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

Stéphane Auray† Zakaria Moussa‡ Arthur Thomas§

June 18, 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 identify expectation-driven oil supply disturbances without external instruments. Applied to global oil market data, the model supports a non-fundamental representation and uncovers anticipatory responses of oil prices and inventories consistent with rational-expectations storage behavior. The identified shocks explain a substantial fraction of oil price fluctuations, particularly during the late 1970s–1980s and the 2014–15 oil price collapse, while the COVID-19 episode appears predominantly demand-driven. Finally, we show that oil supply news shocks are primarily OPEC-driven and generate stagflationary effects on both global and U.S. economic activity.

JEL classification: C32, D84, E32, Q41, Q43
Keywords: Oil supply news shocks, Global oil market, Max-Share, Non-fundamentalness, Structural Noncausal 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

A large body of research in macroeconomics and finance highlights the importance of expectations in shaping economic outcomes and asset prices. Information about future economic conditions can affect investment, consumption, production, inventory accumulation, and financial markets well before the underlying events materialize (Beaudry and Portier, 2006; Barsky and Sims, 2011; Forni et al., 2014; Forni and Gambetti, 2016; Kurmann and Sims, 2021; D'Amico and King, 2023; Chahrour et al., 2024, among others). In oil markets, one important source of anticipation arises from news about future oil supply conditions. Oil supply news shocks can be defined as exogenous revisions in the information available to economic agents regarding the future path of global oil production.

Expectations about future oil supply are influenced by a broad range of factors, including geopolitical developments, OPEC production decisions, investment in extraction capacity, technological innovations, and discoveries of new reserves. Information about these factors is often revealed gradually, prompting market participants to update their beliefs about future market conditions before any actual change in production occurs. As a result, expectations-driven shocks may generate immediate adjustments in oil prices, inventories, production plans, and broader macroeconomic activity, even though the underlying supply changes have not yet materialized.

Despite their potentially important role in oil market dynamics, empirical evidence on oil supply news shocks remains relatively scarce. A key reason is that anticipation effects generate non-fundamental representations of the data. Because economic agents react to information about future supply conditions before these conditions become observable, current variables may fail to span the relevant information set. As a result, structural news shocks are generally not recoverable from standard VAR innovations, creating a fundamental identification challenge. One notable exception is Arezki et al. (2017), who exploit the timing of giant oil discoveries as a quasi-natural experiment to show that anticipated future increases in oil supply can generate significant macroeconomic effects well in advance of actual production.

More recently, Känzig (2021) proposed a high-frequency identification strategy combining OPEC announcement windows with a 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 question of whether high-frequency surprises can be treated as clean measures of low-frequency structural shocks, given the aggregation and measurement issues emphasized by Kilian (2024). Moreover, Kilian and Zhou (2023) and Degasperi et al. (2025) argue that these OPEC-related surprises are more naturally interpreted as shifts in broader oil price expectations than as pure oil supply news.

More fundamentally, the identification of news shocks faces a recoverability problem. Agents react to information about future fundamentals before these fundamentals materialize. As a result, the econometrician's information set may be insufficient to recover the underlying structural innovations, generating a non-fundamental representation of the data. In the context of oil markets, Plagborg-Møller and Wolf (2022) show that Känzig's proxy-SVAR suffers from precisely this problem. Recent contributions by Chahrour and Jurado (2021), Plagborg-Møller and Wolf (2022), and Forni et al. (2025) have proposed alternative frameworks that either test recoverability, derive partial identification results, or explicitly incorporate forward-looking dynamics. However, the oil market literature still lacks a unified empirical framework capable of jointly addressing the non-fundamental nature of news shocks and the recoverability problem while delivering a structural identification of oil supply news shocks.

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) provides a coherent solution to the recoverability problem. The noncausal specification explicitly allows variables to depend on both past and future shocks. Rather than assuming that agents observe future shocks perfectly, it captures the idea that economic decisions may respond to information about future disturbances before those disturbances fully materialize. This feature makes NC-VAR models particularly well suited for studying news shocks and other forms of anticipatory behavior that generate non-fundamental representations. 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). 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 (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. However, their framework does not recover the causal–noncausal decomposition itself, which requires non-Gaussian information.[1]

Exploiting non-Gaussianity makes it possible to break the observational equivalence between fundamental and non-fundamental representations, a result that is now well

[1]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).

established in the literature. Following Lanne and Saikkonen (2013), we use this property to recover the underlying structural shocks and the associated two-sided impulse response functions from the NC-VAR representation. This step allows anticipated shocks to be inferred directly from the data without relying on external instruments.

Recoverability, however, does not by itself determine the economic interpretation of the recovered shocks. Once the structural innovations have been identified, an additional identification strategy is required to isolate the particular shock corresponding to revisions in expectations about future oil supply. In other words, recovering the structural shocks and identifying the oil supply news shock are conceptually distinct problems. The first concerns the extraction of structural innovations from a non-fundamental system, whereas the second concerns the selection of the economically relevant shock among the set of recovered structural disturbances.

Our second contribution is to provide an economically meaningful identification of the recovered structural shocks. To this end, we adapt the Max-Share identification strategy of Chahrour et al. (2023) to a noncausal VAR framework estimated on standard oil market variables, global oil production, global economic activity, the real price of oil, and global oil inventories. Within this framework, the oil supply news shock is defined as the structural innovation that maximizes the finite-horizon forecast error variance share of global oil production. The Max-Share criterion therefore provides an economically grounded way to isolate the expectation-driven component of oil supply fluctuations from the set of recovered structural shocks.

This combination of a noncausal VAR and Max-Share identification yields a unified framework that jointly addresses non-fundamentality, recoverability, and economic identification. Unlike proxy-SVAR approaches, it does not rely on external instruments and remains valid in the presence of non-fundamentalness, since the structural shocks are recovered directly from the estimated two-sided representation. In this sense, the framework allows oil supply news shocks, typically non-identifiable in standard VARs, to be recovered and identified using only information contained in the data.

Having established a framework for recovering and identifying news shocks, we next investigate their economic origins. Our third contribution is to trace the origins of the identified oil supply news shocks by exploiting institutional differences between OPEC and non-OPEC producers. OPEC's coordinated production management contrasts sharply with the market-based behavior of non-OPEC producers, implying different supply elasticities, adjustment speeds, and information structures (Pierru et al., 2018, 2020; Baumeister and Hamilton, 2024). This decomposition allows us to identify the sources of aggregate oil supply news and to assess whether anticipation effects originate primarily from coordinated production decisions or from broader market forces.

Our empirical results yield three key findings. first, the data provide strong support for a noncausal, fat-tailed specification, lending support to a non-fundamental representation of the global oil market. This finding implies that oil supply news shocks can be recovered directly from the data through the estimated two-sided representation, without relying on external instruments or proxy variables.

Second, the identified oil supply news shocks display clear anticipatory features: real oil prices and inventories react before the associated adjustments in physical market fundamentals, consistent with forward-looking behavior and storage dynamics. Historical decompositions show that these shocks explain a significant and time-varying share of both oil price and inventory fluctuations, particularly during the late 1970s–1980s and around the 2014 oil price collapse, whereas the COVID-19 episode appears largely demand-driven. Inventory dynamics provide additional evidence in favor of the news-shock interpretation through precautionary stock accumulation ahead of major supply-related events.

Third, our decomposition of global oil supply further reveals important differences between OPEC and non-OPEC producers. For OPEC production, the Max-Share criterion isolates shocks that are naturally interpreted as expectation-driven disturbances related to cartel behavior and future production management. In contrast, when applied to non-OPEC production, the procedure tends to capture a mixture of supply-capacity adjustments and demand-driven fluctuations. These findings suggest that aggregate oil supply news shocks are primarily an OPEC-driven phenomenon, while non-OPEC production contains substantially more information about global business-cycle conditions than about exogenous supply news. More broadly, our findings suggest that accounting for non-fundamentalness is crucial for understanding how expectations shape oil market dynamics and that standard fundamental VAR representations may substantially understate the role of anticipated supply disturbances.

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 news shocks and demonstrates how the NC-VAR framework resolves the recoverability problem. Section 3 details the NC-VAR specification, impulse response analysis, historical decomposition, and the Bayesian estimation procedure. Section 4 describes the data and presents the Max-Share identification strategy. Section 5 reports the baseline results for the global oil market. Section 6 examines the macroeconomic effects of oil supply news shocks at both the global and U.S. levels. Section 7 investigates the origins of oil supply news shocks by contrasting OPEC-driven and non-OPEC production dynamics. Section 8 concludes.

2 Stylised model of the oil market with news shocks

In this section, we develop a stylised rational expectations model of the global oil market in which movements in the real price of oil are driven by expectations about future oil supply.

The model provides a transparent link between anticipated oil supply disturbances and the joint dynamics of oil production and prices, and it clarifies how non-fundamentalness arises from the lagged effects of news shocks on observables. By embedding this structure into a noncausal VAR, we show that oil supply news shocks are recoverable in the sense of Chahrour and Jurado (2021), despite the failure of standard causal VARs to span the underlying structural innovations.

We begin by specifying the process for global oil production, denoted by q_t, as an autoregressive process with both contemporaneous and anticipated components:

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

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.[2] The parameter $\chi \geq 0$ governs the contemporaneous response of production to the shock. When $\chi < 1$, the anticipated component ϵ_{t-l} receives a relatively larger weight than the contemporaneous innovation, implying that expectations about future supply conditions play a central role in driving current outcomes. Because market participants observe information about future supply disturbances before these disturbances are fully reflected in production, their information set is richer than that available to the econometrician. Consequently, current and past realizations of production do not fully reveal the underlying structural shock. The reduced-form representation is therefore non-fundamental.

The real price of oil, denoted byp_t , is determined by a standard forward-looking equilibrium condition in which the price equals the present discounted value of expected future supply, up to an exogenous disturbance:

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

where $\beta < 1$ is a discount factor, $\mathbb{E}_t[\cdot]$ is the conditional expectation given the information set $\{q_t, p_t, \epsilon_t, \nu_t\}$, and ν_t is an exogenous disturbance orthogonal to the news shock. Solving equation (2) forward in the usual way delivers the present value solution:

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

so that the price at time t reflects the entire discounted path of expected future production.

To make the role of anticipation explicit, we specialize to the case $l = 2$, leaving the general case for Online Appendix B). Using the production process (1), we obtain closed-form expressions for $\mathbb{E}_t(q_{t+j})$ at all horizons $j \geq 0$. Substituting these expectations into

[2] $l \geq 1$ is necessary for non-fundamentalness to be treated within the NC-VAR framework. If $l = 0$, a solution exists in the noncausal VARMA framework (Gouriéroux et al., 2020).

(3) and collecting terms in lagged production and shocks, the equilibrium price simplifies to:

$$p_t = -\frac{\rho}{1-\rho\beta} q_{t-1} - \frac{\chi + \beta^2}{1-\rho\beta}\epsilon_t - \frac{\beta}{1-\rho\beta}\epsilon_{t-1} - \frac{1}{1-\rho\beta}\epsilon_{t-2} + \nu_t \tag{4}$$

The pair $(q_t, p_t)'$ thus follows a bivariate autoregressive–moving average representation where the news shock enters both current and lagged components of the price.

We can write the joint dynamics of production and prices in compact matrix form as:

$$\begin{bmatrix} q_t \\ p_t \end{bmatrix} = \begin{bmatrix} \rho & 0 \\ -\frac{\rho}{1-\rho\beta} & 0 \end{bmatrix} \begin{bmatrix} q_{t-1} \\ p_{t-1} \end{bmatrix} + \underbrace{\begin{bmatrix} \chi + L^2 & 0 \\ -\frac{\chi+\beta^2}{1-\rho\beta} - \frac{\beta}{1-\rho\beta}L - \frac{1}{1-\rho\beta}L^2 & 1 \end{bmatrix}}_{=B(L)} \begin{bmatrix} \epsilon_t \\ \nu_t \end{bmatrix} \tag{5}$$

where L denotes the lag operator. The process $(q_t, p_t)'$ is fundamental if and only if the determinant of the moving average polynomial $B(L)$ has no zeros inside or on the unit circle, that is, if $|B(z)| \neq 0$ for all $|z| \leq 1$. In the present case, the determinant is given by

$$|B(z)| = \chi + z^2,$$

whose roots are $z = \pm i\sqrt{\chi}$. These roots have modulus $\sqrt{\chi} < 1$ whenever $\chi < 1$. Hence, as soon as the anticipated component dominates, the MA representation is non-fundamental, and no purely causal VAR can invert the observed process to recover the structural shocks (Gambetti and Moretti, 2017; Kilian and Lütkepohl, 2017; Gouriéroux et al., 2020; Nelimarkka, 2017a). In other words, the news shock is not spanned by past observables alone.

To study recoverability within a noncausal framework, we now consider the benchmark case $\chi = 0$ which allows us to derive an explicit noncausal representation. In this case, the production process reduces to

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

so that the news shock affects production only with a two-period delay. Substituting this expression into (4) and eliminating ϵ_t and its lags in favour of current and future values of q_t, we obtain the following noncausal representation for the price:

$$p_t = -q_t - \beta q_{t+1} - \frac{\beta^2}{1-\rho\beta} q_{t+2} + \nu_t \tag{6}$$

Equations $q_t = \rho q_{t-1} + \epsilon_{t-2}$ and (6) together describe a system in which production is purely backward-looking while prices depend on both current and future production, capturing the forward-looking nature of oil prices in the presence of anticipated supply disturbances.

This system can be recast as a multiplicative noncausal VAR(1,2). Let $y_t = (q_t, p_t)'$. Then the joint dynamics can be written as:

$$\underbrace{\left(I_2 - \begin{bmatrix} \rho & 0 \\ -\rho & 0 \end{bmatrix} L\right)}_{\Pi(L)} \underbrace{\left(I_2 - \begin{bmatrix} 0 & 0 \\ -\beta & 0 \end{bmatrix} L^{-1} - \begin{bmatrix} 0 & 0 \\ -\frac{\beta^2}{1-\rho\beta} & 0 \end{bmatrix} L^{-2}\right)}_{\Phi(L^{-1})} \begin{bmatrix} q_t \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} \epsilon_{t-2} \\ \nu_t \end{bmatrix} \quad (7)$$

where $\Pi(L)$ and $\Phi(L^{-1})$ are, respectively, the causal and noncausal polynomial matrices. The lead terms in $\Phi(L^{-1})$ encode the anticipation effects: they capture the dependence of current prices on future realizations of production through the anticipated component of the news shock. Since both $\Pi(L)$ and $\Phi(L^{-1})$ are invertible, the process y_t admits a two-sided MA representation.

Proposition 1. *Suppose that $|\rho| < 1$ and $|\beta| < 1$. Then both $\Pi(L)$ and $\Phi(L^{-1})$ have all roots outside the unit circle, and the process $y_t = (q_t, p_t)'$ admits a two-sided MA representation*

$$\Psi(z) = \Phi(z^{-1})^{-1}\Pi(z)^{-1}.$$

Moreover, the oil supply news shock ϵ_t is recoverable from the observable process, in the sense that

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

Proof. The determinant of the causal polynomial matrix is

$$\det \Pi(z) = 1 - \rho z,$$

which does not vanish for $|z| \leq 1$ when $|\rho| < 1$. Hence, $\Pi(L)$ is invertible in the closed unit disc. The noncausal polynomial $\Phi(z)$ in (7) is lower triangular with unit diagonal elements, so that $\det \Phi(z) = 1$ for all z. It is therefore invertible everywhere.

Since both polynomial matrices are invertible on and inside the unit circle, the two-sided inverse

$$\Psi(z)^{-1} = \Pi(z)\Phi(z^{-1})$$

is well defined, and the process y_t admits a two-sided moving-average representation with absolutely summable coefficients.

The structural shocks $(\epsilon_{t-2}, \nu_t)'$ can thus be expressed as two-sided linear combinations of the observables $(q_s, p_s) : s \in \mathbb{Z}$. By Proposition 1 of Chahrour and Jurado (2021), the existence of such a two-sided representation with invertible spectral density implies that the shocks are recoverable from the entire path of observables.

In particular, the news shock ϵ_t belongs to the closed linear span generated by $(q_s, p_s) : s \in \mathbb{Z}$ and is therefore measurable with respect to the information contained in the complete path of observables. Hence, the news shock is recoverable, which establishes the result. □

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. While a purely causal VAR fails to recover the news shock because it relies exclusively on past observables, the two-sided moving average representation induced by the NC-VAR spans the entire history and future of the observable process and thus fully captures the anticipated component of the shock. Our approach differs from Chahrour and Jurado (2021) in two important respects. First, instead of taking the causal–noncausal decomposition as given and verifying a rank condition on the spectral characteristic matrix, we estimate the lead–lag structure directly by maximum likelihood under non-Gaussian innovations, thereby ensuring recoverability by construction. Second, as shown in Online Appendix A, the multiplicative VMAR(r, s) representation (7) cannot be recast into the VAR(n_1, n_2, p) framework of Davis and Song (2020) and Gourieroux and Jasiak (2017), because the lead coefficient matrix Φ_2 is rank-deficient. This rank deficiency is not a numerical artefact but a structural feature of the triangular transmission mechanism of news shocks, in which production remains purely backward-looking while prices load on future production.

In the more general case in which both the anticipation horizon l and the contemporaneous response parameter χ are unknown, the exact two-sided moving average representation is infinite and cannot be written in closed form. This motivates the use of a flexible NC-VAR(r, s) specification in empirical applications, where the orders r and s are chosen to approximate the infinite two-sided representation up to a truncation horizon, and are ultimately determined by the data. The formal derivation for arbitrary $l \geq 1$ and $\chi \in [0, 1)$, together with the associated recoverability conditions, is provided in Online Appendix B.

3 Econometric framework

This section formalizes the econometric framework used in the empirical analysis. We first specify the mixed causal–noncausal VAR and the conditions under which it admits a two-sided moving-average representation. We then explain how impulse response functions and historical decompositions are computed in this setting, emphasizing how the noncausal structure generates anticipatory dynamics that are central for identifying oil supply news shocks.

3.1 The noncausal structural VAR

We adopt the Bayesian NC-VAR(r,s) (or VMAR(r,s)) model of Lanne and Saikkonen (2013) and Lanne and Luoto (2016), in which the n-dimensional vector y_t evolves according to

$$\Pi(L)\,\Phi(L^{-1})\,y_t = \epsilon_t, \tag{8}$$

where L is the backward shift operator, $\Pi(L) = I_n - \Pi_1 L - \cdots - \Pi_r L^r$ is the causal lag polynomial, $\Phi(L^{-1}) = I_n - \Phi_1 L^{-1} - \cdots - \Phi_s L^{-s}$ is the noncausal lead polynomial, and ϵ_t is a sequence of i.i.d. random vectors with zero mean and positive-definite covariance matrix Σ. Stationarity and the existence of a unique two-sided moving-average representation obtain under the standard invertibility conditions

$$\det \Pi(z) \neq 0 \quad \text{and} \quad \det \Phi(z) \neq 0 \qquad \text{for all } |z| \leq 1,$$

which ensure that both $\Pi(L)$ and $\Phi(L^{-1})$ are invertible on the closed unit disc.

Under these conditions, the process admits the two-sided MA representation

$$y_t = \sum_{j=-\infty}^{\infty} \Psi_j \epsilon_{t-j}, \qquad \Psi(z) = \Phi(z^{-1})^{-1}\,\Pi(z)^{-1}, \tag{9}$$

where $\Psi(z) = \sum_{j=-\infty}^{\infty} \Psi_j z^j$ is a matrix-valued Laurent series with absolutely summable coefficients. Premultiplying (8) by $\Phi(L^{-1})^{-1}$ yields

$$y_t = \Phi_1 y_{t+1} + \cdots + \Phi_s y_{t+s} + \sum_{j=0}^{\infty} M_j \epsilon_{t-j}, \tag{10}$$

where the matrices M_j collect the purely backward-looking component. Non-zero lead matrices Φ_i imply that y_t depends on future values of the process, so that past observables do not span the innovation space: the reduced-form representation is non-fundamental.

A crucial implication is that Gaussianity would render the NC-VAR observationally indistinguishable from a purely causal VAR(p) with $p = r + s$, because causal and noncausal representations share the same second-order properties under Gaussianity (Lanne and Saikkonen, 2013). In that case neither the lead–lag decomposition nor the structural shocks would be identified. To break this equivalence, we assume that ϵ_t follows a multivariate t-distribution,

$$\epsilon_t = \omega_t^{-1/2} \eta_t, \tag{11}$$

where $\eta_t \sim \mathcal{N}(0, \Sigma)$, $\lambda\omega_t \sim \chi^2_\lambda$, and λ denotes the degrees-of-freedom parameter. Small values of λ generate fat tails and excess kurtosis, while $\lambda \to \infty$ recovers the Gaussian case. The scalar factor $\omega_t^{-1/2}$ introduces a common stochastic volatility component that produces the higher-order dependence necessary to (i) ensure uniqueness of the mixed causal–noncausal representation and (ii) identify the lead–lag structure and the associated structural shocks (Gourieroux and Jasiak, 2017; Gouriéroux et al., 2020; Gourieroux and Jasiak, 2025, 2023).

We estimate the NC-VAR by Bayesian methods following Lanne and Luoto (2016). A Gibbs sampler alternates between draws of the dynamic coefficients, the covariance matrix Σ, the latent volatility factors $\{\omega_t\}$, and the degrees-of-freedom parameter λ, subject to the root-exclusion restrictions on $\Pi(L)$ and $\Phi(L^{-1})$. Minnesota-type priors are imposed on the lag and lead coefficients, with tighter shrinkage on leads than on lags to reflect the prior belief that most dynamics are causal and to reduce the risk of multimodality in the posterior distribution of the NC-VAR (Lanne and Luoto, 2016). Detailed prior settings and algorithmic steps are reported in Online Appendix D.

3.2 Impulse response function computation

Impulse response functions (IRFs) in the NC-VAR framework are derived from the structural two-sided MA representation implied by (9). Let the reduced-form innovations satisfy

$$\epsilon_t = \omega_t^{-1/2}\eta_t = \bar{B}\,\bar{u}_t,$$

where $\eta_t \sim \mathcal{N}(0, \Sigma)$, $\bar{u}_t$ is the vector of structural shocks, and $\bar{B}$, the rotation matrix identified using the Max-share approach developed in Section 4. Substituting into (9) yields

$$y_t = \sum_{h=-\infty}^{\infty} \Psi_h \bar{B}\,\bar{u}_{t-h} = \sum_{h=-\infty}^{\infty} \Theta_h \bar{u}_{t-h}, \tag{12}$$

where $\Theta_h = \Psi_h \bar{B}$. The response of variable i at horizon h to a unit realization of the j-th structural shock $\bar{u}_{j,t}$ is therefore

$$\frac{\partial y_{i,t+h}}{\partial \bar{u}_{j,t}} = e_i' \Psi_h \bar{B} e_j = e_i' \Psi_h \tilde{A} w_j = \Theta_{i,j,h}, \tag{13}$$

with e_i and e_j the canonical basis vectors and w_j the j-th column of W. In contrast to standard causal VARs, the horizon index h ranges over all integers $h \in \mathbb{Z}$. Negative horizons $h < 0$ capture anticipatory responses: they measure how variables move in advance of the shock, reflecting the noncausal component of the NC-VAR.

Numerically, the matrices Ψ_h are obtained by inverting the lag and lead polynomials in (9). Writing

$$\Psi(z) = \Phi(z^{-1})^{-1}\Pi(z)^{-1}$$

for $|z| = 1$, we expand $\Psi(z)$ as a Laurent series and recover the coefficients $\{\Psi_h\}$ via recursion. For $h \geq 0$, Ψ_h is generated by the inverse lag polynomial $\Pi(z)^{-1}$; for $h <$

0, the coefficients stem from the inverse lead polynomial $\Phi(z^{-1})^{-1}$. Under the root-exclusion conditions on Π and Φ, the coefficients decay geometrically as $|h|$ increases, so that truncating the series outside $[-H_{\text{irf}}, H_{\text{irf}}]$ yields a negligible approximation error.

Cumulated IRFs reported in the empirical analysis are defined as

$$\text{IRF}^{\text{cum}}_{i,j}(H) = \sum_{h=-s}^{H} \Theta_{i,j,h},$$

where the lower bound $-s$ corresponds to the maximum lead order of the NC-VAR. To quantify uncertainty, we compute $\{\Psi_h^{(m)}, \Theta_h^{(m)}\}$ for each posterior draw m from the Gibbs sampler and construct pointwise credible sets by taking the relevant quantiles across draws.

3.3 Historical decomposition

Historical decompositions in the NC-VAR rest directly on the structural two-sided MA representation (12). For each structural shock j, the contribution of that shock to variable i at date t is defined as

$$HD^{(j)}_{i,t} = \sum_{h=-\infty}^{\infty} \Theta_{i,j,h}\, \bar{u}_{j,t-h}, \tag{14}$$

where $\Theta_{i,j,h}$ denotes the (i,j)-th element of Θ_h and $\bar{u}_{j,t}$ is the j-th structural shock. Summing over all shocks reproduces the observed series,

$$y_{i,t} = \sum_{j=1}^{k} HD^{(j)}_{i,t},$$

so that the decomposition is exhaustive by construction.

In practice, (14) must be truncated. Under the root-exclusion conditions on Π and Φ, the structural coefficients Θ_h decay at a geometric rate as $|h| \to \infty$, which guarantees that truncation at a finite horizon H_{irf} entails only a small approximation error. We therefore compute

$$\widehat{HD}^{(j)}_{i,t} = \sum_{h=-H_{\text{irf}}}^{H_{\text{irf}}} \Theta_{i,j,h}\, \bar{u}_{j,t-h}, \tag{15}$$

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

A distinctive advantage of the noncausal framework is that it allows us to decompose the historical contribution of each shock into anticipation and realization components:

$$HD^{(j)}_{i,t} = \underbrace{\sum_{h=1}^{H_{\text{irf}}} \Theta_{i,j,-h}\, \bar{u}_{j,t+h}}_{\text{anticipation effects}} + \underbrace{\sum_{h=0}^{H_{\text{irf}}} \Theta_{i,j,h}\, \bar{u}_{j,t-h}}_{\text{realization effects}}. \tag{16}$$

The first term measures how future shocks, already anticipated by agents, influence current outcomes; the second corresponds to the conventional backward-looking contribution familiar from causal VARs. This decomposition is particularly informative for forward-looking variables such as the real oil price and global oil inventories, which display significant lead terms in the estimated NC-VAR.

Uncertainty in the historical decomposition is assessed in a fully Bayesian manner. For each posterior draw $m = 1, \ldots, M$ we obtain reduced-form residuals $\epsilon_t^{(m)}$, the structural impact matrix $\bar{B}^{(m)}$, the IRFs $\Theta_h^{(m)}$, and hence structural shocks $\bar{u}_t^{(m)} = (\bar{B}^{(m)})^{-1}\epsilon_t^{(m)}$. Applying (15) and (16) for each draw yields a posterior distribution of $HD_{i,t}^{(j)}$ and its anticipation and realization components, from which we report pointwise medians and credible sets. For variables whose estimated lead coefficients are not significantly different from zero, the anticipation term is negligible and the decomposition collapses to the standard backward-looking one.

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 NC-VAR specification, 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

Our baseline estimation sample spans January 1974 to July 2025.[3] We estimate a monthly NC-VAR for the vector

$$y_t = \left[\Delta wop_t,\ \Delta wip_t,\ RRAC_t,\ \Delta stocks_t\right]', \tag{17}$$

where Δwop_t is the monthly log difference of world oil production from the U.S. Energy Information Administration's Monthly Energy Review, Δwip_t is the monthly log difference of world industrial production, proxied by the OECD+6 industrial production index of Baumeister and Hamilton (2019), $RRAC_t$ is the real refiners' acquisition cost for imported crude oil (deflated by the U.S. consumer price index), and $\Delta stocks_t$ is the monthly log difference of global oil stocks constructed following Kilian and Murphy (2014). The

[3]Online Appendix C provides full details on data sources, construction, and transformations.

inclusion of global oil stocks is crucial: it both alleviates the informational deficiency of small VAR systems (Kilian and Murphy, 2014) and provides a direct observable through which anticipatory behavior in response to supply news can be assessed, since agents are expected to build (draw down) inventories in anticipation of future shortfalls (surpluses).

Inference is based on 10,000 posterior draws retained after a burn-in of 50,000 iterations. Section 5 discusses the selection of the causal lag order r and noncausal lead order s, while Online Appendix D reports the prior settings used in the baseline specification. In line with the oil-market VAR literature, we explore specifications with relatively rich dynamics in both lags and leads to ensure that the model can capture the slow adjustment of oil quantities and the potentially long anticipation horizons implied by forward-looking behavior.

4.2 Oil supply news shock identification

To identify the oil supply news shock, we build on the Max-Share methodology used in the news-shock literature (see, among others, Beaudry and Portier, 2006; Barsky and Sims, 2011; Forni et al., 2014; Beaudry and Portier, 2014; Kurmann and Sims, 2021) and on its recent implementation in a non-fundamental environment by Chahrour et al. (2023). The key idea is to select, within the set of admissible structural representations of the NC-VAR, the rotation that assigns to one shock the largest contribution to anticipated fluctuations in global oil production over a medium-run horizon.

Starting from the multivariate t-specification in (11), we write the reduced-form innovations as

$$\epsilon_t = \omega_t^{-1/2}\eta_t = \bar{B}\,\bar{u}_t, \tag{18}$$

where $\eta_t \sim \mathcal{N}(0,\Sigma)$, $\bar{u}_t = \omega_t^{-1/2}u_t^*$ with $u_t^* \sim \mathcal{N}(0, I_k)$, and $\bar{B} = \tilde{A}W$ for $\Sigma = \tilde{A}\tilde{A}'$ and an orthogonal matrix W. The common scalar volatility factor $\omega_t^{-1/2}$ induces tail dependence across the components of $\bar{u}_t$. As a result, second-moment restrictions alone do not uniquely identify the orthogonal rotation W, and additional information from higher-order moments or structural restrictions is required to pin down the structural decomposition.

We first impose restrictions on the noncausal polynomial $\Phi(L^{-1})$ to ensure that the forward-looking component of the NC-VAR enters the system in an economically interpretable way. Specifically, we require $\Phi(L^{-1})$ to have a block-triangular structure of the form

$$\Phi(L^{-1}) = I - \Phi_1 L^{-1} - \cdots - \Phi_s L^{-s}, \qquad \Phi_i = \begin{bmatrix} 0 & \mathbf{0}_{1\times(n-1)} \\ \Phi_{21,i} & \Phi_{22,i} \end{bmatrix}, \quad i = 1,\ldots,s, \tag{19}$$

and parameterize the unrestricted lead coefficients via

$$\phi = \text{vec}(\Phi) = R_\phi \phi_r.$$

Here ϕ_r collects the free elements of Φ, R_ϕ is a deterministic selection matrix, and $s^{=}ns$ zero restrictions are imposed on the first row of each Φ_i. In the ordering of our system, this implies that global oil production is purely backward-looking: it does not depend on future values of any variable at any lead horizon. We also set to zero the lead coefficients that would make world industrial production depend on future production, so that forward-looking dynamics operate exclusively through the real oil price and inventories. These restrictions are consistent with standard oil-market theory, which treats production and real activity as sluggish, and prices and inventories as the natural carriers of anticipatory behavior. Their implementation within the Bayesian estimation algorithm is detailed in Online Appendix D. Robustness exercises in Online Appendix E.3 show that relaxing the zero restrictions leaves the main results essentially unchanged.

The Max-Share criterion then pins down the first column of W. Using the two-sided MA representation (9), the time-t revision in the conditional forecast of global oil production at horizon τ can be written as

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

where $e_1 = [1\ 0\ 0\ 0]'$ selects global oil production. We identify the oil supply news shock as the shock that maximizes the contribution to the variance of these forecast revisions over a finite horizon $[H_1, H_2]$. Formally, we choose

$$w_1^{\text{news}} = \arg\max_{w_1} \sum_{\tau=H_1}^{H_2} e_1' \Psi_\tau \tilde{A} w_1 w_1' \tilde{A}' \Psi_\tau' e_1 \quad \text{subject to} \quad w_1' w_1 = 1. \tag{20}$$

The maximizer w_1^{news} is the eigenvector associated with the largest eigenvalue of

$$S = \sum_{\tau=H_1}^{H_2} (\Psi_\tau \tilde{A})' e_1 e_1' (\Psi_\tau \tilde{A}),$$

and the corresponding impact vector is $\gamma_1 = \tilde{A} w_1^{\text{news}}$, the first column of $\bar{B}$.

In contrast to Barsky and Sims (2011), we do not impose that the news shock be orthogonal to the current innovation in global oil production. There are two reasons for this choice (see Chahrour et al., 2023; Kurmann and Sims, 2021). First, because the NC-VAR features a two-sided MA representation, impact restrictions are difficult to interpret: the timing of the structural shock relative to observables is not known a priori, and constraining the impact response risks misaligning the identification with the underlying lead–lag structure. Second, relaxing the impact restriction improves robustness to measurement error and data revisions in production, and mitigates small-sample biases inherent in long-run identification schemes. To avoid rewarding purely short-run movements, we instead set $H_1 > 0$, so that the Max-Share criterion targets medium-run forecast revisions in production rather than impact effects. The resulting medium-run identification is formally analogous to the Max-Share procedures of Uhlig (2004) and Francis et al. (2014), adapted to the noncausal setting following Chahrour et al. (2023).

5 Baseline evidence on oil supply news shocks

In this section, we present the baseline evidence on oil supply news shocks obtained from the noncausal VAR specification developed above. The analysis focuses on the impulse response functions and historical decompositions implied by the estimated NC-VAR, highlighting how the presence of noncausal dynamics provides a coherent and economically interpretable account of forward-looking behavior in global oil markets.

We begin by assessing whether the data support a non-fundamental representation. As stressed by Lanne and Saikkonen (2013), Gourieroux and Jasiak (2017) and Davis and Song (2020), a non-Gaussian assumption on the reduced-form innovations is a necessary and sufficient condition for uniqueness of the NC-VAR representation and is therefore indispensable for structural identification. Figure A1, in Online Appendix E.1, reports the estimated marginal posterior density of the degrees of freedom (DOF) λ. The posterior distribution is sharply concentrated at low values, with a posterior mean around 3.75 and virtually zero probability mass above 6. This is in stark contrast with the prior mean, which we set equal to 8, and indicates that the data strongly favor fat-tailed innovations. In other words, the normality assumption is decisively rejected in favor of a multivariate t-distribution, thereby validating the use of a non-Gaussian NC-VAR(r, s) specification and ensuring uniqueness of the causal–noncausal decomposition.

This outcome has important implications for the representation of oil market dynamics. Since the NC-VAR with non-Gaussian errors is observationally distinct from any purely causal VAR of higher order, the evidence on λ suggests that a noncausal representation is statistically preferred to purely causal alternatives and provides a natural framework for modeling forward-looking behavior. This finding is fully consistent with the theoretical arguments of Chahrour and Jurado (2021) and Gouriéroux et al. (2020) who show that non-fundamentalness naturally arises in environments where expectations respond to information before physical quantities adjust.

Further support for noncausality comes from the lag and lead selection exercise reported in Tables A2 and A3 in Online Appendix E.2. Across a wide range of candidate NC-VAR(r, s) specifications, the preferred models systematically feature a strictly positive number of leads. The specification that maximizes the marginal likelihood is a VAR$(3, 4)$, while the Akaike Information Criterion selects a VAR$(1, 6)$. Although the exact decomposition between lags and leads is not unique, the message is unambiguous: specifications that restrict the process to be purely causal are dominated by models that allow for noncausal dynamics, regardless of the chosen selection criterion. The purpose of this exercise is not to pin down a unique pair (r, s), but rather to assess whether forward-looking terms are empirically warranted. On this dimension, the evidence is clear: noncausal components are essential for an adequate description of the data.

For the structural analysis, we therefore adopt a NC-VAR(12, 12) specification, which is deliberately more generous in terms of dynamic order than the information criteria would suggest. This choice is guided by two considerations. First, it aligns with standard practice in the oil market literature, where twelve monthly lags are routinely included to capture the slow adjustment dynamics of oil variables. Second, allowing for twelve leads ensures that the model can accommodate potentially long anticipation horizons, while maintaining symmetry between the causal and noncausal components. Our objective is thus not to estimate the most parsimonious NC-VAR, but to employ a sufficiently rich dynamic structure capable of capturing both delayed and anticipatory responses in the oil market, and thereby to provide a robust platform for the identification and interpretation of oil supply news shocks.

5.1 How does the oil supply news shock propagate through the oil market?

5.2 Transmission of oil supply news shocks in the oil market

We now examine how the identified oil supply news shock propagates through the oil market. The shock of interest is the first structural shock in the system, identified via the Max-Share procedure as the innovation that maximizes the contribution to forecast revisions of global oil production over a medium-run horizon. Figure 1 reports the cumulative impulse responses of the four baseline variables to a negative oil supply news shock, normalized so that the real price of oil increases by 10 percent on impact. This normalization does not impose any sign restriction ex ante; the sign and magnitude of the price response are entirely determined by the data through the identification scheme. The figure also displays 68 and 90 percent posterior credible sets, and, owing to the noncausal structure, responses are shown both at negative horizons (lead terms) and positive horizons (lag terms).

The news shock is identified using a truncation horizon of $[H_1, H_2] = [24, 24]$, meaning that the news shock is chosen to maximize the contribution to forecast revisions of global oil production at the two-year horizon.[4] This choice is motivated by the medium-run nature of oil supply adjustments. Although production may respond relatively quickly to new information, the full adjustment of global supply typically unfolds over several quarters because of production planning, capacity constraints, field operations, and strategic decisions by major producers. The 24-month horizon therefore targets the medium-run adjustment of global oil supply, allowing the identification procedure to isolate information about future supply conditions rather than short-run fluctuations. Accordingly, the identified disturbance should be interpreted as a shock conveying information about future oil supply conditions rather than as a conventional technology news shock with necessarily delayed effects.

Figure 1: Impulse-response functions to the oil supply news shock

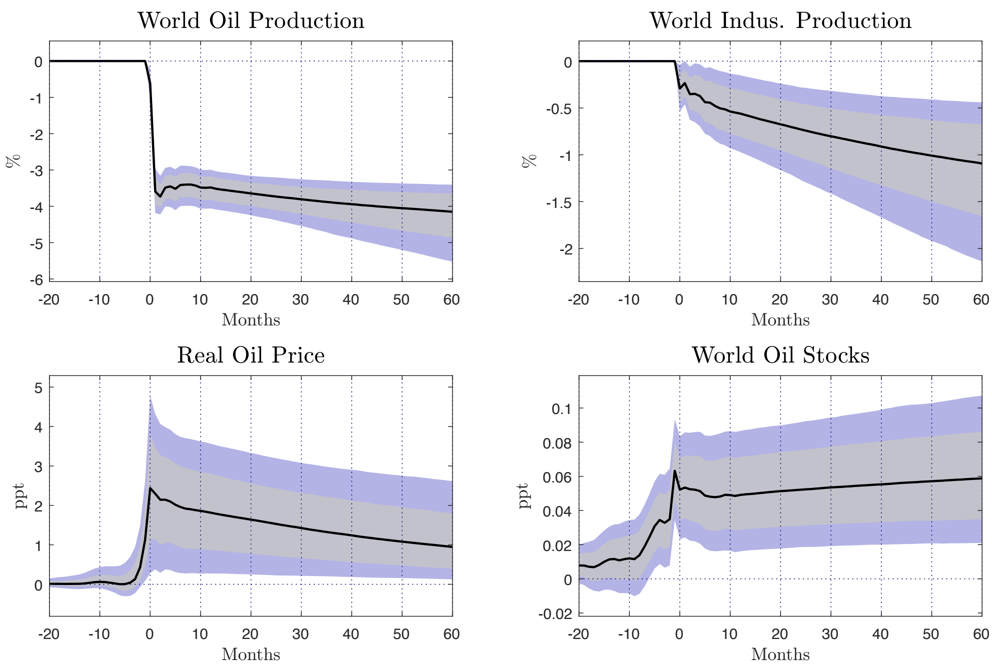

Notes: The black solid lines are the posterior median responses of the 4-variable baseline model from Baumeister and Hamilton (2019). The solid lines are the posterior median cumulated impulse responses of the NC-VAR(12,12). Shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because 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.

Turning to the dynamics themselves, the restrictions imposed on the lead coefficients imply that global oil production exhibits no anticipatory response to the news shock, while a contemporaneous response is allowed. Global oil production responds rapidly and declines persistently following the shock, eventually reaching about 4 percent below baseline. Although production is allowed to react contemporaneously under our identification strategy, this does not contradict the interpretation of the shock as conveying informa-

[4]Although the baseline identification uses [24, 24], we show in Section 5.5 that alternative truncation windows yield very similar impulse responses.

tion about future oil supply conditions. In the oil market, major producers-particularly OPEC members-can rapidly adjust output through the mobilization of spare production capacity or revisions to announced production targets. As a result, news about future supply shortfalls may affect current extraction decisions well before the underlying supply disturbance is fully realized.

The speed of the response suggests that the identified shock is more likely associated with news capable of influencing production decisions in the near term, such as geopolitical developments, strategic announcements, or anticipated supply disruptions (Känzig, 2021; Pinchetti, 2025), rather than information related to long-run changes in productive capacity arising from major discoveries or technological innovations (Arezki et al., 2017). At the same time, the persistence of the production decline indicates that the complete adjustment of global oil supply remains a medium-run process, consistent with the operational, investment, and strategic constraints characterizing the oil industry.

The most forward-looking variables, namely the real price of oil and global oil inventories, display pronounced anticipatory responses. On the lead side of the NC-VAR representation, both variables respond strongly and significantly before the shock fully materialises in production. On the causal side, after rescaling the responses so that the impact effect on the real price equals 10 percent, the price remains about 5 percent above baseline over the medium run. Inventories increase by roughly 3 percent on impact and continue to trend upwards, confirming the interpretation of the identified disturbance as a genuine news shock. Agents anticipate higher future scarcity and respond by accumulating precautionary inventories, rather than drawing them down as they would following an unanticipated physical disruption. This behavior is consistent with the predictions of rational-expectations storage models.

Finally, global economic activity displays a delayed but persistent contraction, reaching a trough of about 1 percent only after several years. This gradual response is consistent with a news-driven supply disturbance that initially operates through expectations, precautionary pricing, and inventory accumulation before fully propagating to physical supply conditions and, ultimately, to the broader economy.

Overall, Figure 1 provides direct evidence that allowing for noncausal dynamics fundamentally changes the interpretation of oil market fluctuations. Prices and inventories react in anticipation of future scarcity, while production and economic activity adjust subsequently, producing a pattern that closely matches the predictions of forward-looking storage models. These responses are consistent with the interpretation of the identified disturbance as an aggregate oil supply news shock. In Section 7, we show that this aggregate shock is primarily driven by information embedded in OPEC production decisions.

5.3 How much the oil supply news shock explains real oil price fluctuations over time?

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 turn to the historical decomposition of real oil prices, which allows us to quantify the contribution of oil supply news shocks to price fluctuations over different historical episodes.

Figure 2 plots the cumulative (causal and noncausal) contribution of the oil supply news shock alongside the demeaned real price for the period 1975–2025. The figure indicates that oil supply news shocks account for a sizeable and highly time-varying share of real oil price fluctuations over the last five decades. Their contribution is concentrated around major episodes involving revisions in expectations about future supply conditions, while remaining limited during episodes that are predominantly demand-driven. Around the Iranian Revolution, news shocks exert a strongly negative in the mid to late 1970s, consistent with downward revisions to expected future supply that had already been priced in before the full shock materialised. This is followed by a substantial positive contribution of supply news from the Iran-Iraq War through the mid-1980s, even as real oil prices gradually decline, indicating that other shocks, such as demand factors or inventory adjustments, offset part of the news-driven component.

The period 1980–1985 is particularly informative. Rapid growth in non-OPEC production and growing doubts about OPEC's ability to sustain production discipline likely contributed to persistent upward revisions of expected future supply. The historical decomposition attributes an important share of the subsequent decline in oil prices to these supply-news effects, culminating in the 1986 collapse following Saudi Arabia's abandonment of production constraints.

By contrast, during the Gulf War of 1990–91, oil supply news plays only a modest role in explaining the temporary price spike. This is consistent with a narrative in which market participants anticipated compensating production from core OPEC producers and thus did not revise long-term supply expectations as dramatically as spot prices might suggest. The decomposition also shows that oil supply news shocks contribute meaningfully to the trend reversal in oil prices in the late 1990s, culminating around the Venezuelan crisis in the early 2000s, and again play a significant role around the 2014 oil price collapse, where their contribution turns sharply negative.

The onset of the COVID–19 pandemic represents a crucial validation episode for our identification strategy. The unprecedented collapse in oil prices in March and April 2020 was driven primarily by a sudden and severe contraction in global demand. An appro-

Figure 2: Contribution of oil supply news shocks to real oil price fluctuations

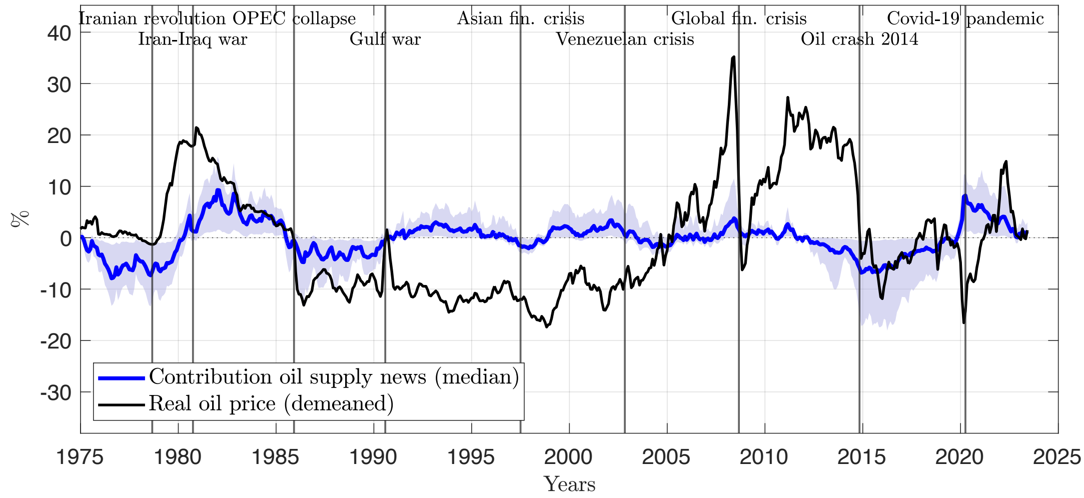

Notes: The solid blue line shows the average contribution of the oil supply news shock to the real crude oil price, estimated using the NC-VAR(12,12). The light blue shaded area represents the 90% credible intervals. The solid black line depicts the demeaned real oil price.

priately identified oil supply news shock should therefore not account for this collapse. The historical decomposition confirms this interpretation: the contribution of oil supply news shocks during early 2020 is positive and small, while the bulk of the price collapse is attributed to other shocks, consistent with a demand-driven narrative. At the same time, oil supply news shocks contribute positively to the subsequent price recovery, reflecting expectations of production cuts and supply management in response to the demand contraction. This pattern supports the view that the identified disturbance captures expectations about future supply conditions rather than misclassifying demand shocks as supply news. The COVID-19 episode therefore provides an important validation of the identification strategy.

Further insight into the nature of the identified disturbance comes from decomposing the contribution of the news shock into lead and lag components, as reported in Figure A6 in Online Appendix F. This result highlights an important distinction between the existence of anticipatory dynamics and their quantitative contribution to observed price fluctuations. The decomposition reveals that the lag component accounts for the dominant share of the price fluctuations associated with the news shock, while the lead component is quantitatively smaller over most of the sample. This outcome is economically intuitive. Oil prices reflect not only anticipatory forces, but also contemporaneous market conditions, endogenous propagation mechanisms and interactions with other shocks. As a result, the purely anticipatory component of the disturbance, captured by the lead terms, appears more limited in magnitude than the full contribution. Nevertheless, the lead component becomes economically meaningful during major geopolitical and supply-related episodes, indicating that expectations about future supply conditions do affect prices before the underlying physical disturbance materializes.

Overall, the historical decomposition confirms that oil supply news shocks represent a quantitatively important source of oil price fluctuations. Their contribution is concentrated around episodes characterized by major revisions in expectations about future supply conditions, whereas events dominated by demand disturbances, such as the COVID-19 collapse, are largely attributed to other shocks. These findings provide further support for the interpretation of the identified disturbance as a genuine oil supply news shock and underscore the importance of accounting for anticipatory behavior in oil market analysis.

5.4 How much do oil supply news shocks explain oil inventory fluctuations?

We now turn to global oil inventories, which provide perhaps the most direct evidence that the identified disturbance operates through expectations about future supply conditions. In storage-based models of the oil market, inventories constitute the primary margin through which agents respond to anticipated scarcity. As such, their behavior offers a particularly informative test of the news-shock interpretation. Figure 3 reports the historical decomposition of global oil inventories associated with the identified oil supply news shock, distinguishing between the lead and lag components of the bilateral moving-average representation. The figure shows that the contribution of the news shock to inventory fluctuations is economically large and statistically significant during major geopolitical and oil-market episodes, including the Iranian Revolution, the Iran–Iraq War, the Gulf War, the Global Financial Crisis, the 2014 price collapse and the COVID–19 pandemic.

The decomposition indicates that the lag component is quantitatively dominant over much of the sample, reflecting the fact that a substantial fraction of inventory dynamics is driven by the endogenous propagation of the news shock through the oil market. However, the lead component is far from negligible. It becomes particularly prominent during episodes of heightened geopolitical uncertainty, when market participants revise expectations about future supply and adjust their precautionary inventory holdings ahead of the actual supply disruption. This finding is particularly noteworthy because inventories represent a physical quantity directly chosen by market participants. Unlike prices, which are affected by many contemporaneous forces, inventory accumulation reflects an active decision to transfer oil across time in response to expected future supply conditions.

In sum, the inventory decomposition provides some of the strongest evidence in favor of the news-shock interpretation. The fact that inventories respond in advance of major supply-related events and display a larger anticipatory component than oil prices is consistent with the predictions of rational-expectations storage models, in which stock accumulation serves as the primary channel through which expectations about future scarcity affect current market outcomes.

Figure 3: Contribution of oil supply news shocks to global oil inventories

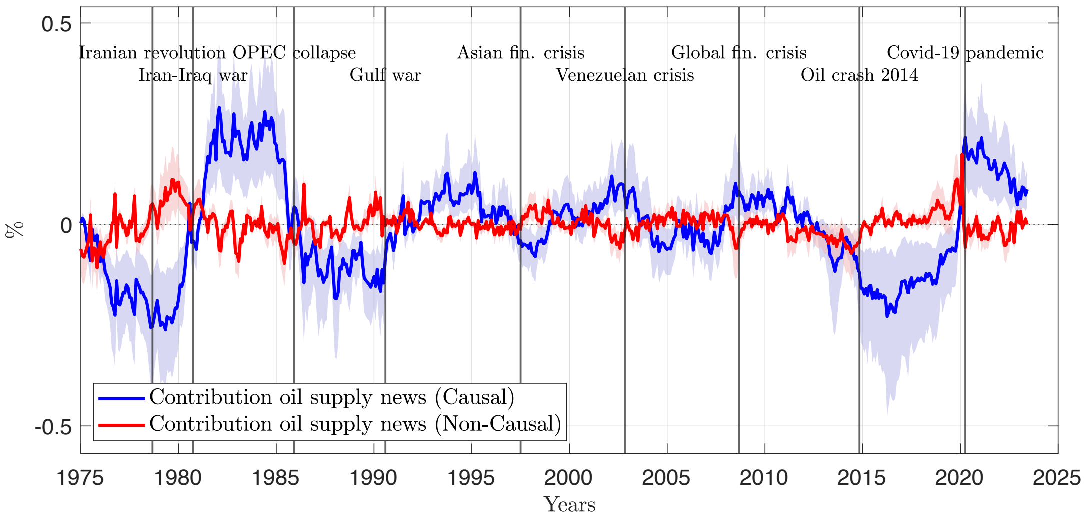

Notes: The figure reports the historical decomposition of global oil inventories associated with the identified oil supply news shock in the NC-VAR model. The decomposition separates the contribution of the lead component (red) and the lag component (blue) of the bilateral moving-average representation. Shaded areas denote the associated 90% credible intervals. Vertical lines indicate major geopolitical and oil-market events. Positive contributions of the lead component are interpreted as reflecting anticipatory inventory accumulation in response to expected future supply disruptions. Sample: 1975–2025.

5.5 Robustness and identification Diagnostics

How sensitive are the results to the choice of truncation horizons?

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 A3, in Online Appendix E.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.

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 (H_1,H_2) 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. Although the magnitude and timing of the oil price response vary somewhat across truncation windows, its qualitative behavior remains unchanged: prices increase significantly and persistently following the identified shock.

How robust are the results across different sample periods?

Figure A4 in Online Appendix E.5 shows that re-estimating the model using a smaller sample size yields qualitatively similar results to the benchmark specification 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. This result is not surprising, as it can be explained by the declining contribution of oil supply news shocks to real oil price fluctuations in the second subsample, consistent with the historical decomposition reported in Figure 2 and discussed in 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.

How distinct is the identified news shock from alternative supply shocks?

A potential concern with Max-Share identification is that the identified shock may partly reflect other disturbances that generate similar production dynamics. Following Dou et al. (2025), contamination may arise when the target shock is not sufficiently orthogonal to other shocks that generate similar responses of the target variable. We therefore perform two complementary exercises. First, we verify that the identified news shock is internally well separated from the second max-share shock within the NC-VAR system. Second, we compare our identified news shock with the oil supply surprise shock of Baumeister and Hamilton (2019). Specifically, we project the impulse response of global oil production associated with the Hamilton and Baumeister shock onto the impulse response generated by our Max-Share news shock. Large projection coefficients indicate that the two shocks generate similar production dynamics and therefore provide evidence of potential overlap. However, as emphasized by Dou et al. (2025), such overlap should not be interpreted mechanically as evidence of structural equivalence and must be assessed jointly with the broader multivariate responses of the system.

Table 1: Contamination diagnostics: propagation horizons ($H = [0, 20]$)

Contaminating shock	ζ	[16%, 84%]	$\hat{\beta}$
Contemporaneous (2nd eig.)	0.00%	[0.00, 0.00]	-0.000
BH Supply	39.53%	[36.79, 42.61]	0.547

Notes: $\hat{\beta}$ measures the projection of a candidate shock's oil production response onto that of the NC-VAR max-share news shock; ζ is the associated contamination weight following Dou et al. (2025). 'Contemporaneous (2nd eig.)" denotes the second eigenvector of the bilateral max-share problem over $[-12, 20]$. 'BH Supply" refers to the Baumeister and Hamilton (2019) oil supply shock and is evaluated via an external projection diagnostic. Statistics are computed over $h \in [0, 20]$ using only oil production responses, thus capturing similarity in production dynamics rather than full structural equivalence. All shocks are normalized to a negative production response. Sample: 1975–2017.

Table 1 reports the contamination diagnostics based on the projection procedure of Dou et al. (2025). We begin by examining the internal identification diagnostic associated with the second eigenvector of the max-share eigenproblem. The estimated projection coefficient is essentially zero ($\hat{\beta} = 0$), implying a contamination weight of $\zeta = 0$. This result indicates that the identified NC-VAR news shock is orthogonal to the closest competing direction within the max-share space and therefore appears to be well separated from alternative shocks generated by the same identification procedure.

We next consider the external diagnostic based on the Baumeister-Hamilton surprise supply shock. By construction, the Dou et al. (2025) diagnostic relies on projections of impulse responses for a single target variable, which in our case is global oil production. Since both negative oil supply news shocks and negative surprise supply shocks induce similar declines in oil production, a non-negligible projection coefficient is mechanically expected and should not be interpreted in isolation as evidence of contamination. In fact, a near-zero projection would be difficult to reconcile with the notion of a supply news shock, as it would imply that information about future oil supply contains little predictive content for realized supply conditions.

Consistent with this reasoning, the overlap with the Baumeister-Hamilton shock is economically meaningful but far from complete, with an estimated contamination weight of approximately 40%. More importantly, the diagnostic is constructed using oil production only, whereas the economic content of the identified shock is determined by the joint response of all oil-market variables. In this respect, the NC-VAR news shock is sharply differentiated from the Baumeister-Hamilton surprise supply shock. While both shocks reduce oil production, the former leads to an increase in oil inventories whereas the latter implies a decline in inventories. This divergence is consistent with standard economic mechanisms: news about future supply disruptions induces precautionary inventory accumulation, whereas an unexpected contemporaneous disruption reduces available stocks. These results therefore suggest that the identified NC-VAR news shock remains economically distinct from a conventional surprise supply shock despite some overlap in production dynamics.

How important are noncausal dynamics for the identified news shock?

To assess the role of noncausal dynamics in the identification of oil supply news shocks, we compare the impulse responses obtained from the baseline NC-VAR(12,12) specification with those generated by a purely causal VAR(12) benchmark. Figure A5, in Online Appendix E.6 reports the corresponding responses.

The comparison reveals substantial differences between the two specifications. In particular, allowing for noncausal dynamics generates stronger and more persistent responses of oil inventories and oil prices, while also introducing anticipatory dynamics prior to impact that are mechanically absent from the causal VAR. These differences are precisely what one would expect if information about future supply conditions is not fully captured by a purely causal representation. By construction, the causal VAR cannot accommodate anticipatory responses, whereas the NC-VAR allows expectations to affect current market outcomes before physical quantities adjust. The comparison therefore highlights the importance of accounting for non-fundamentalness when identifying and interpreting oil supply news shocks.

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

While the macroeconomic consequences of unexpected oil supply disruptions have been extensively studied, much less is known about the effects of shocks that operate through expectations about future oil supply conditions. In this section, we first examine how oil supply news shocks affect the global economy before turning to their effects on U.S. macroeconomic variables.

Figure 4 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 the world stock price index (MSCI). The additional variables are included in the system one at a time in order to avoid estimating an excessively large model, which would raise important concerns regarding parameter proliferation and estimation precision in a noncausal VAR framework. Given the dimensionality of the data, we follow a standard augmentation strategy and include each additional variable separately in the baseline

Figure 4: Reactions of global macroeconomic variables to a news shock

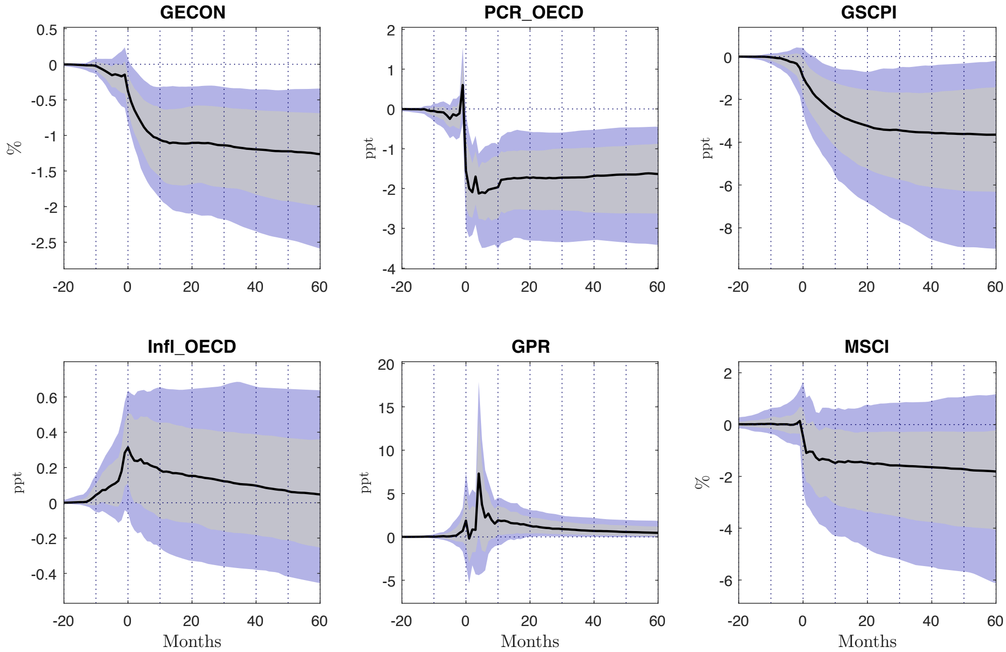

Notes: The solid black lines are the posterior median cumulated impulse responses of the NC-VAR(12,12) and shaded areas represent the 90 and 68 % credible sets. Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds to the lead terms of the MA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

system.

The results reveal a broad and persistent deterioration in global economic conditions following a negative oil supply news shock. GECON gradually declines, reaching around 1 percent below baseline after roughly two years and remaining persistently depressed thereafter, echoing the slow but sizeable contraction in global industrial production. OECD passenger car registrations display a remarkably similar response, which is in line with the sensitivity of durable consumption and particularly car demand to expected future fuel costs and heightened macro-uncertainty. Particularly noteworthy is the response of the Global Supply Chain Pressure Index (GSCPI), which sharply contrasts with the response typically observed following 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 effect starting already in the lead terms and stabilizing at a negative level. Finally, the MSCI exhibits a modest decline, but the response remains statistically insignificant throughout the horizon. The effect of anticipated oil supply shocks on stock prices appears to be as negligible as that of unanticipated shocks (Kilian and Park, 2009).[5] 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.

Taken together, these results suggest that oil supply news shocks generate stagflationary forces at the global level. Expectations of future supply scarcity weaken economic activity, depress confidence, and increase geopolitical uncertainty, while their effects on global financial markets remain comparatively limited.

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

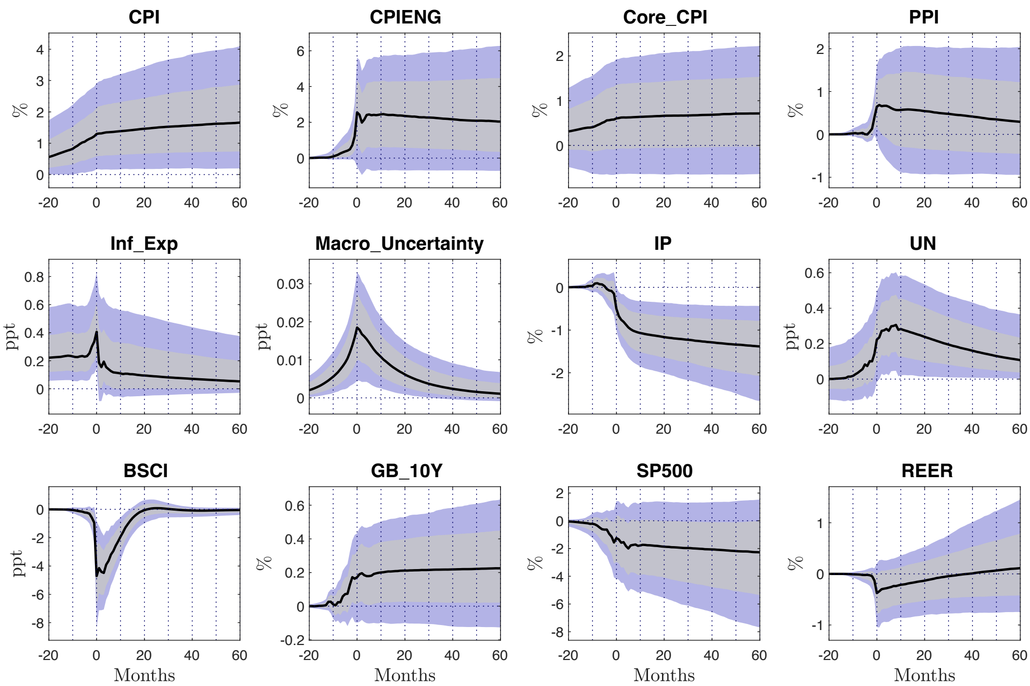

Notes: The solid lines are the posterior median cumulated impulse responses of US macroeconomics variables. Blue and grey shaded areas are the 90 and 68 % credible sets of the NC-VAR(12,12). Because of noncausality, the impulse responses are located on both sides of zero. The negative side corresponds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

As for the effect of the oil supply news shock on the U.S. economy, Figure 5 reports

[5]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.

the responses of a broad set of U.S. macroeconomic and financial indicators,[6] covering prices, inflation expectations, real activity, labor market conditions, confidence measures, financial variables, and exchange rates.[7]

As shown in the top panel of Figure 5, 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 prevents 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-

[6]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

[7]The impulse responses of the U.S. variables reported in Figure 5 are obtained by augmenting the baseline four-variable NC-SVAR model with one additional variable at a time. While this strategy ensures tractability, it is more restrictive in the case of U.S. variables, as it does not allow for a full characterization of domestic interactions and feedback effects within the U.S. economy. Accordingly, the impulse responses reported here should be viewed as conditional responses within a parsimonious system, rather than as structural responses derived from a fully specified U.S. macroeconomic model.

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.

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

While global oil production delivers a clear supply-news shock, it remains silent on its underlying sources. To shed light on the origin of the identified news, we reapply the same identification strategy using disaggregated production measures, sequentially replacing global oil production with OPEC and non-OPEC production.

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.[8]

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), 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 en-

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

dogenous 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 6: Impulse-response functions to the OPEC and Non-OPEC oil news shock

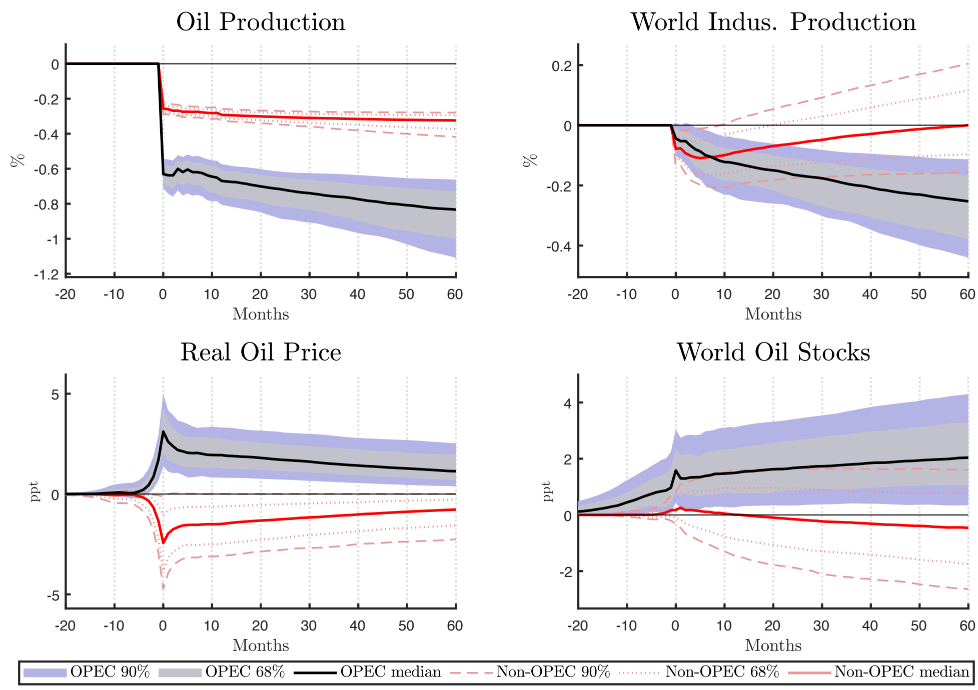

Notes: Black lines show the posterior median cumulated impulse responses of a 4-variable NC-VAR(12,12) identifying the oil supply news shock via OPEC production, while red lines show a separate NC-VAR(12,12) using Non-OPEC production. Shaded areas and dashed/dotted lines indicate 90% and 68% credible 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.

Figure 6 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 response of oil prices. While OPEC-related news raises oil prices, non-OPEC news generates a significant price decline. 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 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 evidence suggests that the aggregate oil supply news shock identified in the baseline model is predominantly an OPEC-driven phenomenon. While OPEC production embeds information about future supply scarcity and coordinated production management, non-OPEC production appears to convey information that is more closely related to cyclical demand conditions and endogenous supply responses.

Figure 7 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, non-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 A7 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

Figure 7: Contribution of OPEC an Non-OPEC oil news shocks to oil price fluctuations

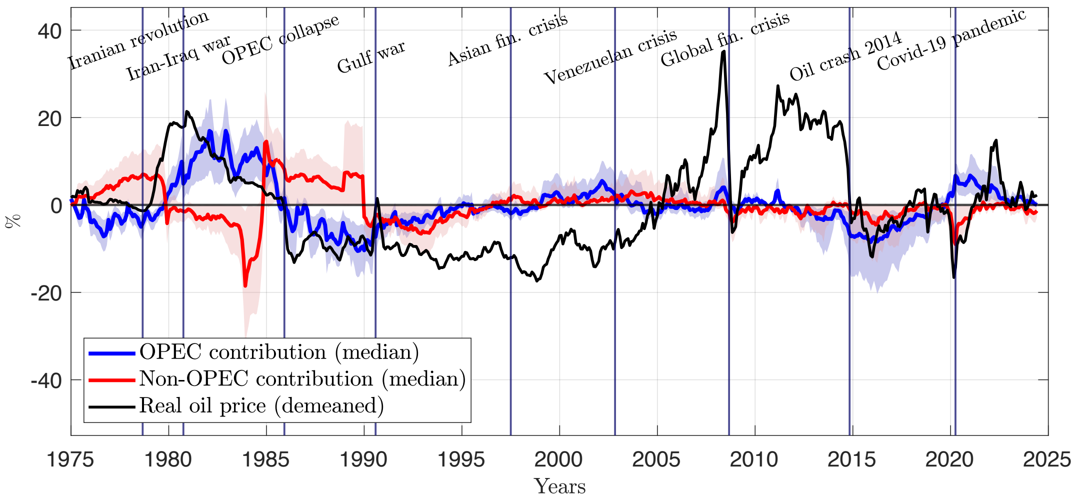

Notes: The solid blue line shows the median contribution of the oil supply news shock identified from OPEC oil production, while the solid red line shows the contribution identified from Non-OPEC oil production. Light blue and light red shaded areas indicate the 90% credible sets for OPEC and Non-OPEC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

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.

The close similarity between Figure 4 and Figure A7, in Online Appendix G.1 strengthens the argument that the baseline news shock identified from aggregate production is effectively an OPEC-driven expectations shock at the global level.

In contrast, Figure A8, in Online Appendix G.1, 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 A9, in Online Appendix G.2, reproduces the U.S. macro-financial responses to an OPEC-based news shock and confirms the stagflationary pattern documented in Figure 5, and clarifies that the US stagflation response is especially pronounced when the news shock is explicitly tied to OPEC.

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 A10, in Online Appendix G.2, 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.

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 reinforce the conclusion that the aggregate oil supply news shock identified in the baseline model is fundamentally an OPEC-driven phenomenon. While OPEC-based news shocks generate persistent stagflationary pressures and broad macro-financial repercussions, non-OPEC-based shocks appear considerably closer to demand-driven disturbances, producing weaker and often qualitatively different responses across both global and U.S. variables.

8 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.

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

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.

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., Cormun, V., De Leo, P., Guerrón-Quintana, P.A., Valchev, R., 2024. Exchange Rate Disconnect Revisited. Working Paper 32596. National Bureau of Economic Research.

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

D'Amico, S., King, T.B., 2023. What does anticipated monetary policy do? Journal of Monetary Economics 138, 123–139.

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.

Dou, L., Ho, P., Lubik, T.A., 2025. Max-Share Misidentification. Technical Report. Federal Reserve Bank of Richmond Working Paper 25-02.

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

Forni, M., Gambetti, L., 2016. Government spending shocks in open economy VARs. Journal of International Economics 99, 68–84.

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.

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.

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.

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., 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.

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.

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.

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.