi = R_\phi \phi_r \\
\Sigma &\sim \mathcal{IW}(\underline{S}, \underline{\nu}) \\
\lambda &\sim \text{Exp}(\underline{\lambda})
\end{aligned}$$

where $\mathbf{1}(\cdot)$ is an indicator function equal to 1 when the polynomial to which π or ϕ is mapped has all roots outside the unit circle, and $\mathcal{IW}$ denotes the inverse Wishart distribution. We impose a Minnesota-Litterman type prior on the dynamic coefficients. The prior means are set to zero: $\underline{\pi} = 0$ and $\underline{\phi}_r = 0$. The prior covariance matrices V_π and V_{ϕ_r} are diagonal with elements:

$$\begin{aligned}
\sigma_{\pi,iil} &= \frac{\gamma_{1,\pi}}{l^{\gamma_3}}, \quad \sigma_{\pi,ijl} = \gamma_2 \frac{\gamma_{1,\pi}}{l^{\gamma_3}} \frac{\hat{\sigma}_i}{\hat{\sigma}_j}, \quad i, j = 1, \ldots, n, \quad l = 1, \ldots, r \\
\sigma_{\phi_r,iil} &= \frac{\gamma_{1,\phi}}{l^{\gamma_3}}, \quad \sigma_{\phi_r,ijl} = \gamma_2 \frac{\gamma_{1,\phi}}{l^{\gamma_3}} \frac{\hat{\sigma}_i}{\hat{\sigma}_j}, \quad i, j = 1, \ldots, n, \quad l = 1, \ldots, s
\end{aligned}$$

where $\hat{\sigma}_i$ is the residual standard error from a univariate autoregression with r lags on the ith variable, $\gamma_{1,\pi}$ and $\gamma_{1,\phi}$ control for overall tightness, γ_2 controls for relative tightness across equations, and γ_3 is a decay parameter for more distant lags and leads.

In our baseline estimation, we set $\gamma_{1,\pi} = 0.2$, $\gamma_{1,\phi} = 0.15$, $\gamma_{2,\pi} = 0.5$, $\gamma_{2,\phi} = 0.3$, $\gamma_{3,\pi} = 1.3$ and $\gamma_{3,\phi} = 1$. The tighter prior on lead coefficients ($\gamma_{1,\phi} < \gamma_{1,\pi}$) shrinks them

more heavily towards zero, reflecting the prior belief that lag terms are more important in determining the dynamics. This asymmetric shrinkage is crucial for achieving unimodality of the posterior distribution. For the remaining hyperparameters, we set $\underline{S} = (\underline{\nu} - n - 1)\text{diag}(\hat{\sigma}_1^2, \ldots, \hat{\sigma}_n^2)$ with $\underline{\nu} = n + 2$, and $\underline{\lambda} = 8$.

D.4 Gibbs sampler algorithm

The Gibbs sampler iterates through the following steps. Define the matrices Y and U by stacking $v_t^* = \omega_t^{1/2} v_t(\phi)'$ and $U_t^* = \omega_t^{1/2}[v_{t-1}(\phi)' \cdots v_{t-r}(\phi)']'$, respectively, for $t = r+1, \ldots, T-s$. Similarly, define Y^* and X^* by stacking $y_t^* = \omega_t^{1/2}\Pi(L)y_t$ and $X_t^* = \omega_t^{1/2}\Pi(L)X_t$, where $X_t = I_n \otimes [y_{t+1}' \cdots y_{t+s}']'$.

Step 1: Draw ϕ_r. The full conditional posterior distribution of ϕ_r is:

$$\phi_r | y, \pi, \Sigma, \omega \sim \mathcal{N}(\bar{\phi}_r, \bar{V}_{\phi_r}) \cdot \mathbf{1}(\phi), \quad \phi = R_\phi \phi_r$$

where

$$\begin{aligned} \bar{V}_{\phi_r}^{-1} &= V_{\phi_r}^{-1} + R_\phi' X^{*\prime} \Omega X^* R_\phi \\ \bar{\phi}_r &= \bar{V}_{\phi_r}\left(V_{\phi_r}^{-1}\underline{\phi}_r + R_\phi' X^{*\prime}\Omega Y^*\right) \end{aligned}$$

and $\Omega = I_{T-r-s} \otimes \Sigma^{-1}$.

Step 2: Draw π. The conditional distribution of π is:

$$\pi | y, \phi, \Sigma, \omega \sim \mathcal{N}(\bar{\pi}, \bar{V}_\pi) \cdot \mathbf{1}(\pi)$$

where

$$\begin{aligned} \bar{V}_\pi^{-1} &= V_\pi^{-1} + \Sigma^{-1} \otimes U'U \\ \bar{\pi} &= \bar{V}_\pi\left(V_\pi^{-1}\underline{\pi} + \text{vec}(U'Y\Sigma^{-1})\right) \end{aligned}$$

Step 3: Draw Σ. Defining $\bar{S} = \underline{S} + E'E$ where $E = Y - U\Pi$, and $\bar{\nu} = \underline{\nu} + T - s - r$, the conditional posterior distribution for Σ is:

$$\Sigma | y, \pi, \phi, \omega \sim \mathcal{IW}(\bar{S}, \bar{\nu})$$

Step 4: Draw ω. The volatility parameters are drawn from:

$$\left(\lambda + \epsilon_t(\vartheta)' \Sigma^{-1} \epsilon_t(\vartheta)\right) \omega_t | y, \pi, \phi, \Sigma, \lambda \sim \chi^2_{\lambda+n}, \quad t = r+1, \ldots, T-s$$

Step 5: Draw λ. The degrees-of-freedom parameter is drawn using a Metropolis-within-Gibbs step from the kernel:

$$p(\lambda | y, \omega) \propto \left(2^{\lambda/2}\Gamma(\lambda/2)\right)^{-(T-r-s)} \lambda^{\lambda(T-r-s)/2} \left(\prod_{t=r+1}^{T-s} \omega_t^{(\lambda-2)/2}\right) \exp\left[-\left(\frac{1}{\underline{\lambda}} + \frac{1}{2}\sum_{t=r+1}^{T-s} \omega_t\right)\lambda\right]$$

We use a univariate normal distribution with mean equal to the mode and variance equal to the inverse of the second derivative of the log-kernel as the proposal distribution. The standard Metropolis-Hastings acceptance probability is then computed.

D.5 Posterior inference

In our empirical application, results are based on 10,000 posterior draws obtained after a burn-in period of 50,000 draws. For each retained draw $(\pi^{(m)}, \phi^{(m)}, \Sigma^{(m)}, \lambda^{(m)})$, $m = 1, \ldots, 10000$, we compute the two-sided MA coefficients $\{\Psi_h^{(m)}\}_{h=-s}^{H}$ from equation (25), perform the Max-Share identification to obtain $W^{(m)}$, and compute the structural impulse responses $\Theta_h^{(m)} = \Psi_h^{(m)} \tilde{A}^{(m)} W^{(m)}$. The reported point estimates are posterior medians, and the credible sets are constructed from the relevant quantiles across draws. The low posterior mean of the degrees-of-freedom parameter λ (typically below 5, as shown in Figure 1) provides strong evidence for fat tails in the error distribution, confirming the appropriateness of the multivariate t-distribution assumption and the identification of a unique NC-VAR specification.

E Robustness Check

E.1 Moving restrictions on real economic activity

Figure 13: Impulse response functions to oil news shocks without restrictions on WIP

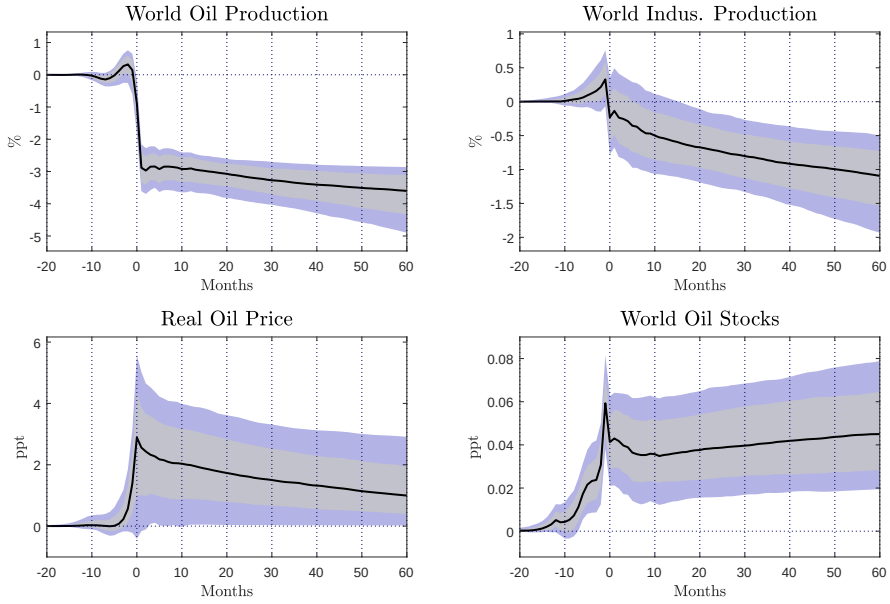

The figure displays posterior median impulse responses and associated 68 % and 90 % credible sets. Solid lines and shaded areas correspond to the earlier period, while dashed and dotted lines represent the later period. Owing to noncausality, impulse responses are defined on both sides of zero, with negative horizons capturing the lead terms of the NC-VAR moving-average representation.

E.2 Time backtest

Figure 14: Impulse response functions to oil news shocks across different sub-samples.

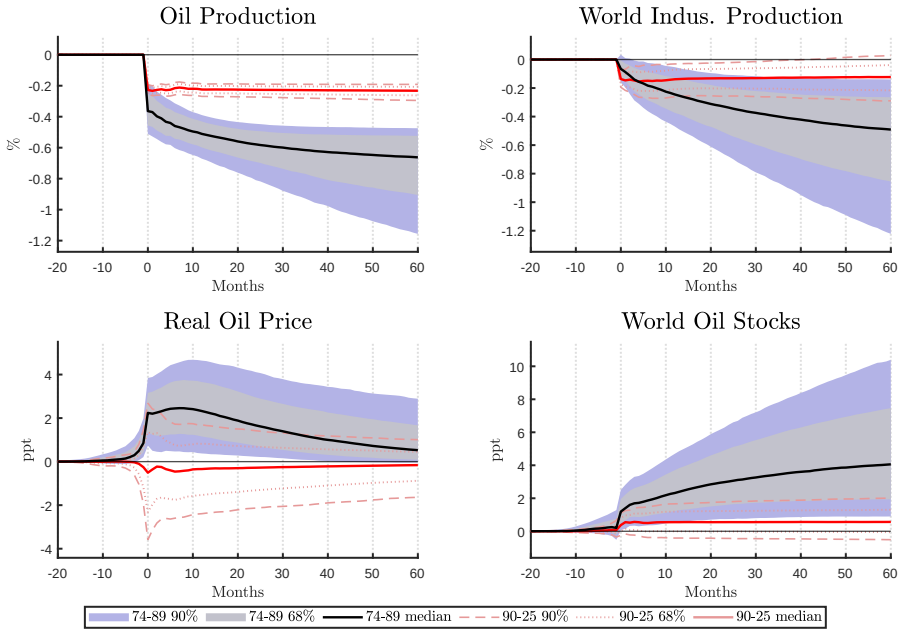

The figure displays posterior median impulse responses and associated 68 % and 90 % credible sets for the 1974–1989 and 1990–2025 samples. Black solid lines and shaded areas correspond to the earlier period, while red dashed and dotted lines represent the later period. Owing to noncausality, impulse responses are defined on both sides of zero, with negative horizons capturing the lead terms of the NC-VAR moving-average representation.