Energy Economics 132 (2024) 107488

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneeco

Unraveling the structural sources of oil production and their impact on CO2 emissions

Helmut Herwartz [a], Bernd Theilen [b,*], Shu Wang [a]

[a] Chair of Econometrics, University of Göttingen, Humboldtallee 3, D-37073 Göttingen, Germany
[b] Departament d'Economia, RESCAT and ECO-SOS, Universitat Rovira i Virgili, Reus, Spain

ARTICLE INFO

JEL classification:
C14
C32
Q43

Keywords:
Fossil fuels
Oil demand and supply shocks
CO2 emissions
Climate change
Energy policy
Global crude oil market
Oil demand and supply elasticity
Structural VAR model

ABSTRACT

In this study, we examine the structural short- and long-run effects of oil supply and demand shocks on the production of crude oil. Among these, oil supply shocks are the major determinant of oil production. Furthermore, we adopt local projections to elicit that oil supply and aggregate demand shocks are significant drivers of CO2 emissions, whereas oil-specific demand shocks have only limited impact on overall emissions in the short-, mid-, and long-term. These findings underscore the limitations of current demand-side policies within a selected group of countries, emphasizing the necessity for global support and binding commitments to effectively reduce emissions from oil production and meet the climate targets set by the Paris Agreement in 2015.

1. Introduction

Since the ratification of the Kyoto Protocol in 1997, the United Nations Framework Convention on Climate Change (UNFCCC) has been striving to mitigate the effects of climate change. The primary objective of the UNFCCC is to curtail greenhouse gas (GHG) emissions, with a key focus on promoting the substitution of fossil fuels with renewable energy sources. Specifically, the reduction of oil demand is crucial, given that, out of the global CO2 emissions totaling 35.26 billion tonnes in 2021, oil ranked as the second-largest contributor after coal, accounting for 11.84 billion tonnes of CO2 emissions (Enerdata, 2023). Indeed, the relationship between oil production and CO2 emissions from fossil fuels, as depicted in Fig. 1, suggests a close correlation. Major policy instruments employed to combat global warming include carbon pricing through taxes or the implementation of jurisdiction-specific cap-and-trade or emission trading systems (ETS) and the promotion of green energy usage through subsidies. The underlying hypothesis for the effectiveness of these policy instruments posits that any reduction in the demand for fossil energy inevitably leads to a decrease in GHG emissions. However, the efficacy of policies aimed at reducing CO2 emissions by decreasing oil demand also depends on the response of oil producers. If oil production exhibits an inelastic response to fluctuations in oil prices, this policy may prove ineffective in reducing GHG emissions globally. In this regard, eyeballing the volatile variations in the international real price of oil and the consistently growing path of oil production in Fig. 2 provides a visual indication of the limited responsiveness of oil production to price changes. Against this background, gaining a profound understanding of the mechanisms governing prices and quantities in oil (and other energy) markets becomes paramount for designing and implementing effective energy and climate protection policies.

In this study, we analyze the influence of demand and supply shocks on global crude oil production in the short-, medium-, and long-term within a structural vector autoregressive (SVAR) model for the global crude oil market, utilizing monthly data spanning from 1974 to 2022. In the macroeconomic literature, SVAR models have become a standard tool for understanding the joint evolution of real oil prices and the level of crude oil production (see Kilian and Zhou, 2023, for a recent review of this rich literature). In the context of the current literature, our study makes three substantial contributions. Firstly, while the major focus of the related literature aims to explain short-run price fluctuations in the oil market (e.g. Kilian and Murphy, 2012; Baumeister and Hamilton, 2019), our focus is on the determinants of mid- and long-term oil production. Secondly, whereas most studies on

* Corresponding author.
E-mail address: bernd.theilen@urv.cat (B. Theilen).

https://doi.org/10.1016/j.eneco.2024.107488
Received 14 September 2023; Received in revised form 12 February 2024; Accepted 13 March 2024
Available online 18 March 2024

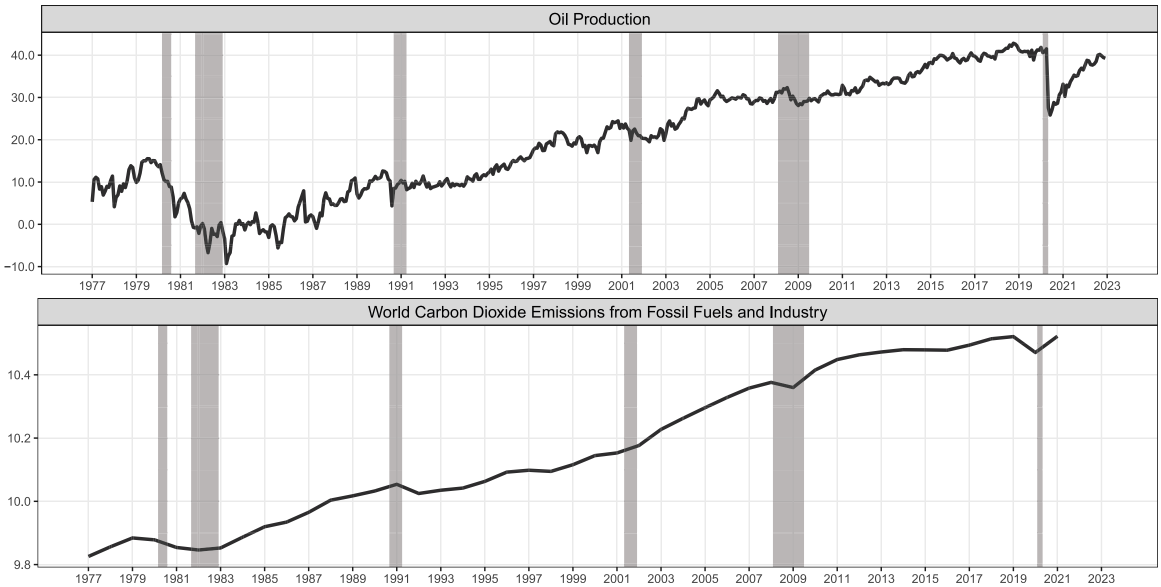

Fig. 1. Oil production and CO2 emission. Upper panel: Reconstructed time series of oil production. The series of oil production is constructed as the sum of cumulative effects of four shocks displayed in Fig. 5 and has been adjusted for truncation bias. Lower panel: World carbon dioxide emissions from fossil fuels and industry (million tonnes in natural logarithm). Shaded areas denote recessions dated by the Business Cycle Dating Committee of the National Bureau of Economic Research. Subjecting the bivariate time series of oil production and CO2 emissions to a formal cointegration analysis yields a Johansen trace statistic of 25.09, which indicates common trending with 5% significance.

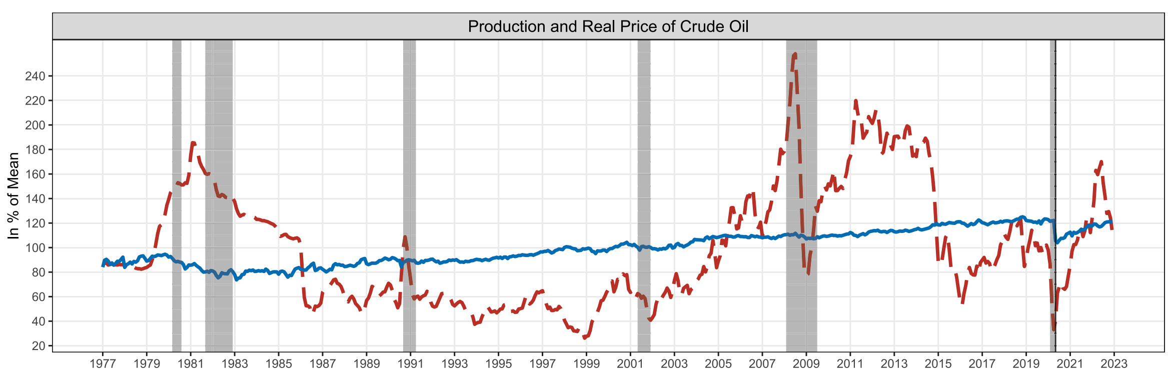

Fig. 2. Production of crude oil including lease condensate (blue solid) and real price of crude oil (red dashed) in percentage of the mean. The vertical line indicates the month of May 2020, which marks the implementation of the OPEC+ decision to reduce oil production by 9.7 MB/D. For further notes see Fig. 1.

the oil market rely on a-priori parametric restrictions, sign restrictions on directional effects of structural shocks, or upper bounds on model implied price elasticities (e.g., Kilian, 2009; Kilian and Murphy, 2012; Baumeister and Hamilton, 2019), we employ a state-of-the-art statistical identification approach that builds upon concepts of independent component analysis, and is fully agnostic with respect to both short- and long-term implications (Hafner et al., 2023). Thirdly, subsequent to identification of soundly defined structural shocks shaping global oil market dynamics, we perform local projections of annual growth of CO2 emissions on structural oil market shocks. This exercise is pivotal to analyze the extent to which structural oil supply and demand shocks impact on CO2 emissions, and allow us to obtain insights into the effectiveness of climate protection policies.

The agnostic approach to identification of the dynamic oil market system results in four shocks with soundly distinct economic features, namely an oil supply shock, an aggregate demand shock, an oil-specific demand shock and an inventory demand shock. Model implications reveal that demand shocks exert only mild and scarcely significant short- and long-term effects on oil production. Additionally, the oil supply curve exhibits inelasticity, leading to oil price fluctuations in response to demand shifts but with minimal repercussions on actual production. Furthermore, a historical decomposition of oil production highlights the limited influence of demand shocks on past trends in oil production. Local projections reveal that oil supply shocks significantly impact CO2 emissions, especially in the short-run. Oil-specific demand and inventory demand shocks have minimal effects on CO2 emissions in the short-, mid-, and long-term, while an aggregate demand shock influences emissions notably in the medium- to longer-term. These findings elucidate the insufficient impact of current demand-side policies, endorsed by a limited number of countries, in achieving the 1.5 °C target set by the 2015 Paris Agreement through efforts to reduce oil demand and promote renewable energy sources. The results also underscore the imperative need for global backing of demand-side policies to genuinely curtail CO2 and other GHG emissions stemming from oil production. This could be accomplished, for example, through the establishment of a global cap-and-trade system, aligning with the collaborative endeavors in recent climate summits to ensure universal adherence to binding commitments on GHG emissions.

The remainder of this study is organized as follows. The next section outlines the structural VARs and the adopted identification approach. Section 3 discusses empirical results. Section 4 analyzes the impact of structural supply and demand shocks on CO2 emissions and addresses policy implications. Section 5 concludes. An Online Appendix provides some estimation results in full detail and a collection of robustness checks.

2. SVARs and kernel maximum likelihood estimation

A structural analysis of the global oil market faces the core problem of identifying the underlying (short-run) supply and demand curves, as unexpected price and quantity changes (see, e.g., Fig. 2) are compatible with different slopes of these curves (Kilian, 2022). As a promising means offering such a structural perspective, Kilian (2009) has pioneered the SVAR analysis of the global oil market and complements notions of structural oil supply shocks with different demand shocks that are traced back to distinct motives for oil consumption. In this section, we start by formalizing the VAR model in reduced and structural form, and outline the identification problem. In addition, we highlight various advances that have been suggested in the literature to identify structural shocks by means of SVARs. Among these approaches, we put particular emphasis on tools of Independent Component Analysis (ICA) and the recent Kernel Maximum Likelihood (KML) approach of Hafner et al. (2023), which has been used for the empirical analysis in this work.

2.1. SVAR model

A K-dimensional VAR model reads in *reduced* and *structural* form, respectively, as

$$y_t = \mu + A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t, \tag{1}$$

$$\text{and } y_t = \mu + A_1 y_{t-1} + \cdots + A_p y_{t-p} + B\xi_t, \quad t = 1, \ldots, T. \tag{2}$$

The models outlined in (1) and (2) are conditional on presample values $y_0, y_{-1}, \ldots, y_{1-p}$, μ is a vector of intercepts and the $K \times K$ matrices $A_1, \ldots, A_p$ capture autoregressive dynamics up to order p. By assumption, the dynamic process is stationary, i.e., $\det(A(z)) \neq 0$ for all $|z| \leq 1$ with $A(z) = I_K - A_1 z - \cdots - A_p z^p$. Let L denote the backshift operator such that $Ly_t = y_{t-1}$. By implication of the stationarity condition, y_t has a Wold representation $y_t = \nu + \sum_{i=0}^{\infty} \Theta_i \xi_{t-i}$, where $\nu = A(L)^{-1}\mu$ and $\Theta(L) = A(L)^{-1}B$ and Θ_i is a matrix of structural impulse responses at the ith horizon. Long-run cumulative responses are given by $\Theta(1) = \sum_{i=0}^{\infty} \Theta_i = A(1)^{-1}B$. While the structural model in (2) formalizes the transmission of structural shocks to the observable system, causal relationships among the model variables are often inferred from the following representation

$$\Pi_0 y_t = c + \Pi_1 y_{t-1} + \cdots + \Pi_p y_{t-p} + \xi_t, \tag{3}$$

where $\Pi_0 = B^{-1}$, $c = B^{-1}\mu$ and $\Pi_l = B^{-1}A_l$ for $l = 1, \ldots, p$.

Although the reduced form model (1) can be consistently estimated by means of LS methods, the causal relationship $u_t = B\xi_t$ and the dynamic transmission of elements in ξ_t to observable system variables in y_t remains latent. Without loss of generality, unit-variance normalization can be imposed, $\mathrm{Cov}[\xi_t] = I_K$. Accordingly, the covariance of the mixed residuals u_t allows for the decomposition $\mathrm{Cov}[u_t] \equiv \Sigma_u = BB'$. However, this moment condition results in $K(K+1)/2$ unique equations, which are insufficient to determine K^2 unknown parameters in B. Owing to its indeterminacy, alternative choices for the matrix B are observationally equivalent in Gaussian SVARs. A space of alternative covariance decompositions can be obtained by rotating the lower triangular Cholesky factor C of the reduced form covariance such that

$$\Sigma_u = BB^\top = CQQ^\top C^\top = CQ(CQ)^\top, \tag{4}$$

where Q denotes an orthogonal rotation matrix ($QQ^\top = I_K$, $|\det Q| = 1$). Hence, to describe the causal relationships within a dynamic system, the selection of an identified model (or a set of identified models) can be seen as the choice of a specific rotation matrix Q (or of a set of rotation matrices) in (4).

2.2. Theory-based vs. statistical identification in a nutshell

The SVAR literature yet comprises several approaches to solve the identification problem that allow for a broad classification into theory-based and statistical approaches (we refer the reader to Kilian and Lütkepohl, 2017, for a textbook treatment of diverse identification schemes). Typical elements of the structural parameter matrix B, denoted b_{ij}, quantify the direction and magnitude of the contemporaneous effects of a (positive) structural unit shock ξ_{jt} on the ith variable in the system. Hence, economic theory might offer plausible restrictions for both characteristics either in the form of strong parametric (exclusion) restrictions (Sims, 1980; Blanchard and Quah, 1989), or as weaker so-called sign restrictions imposed on effect directions and/or relative magnitudes (e.g., Faust, 1998; Uhlig, 2005). As an alternative to theory-guided identification, an analyst might consult statistical identification approaches that derive from informative statistical properties of the supposed structural shocks in ξ_t. On the one hand, non-proportional changes in the variances of structural shocks have been shown to carry informational content to solve the identification problem. Lütkepohl and Netšunajev (2017) review the literature on identification through heteroskedasticity, which has yet suggested a multitude of informative second-order moment structures such as variance shifts, smooth transition models, or conditional heteroskedasticity. Magnusson and Mavroeidis (2014) have suggested a general family of identification schemes utilizing stability restrictions. On the other hand, a rich literature has emerged (see Section 2.3 below for selected references) using the fact that linear combinations of independent and identically non-Gaussian distributed shocks ξ_t can be uniquely recovered from mixed reduced form residuals $u_t = B\xi_t$ (Comon, 1994). It is worth pointing out that statistical identification schemes yield typically a structural parameter matrix B that is only locally identified (i.e., unique up to column permutation and scaling).

Although the two branches of identification in SVARs can be seen to establish alternative modeling strategies, several arguments can be put forth to highlight the scope of viewing theory-based and statistical identification as promising complements. First, theory-based identification schemes are typically just-identifying such that the data cannot object against the identifying assumptions imposed. Hence, theory-based identification is always at the risk to intermingle modeling assumptions and conclusions (Uhlig, 2005). In this regard, statistical identification offers additional information that could enable an explicit testing of otherwise just-identifying (economic) hypothesis. As an example, Lütkepohl and Netšunajev (2014) confirm core identifying assumptions of Kilian and Murphy (2012) based on informative heteroskedasticity for identification and testing. Second, the literature on set identification by means of sign restrictions has pointed to the challenges of eliciting particular models as ‘most reasonable’. In this regard, Herwartz and Wang (2023) have argued in favor of the informational content of independence criteria for the elicitation of a single representative of a set of identified models that all align with sound economic theory. Third, while the supposed informative statistical characteristics of the shocks (i.e., heteroskedasticity or independence & non-Gaussianity) can be diagnosed straightforwardly, statistically identified shocks may not inherently exhibit economic interoperability. In other words, once these shocks are identified, providing meaningful economic labels to them – a process often referred to as ‘shock labeling’ – becomes a vital aspect of structural modeling through statistical identification. The inclusion of theoretical considerations or established narrative features of shocks is paramount in assisting analysts in appropriately labeling these shocks. To delve further into this complementary perspective, we proceed to elaborate on the details of identification by means of ICA.

2.3. KML estimation of the structural parameters

If the elements of ξ_t are mutually independent and at most one of them is Gaussian, the full B matrix can be uniquely recovered from reduced form residuals $u_t = B\xi_t$ (Comon, 1994). Several techniques have been proposed to identify B under the independence assumption: The minimization of non-parametric dependence diagnostics (Herwartz, 2018; Herwartz and Plödt, 2016), pseudo maximum likelihood (PML) based on non-Gaussian fixed densities (Lanne et al., 2017; Gouriéroux et al., 2017), generalized methods of moments (GMM) relying on higher-order moment conditions (Lanne and Luoto, 2021; Keweloh, 2021), ML based on kernel density estimates (KML, Hafner et al., 2023) and discrete location-scale-mixed normal distributions (DLSMN Fiorentini and Sentana, 2023). While PML using fixed densities may suffer from density misspecifications and GMM may be sensitive to heavy-tailed distributions, both DSLMN and KML approaches are consistent for estimating key structural quantities, regardless of the true shock distributions. Through comprehensive Monte Carlo studies, Hafner et al. (2023) demonstrate that the KML approach exhibits superior finite-sample performance compared with a wide range of competing approaches under various data-generating processes. In particular, KML estimation is remarkably robust in the presence of heteroskedastic and co-heteroskedastic components, which is especially crucial for the oil market model given the substantial time span covered by our empirical analysis. Therefore, we consider KML as a promising approach to detect structural innovations underlying the global oil market. We proceed by briefly describing this approach.

Given the decomposition in (4), it is easy to verify that structural shocks can be interpreted as rotated orthogonalized residuals, i.e., $\xi_t = Q^\top \tilde{\xi}_t$ with $\tilde{\xi}_t = C^{-1}u_t = C^{-1}A(L)y_t$, $\tilde{\xi}_t \sim (\mathbf{0}, I_K)$. With f_i denoting the probability density function (pdf) of the ith component of ξ_t, the log-likelihood and gradient are given by

$$\sum_{t=1}^{T}\sum_{i=1}^{K} \log f_i\left(e_i{}^\top Q^\top \tilde{\xi}_t\right) + T\log|\det Q^\top| \quad \text{and} \quad \sum_{t=1}^{T} S(Q^\top\tilde{\xi}_t)\tilde{\xi}_t{}^\top + TQ^{-1}. \tag{5}$$

In (5), e_i is the ith column of an identity matrix with conformable dimension, $S : \mathbb{R}^K \mapsto \mathbb{R}^K$ is the component-wise mapping with $S(\boldsymbol{x}) = (s_1(x_1), \dots, s_K(x_K))^\top$ and $s_i = f_i'/f_i$ is the density score. Given that the true densities of the shocks are unknown, Hafner et al. (2023) recently proposed a non-parametric maximum likelihood approach of utilizing kernel density estimates. Let the orthonormal matrix Q be parameterized by a vector of rotation angles $\theta \in \boldsymbol{\Theta}$, where $\boldsymbol{\Theta}$ is a compact subspace of $\mathbb{R}^{(K(K-1)/2)}$.[1] Moreover, h denotes a bandwidth parameter that converges to zero but $Th/\log(T) \to \infty$ as $T \to \infty$. Then, the kernel density estimator for the ith component $\xi_{it} = e_i{}^\top Q^\top \tilde{\xi}_t$, with $Q = Q(\theta)$ is

$$\hat{f}_{ih\theta}(\xi_{it}) = \frac{1}{Th}\sum_{s=1}^{T} K\left(\frac{\xi_{it}(\theta) - \xi_{is}(\theta)}{h}\right), \; i = 1, \dots, K,$$

where $K(x) : \mathbb{R} \mapsto \mathbb{R}_0^+$ is a uniformly bounded, integrable univariate kernel function with bounded first derivative such that $\int K(u)du = 1$, $|u|K(u) \to 0$ as $|u| \to \infty$ and $K(0) \geq \delta > 0$. In this work, we chose a Gaussian kernel $K(x) = (2\pi)^{-1/2}\exp(-x^2/2)$, such that for a given θ, the rth derivative of the density can be estimated as

$$\frac{d^r}{d\xi_{it}^r}\hat{f}_{ih} = \frac{1}{Th^{r+1}}\sum_{s=1}^{T}(-1)^r \mathcal{H}_r\left(\frac{\xi_{it} - \xi_{is}}{h}\right) K\left(\frac{\xi_{it} - \xi_{is}}{h}\right),$$

where $\mathcal{H}_r(x)$ is the rth Hermite polynomial (Bhattacharya, 1967). Based on the estimated densities and scores, the KML estimator of the structural orthogonal mixing matrix is obtained by maximizing the non-parametric pseudo log-likelihood function

$$\hat{Q} \equiv Q(\hat{\theta}) = \arg\max_{\theta\in\Theta} \sum_{t=1}^{T}\sum_{i=1}^{K} \log \hat{f}_{ih\theta}(\xi_{it}(\theta)). \tag{7}$$

Hafner et al. (2023) show that under regularity conditions of the true densities f_i, $i = 1, 2, \dots, K$, wherein they are uniformly continuous with bounded derivatives almost surely and bounded away from zero, along with relatively mild conditions on the kernel functions $K_h(\cdot)$ and the bandwidth parameter h, the KML estimator converges in probability to the true parameter value as the sample size increases.

3. A structural analysis of the global oil market

In this section we first introduce the analyzed system of global oil market time series variables, and motivate the use of an ICA-based identification scheme for the global oil market. Subsequently, we discuss the key implications of the structural global oil market model identified through statistical independence. We first examine the structural IRFs and demonstrate solid economic foundations of the identified components that allow us to interpret them as shocks in a strict structural sense. Next, we investigate the short- and – more importantly – long-term sensitivities of oil supply and demand to oil prices. Unlike previous studies that mainly focus on short-term price fluctuations, our study aims to understand the structural drivers of oil production over the short-, medium-, and long-term. We present evidence on the structural composition of forecast error variances that are specific to oil production. Additionally, we examine the historical evolution of actual oil production by decomposing it into underlying shocks originating from the supply and demand sides of the global crude oil market. Throughout, we complement point estimates with bootstrap-based inferential results. We follow Brüggemann et al. (2016), who suggest a moving block bootstrap for inference on structural model parameters. As suggested by these authors, the block length is set to 25 being the closest integer to $5.03T^{1/4}$, $T = 563$. Inferential results rely on $M = 1000$ bootstrap replications.

3.1. Data and agnostic identification of oil market shocks

3.1.1. Time series variables

The data are measured at monthly frequency and, following the recent literature, the sample period starts in 1974M1. Adding to available empirical evidence, our sample includes recent observations and ends in 2022M12. In line with leading studies in the related literature (e.g. Kilian and Murphy, 2012, 2014; Baumeister and Hamilton, 2019), our dynamic system consists of quotes of oil production and prices, a measure of (global) economic activity and information on oil inventories. Specifically, for oil production Q_t, we use the (logarithmized) crude oil including lease condensate (MB/D) published by the U.S. Energy Information Administration (EIA). The index of global real economic activity EA_t is constructed by Kilian (2009) based on dry cargo bulk freight rates. This index has several advantages as a measure of current global economic activities (see Kilian and Zhou, 2018). Firstly, it gives proper weight to emerging economies whose real economic activity is often underestimated in official GDP figures and incorporates changes in output composition, import propensity, and country weights. Secondly, it responds quickly to changes in aggregate

[1] More specifically, Q is given in the form of a sequence of Givens rotation matrices (see Gouriéroux et al., 2017; Hafner and Herwartz, 2023) of the form

$$Q(\theta) = \left(\prod_{i=1}^{K-1}\prod_{j=i+1}^{K} \mathcal{G}_{i,j}(\theta_n)\right)^\top.$$

The rotation angles θ_n are defined on the interval $(0, \pi]$ with $n = 1, \dots, K(K-1)/2$ and $\mathcal{G}_{i,j}(\theta_n)$ is the Givens matrix that rotates the subspace spanned by axes i and j while holding other axes fixed. For instance, in the case of $K = 3$, one has

$$\begin{aligned} Q(\theta) &= \left(\mathcal{G}_{1,2}(\theta_3)\mathcal{G}_{1,3}(\theta_2)\mathcal{G}_{2,3}(\theta_1)\right)^\top \\ &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_1 & -\sin\theta_1 \\ 0 & \sin\theta_1 & \cos\theta_1 \end{bmatrix}\begin{bmatrix} \cos\theta_2 & 0 & -\sin\theta_2 \\ 0 & 1 & 0 \\ \sin\theta_2 & 0 & \cos\theta_2 \end{bmatrix}\begin{bmatrix} \cos\theta_3 & -\sin\theta_3 & 0 \\ \sin\theta_3 & \cos\theta_3 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \end{aligned} \tag{6}$$

demand and can be constructed in real-time. Finally, it recognizes that fluctuations in the volume of commodity shipping may be higher than in global GDP or industrial production as it responds to shifts in expectations about future levels of industrial production. The (logarithmized) real oil price series P_t is obtained on the basis of the refiner acquisition cost of imported crude oil provided by the US Department of Energy, and deflated by the US consumer price index. The US consumer price index is collected from FRED, Federal Reserve Bank of St. Louis (code: CPIAUCSL). The data were downloaded in April 2023. Following Kilian and Murphy (2014), we use OECD data as a proxy for global oil inventories. Specifically, we denote with S_t the OECD petroleum stocks that include crude oil and lease condensates (with strategic reserves) provided by the EIA.[2] Since consistent time series data for OECD petroleum stocks are only available since 1989, we use data from Kilian and Murphy (2014) to extrapolate the changes in OECD inventories prior to 1989 on the basis of the growth rates of US petroleum stocks. Except for EA, all time series enter our analysis in the form of (rescaled) changes of natural logarithms. With $\Delta = 1 - L$ denoting the first difference operator, the four-dimensional vector of jointly endogenous variables is $y_t = (\Delta Q_t \times 100, EA_t, \Delta P_t \times 100, \Delta S_t)$. Before being subjected to model estimation and structural analysis, the data are seasonally adjusted using dummy variables for the monthly seasonal pattern.

3.1.2. Identification of global oil market shocks

To identify the structural model (2) within the global oil market model, various approaches have been suggested, including, e.g., the assumption of a triangular B matrix (i.e., setting $Q = I$, Kilian, 2009), or the imposition of weak sign restrictions, with a refinement of admissible models achieved by means of (i) bounds on price elasticities (Kilian and Murphy, 2012), or (ii) narrative information on the relatively minor contribution of aggregated demand shocks to oil price changes in September and October 1980 (outbreak of Iran–Iraq War) and August 1990 (the outbreak of the Persian Gulf War) (Antolín-Díaz and Rubio-Ramírez, 2018). With a focus on identifying the causal relations in (3) directly, Baumeister and Hamilton (2019) have suggested a Bayesian approach that combines sign and exclusion restrictions with informative priors on core marginal effect parameters. Further studies on crude oil markets in this vein are, for instance, Braun (2021) and Valenti et al. (2023). While these informed approaches to the identification of oil market shocks have provided valuable insights into causal relations governing the interdependence, for instance, among oil production and oil prices, it is worth noting that Herwartz and Plödt (2016) have broadly confirmed central findings of Kilian and Murphy (2012, 2014) relying on a purely agnostic and ICA-based identification approach.

Witnessing the informational content of independent components for identification in oil market models, we utilize the non-parametric KML estimator to estimate the structural model. This approach is chosen for its robustness in handling unknown source distributions and its favorable performance in small samples, as demonstrated by Hafner et al. (2023). In particular, the estimates remain consistent in the presence of potential heteroskedasticity of the structural shocks. The implementation of the KML estimator in (7) requires the analyst to opt for a particular bandwidth for density estimation. For this purpose, we utilize the rule-of-thumb suggested by Silverman (1998), and set $h = 1.06T^{1/5}$. This choice accounts for the fact that (orthogonalized) components $\tilde{\xi}_t$ and their rotations have an identity unconditional covariance by construction.

3.2. KML estimation and identifying assumptions

The KML-approach to the identification of B (and hence ξ_t) is fully agnostic, and it is unclear if the unique independent components allow for an interpretation as economically meaningful shocks in a structural sense. Explicit representations of the structural parameter estimates in $\hat{Q}$ and $\hat{B}$ (i.e., impact effects) are documented in Online Appendix A joint with bootstrap means and t-ratios.

The identification of the structural model in (2) relies upon the assumptions that the marginal distribution of at most one element in ξ_t is Gaussian and elements in ξ_t are mutually independent. Testing the null hypothesis of Gaussianity by means of Jarque–Bera tests for each identified component ξ_{it}, $i = 1, 2, 3, 4$ results in highly significant rejections for almost all component estimates with p-values below 0.001, with estimates of ξ_{4t} being the only exception (p-value 10.4%). Furthermore, we test the null hypothesis of mutual independence by means of two non-parametric tests that have been shown in the literature to have power against unspecified forms of dependence. To assess significance of particular independence statistics, both approaches rely on 1000 permutations. On the one hand, the test based on the so-called distance covariance statistic of Székely et al. (2007) results in an insignificant diagnostic of 7.88E−05 (with p-value of 0.441). On the other hand, the Cramér–von Mises statistic of Genest et al. (2007) takes a value of 0.021 which results in a p-value of 67.38%.[3]

3.3. Shock labeling

3.3.1. Impulse response analysis

We next discuss structural estimation results and address, in particular, the so-called shock labeling problem. For this purpose, Fig. 3 shows impulse response estimates for the estimated independent components in $\{\xi_t = \hat{B}^{-1}\hat{u}_t\}_{t=1}^{T}$, where $\hat{u}_t$ are the LS estimates retrieved from the reduced form model in (1).

The first row of Fig. 3 shows impulse responses of the first element of ξ_t on the system variables. It emerges that the estimated dynamic effects of ξ_{1t} are unique in the sense that the short- and medium-term relationship between the responses of produced oil quantities and real oil prices are inverse. Such a sign pattern is not observed for any other component in ξ_t. Therefore, ξ_{1t} is the only component estimate that qualifies as a potential candidate to subsume properties of an *oil supply shock*. An oil supply shock refers to a sudden and unexpected disruption in the supply of oil to the global market. Such shocks are often caused by geopolitical events, natural disasters, or other unforeseen circumstances that disrupt the production, refining, or transportation of oil (Kilian and Murphy, 2014). With this in mind, an exogenous increase of oil production is not unlikely to result in a build up of oil inventories as signified by IRFs in the top right panel of Fig. 3. Put differently, in the argumentation of Kilian and Murphy (2014), the build-up of oil inventories signifies that the direct effects of an enhanced oil production exceed the indirect inventory effects that are channeled through the negative price effect of additional oil supply. Although the estimated positive impact response of economic activity to an oil supply shock aligns with sign restrictions suggested for identification in Kilian and Murphy (2012), the economic impact of supply shocks on longer-term activity is mild and lacks significance (see second panel in the first row of Fig. 3). Throughout, estimated dynamic effects on economic activity are insignificant at a 5% level, which also holds for the respective IRFs in Baumeister and Hamilton (2019).

[2] The monthly time series data for crude oil including lease condensate production, U.S. crude oil composite acquisition cost by refiners and OECD petroleum stocks are downloaded from the https://www.eia.gov website of the U.S. Energy Information Administration. Kilian and Murphy (2014) argue that, unlike inventory growth rates, simple inventory changes can be considered covariance stationary.

[3] While the distance covariance statistic is determined from the components ξ_{it}, $i = 1, 2, 3, 4$, in their metric form, the Cramér–von Mises statistic is derived from (joint) rank statistics. Accordingly, the distance covariance statistic promises some power advantages, whereas the Cramér–von Mises test can be considered more robust in case that particular higher-order (co)moments do not exist.

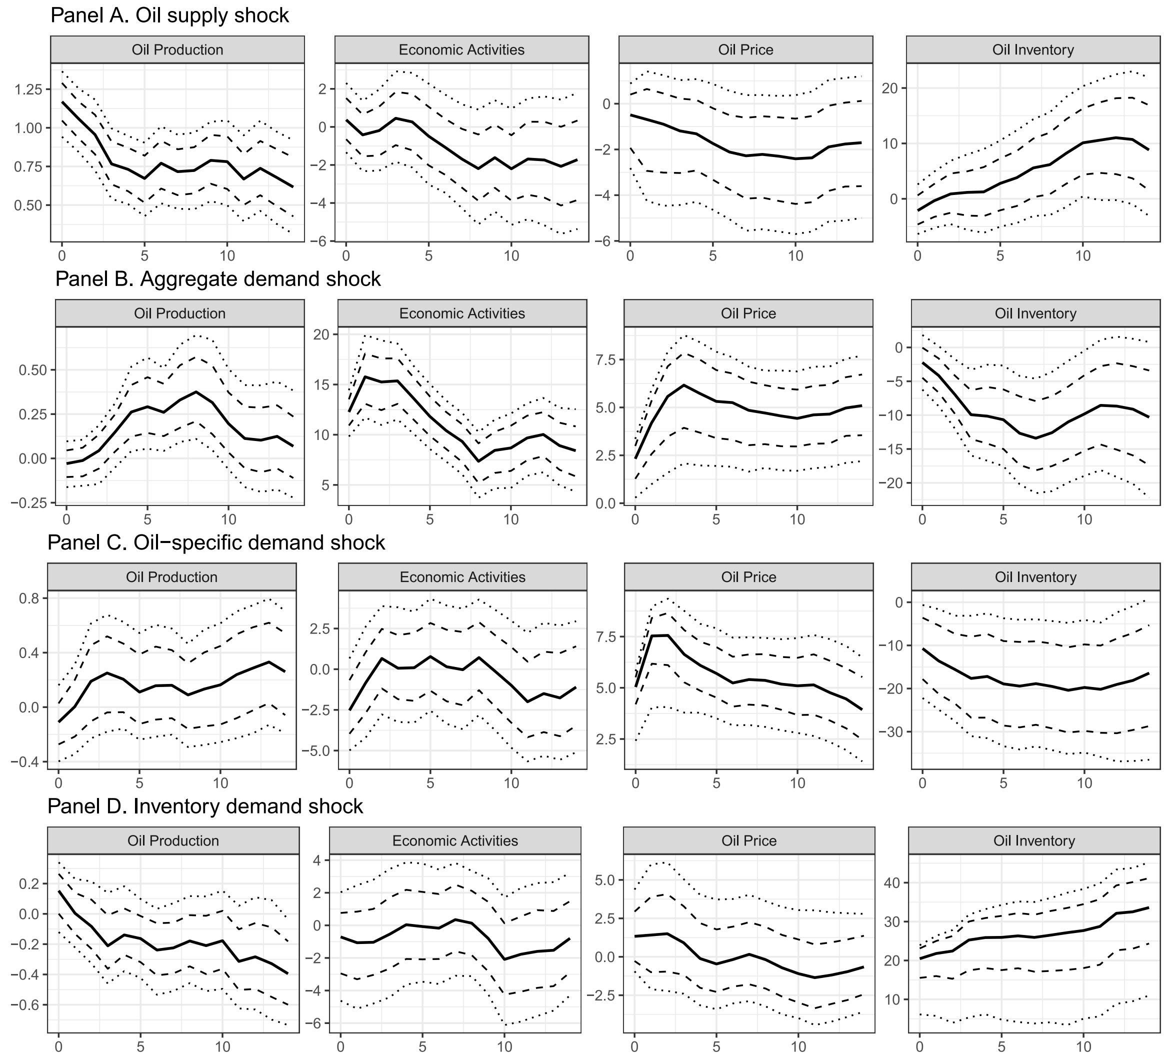

Fig. 3. Structural IRFs to unit shocks ξ_{it}. Panels show medians (solid curves) and confidence bonds with 68% (dashed) and 90% pointwise coverage (dotted) from 1000 replications of a moving block bootstrap resampling scheme. Except for the second column (Economic activity) all panels show accumulated effects.

As noted before, all remaining components $\xi_{it}, i = 2,3,4$, are characterized by moving oil production and real oil prices in the same direction in the short- to medium-term. Accordingly, these components align with stylized properties of demand shocks. Hence, it appears conducive to outline heterogeneous motives to demand oil and, if possible, label identified components $\xi_{it}, i = 2,3,4$, as corresponding structural demand shocks.

The second row of Fig. 3 shows the effects of ξ_{2t}, which predict a marked improvement of economic activity over all horizons and persistent positive effects on the real price of oil. In the literature such a shock has been interpreted as an *aggregate demand shock* (or, alternatively, as a flow oil demand or income shock, see, for example, Kilian and Murphy, 2012; Baumeister and Hamilton, 2019). This shock captures surprise information originating in the global business cycle. Pointing to the informational content of the adopted ICA approach to identification, it is worth highlighting that – despite a very slight violation of the theoretical sign pattern of Kilian and Murphy (2012) on impact – the response profiles generated by ξ_{2t} clearly align with both the theoretical sign pattern of Kilian and Murphy (2012) and related empirical evidence documented in Baumeister and Hamilton (2019). While oil producers increase quantities in response to this shock in the medium-term, their initial responses lack significance. Aligning with its ambiguous effects on stored oil (Kilian and Murphy, 2014), the aggregate demand shock results in a reduction of oil inventories in the medium-term but lacks clear directional effects on inventories upon impact.

The third row of Fig. 3 shows the dynamic effects of ξ_{3t}. This shock exerts the strongest effect on oil prices among all identified demand components $\xi_{it}, i = 2,3,4$. Moreover, it reduces economic activity on impact significantly and invokes an enhancement of oil production in the medium- to longer-term. Owing to its marked (and direct) effect on oil prices, it is intuitive to refer to this shock as an *oil-specific* demand shock. Speculative (or precautionary) motives have been put forth to explicitly characterize this demand shock (e.g., Kilian and Murphy, 2012, 2014). While such an interpretation is widespread in the literature, the work of Baumeister and Hamilton (2019) points more to purely consumptive motives to materialize this shock economically. As an effective way to differentiate between both demand motives, it is promising to consult the evolution of oil inventories in response to ξ_{3t}. While theoretical arguments would predict a positive response of oil inventories to exogenous shocks to precautionary oil demand, the estimated impulse responses are negative. As a specific demand for consuming oil could also be covered by reducing oil inventory, it is plausible to observe that oil inventories shrink in response to an oil-specific consumption demand shock. Unlike the aggregate demand shock, this shock exerts a contractionary, mildly significant impact effect on economic activity. While this impact effect seems at odds with theoretical considerations in Baumeister and Hamilton (2019), dynamic IRFs lack significance throughout, similar to posterior credibility intervals shown in Baumeister and Hamilton (2019).

The dynamic effects of the fourth independent component ξ_{4t} are displayed in the fourth row of Fig. 3. With considering ξ_{3t} to capture exogenous shocks in the consumption of oil, the inventory demand shock ξ_{4t} is best understood to characterize a precautionary (or speculative) demand for oil inventories. As Kilian and Murphy (2012) point out, various kinds of surprise information or upcoming expectations can be associated with such inventory demand shocks, including the potential of political unrest in oil-producing countries, discussions about the depletion of oil reserves (or peak oil effects), or the discovery of new oil fields and wells. According to the estimated model, the build up of inventory has only mild positive price effects during the first quarter after an inventory demand shock occurs. Subsequently the price effects turn (insignificantly) negative. While Baumeister and Hamilton (2019) have documented positive price effects of such a shock, it is worth noticing that a precautionary motive of holding oil might emerge under states of economic uncertainty and pessimistic views about business opportunities. In this case such inventory shocks would relax the crude oil market from price pressures.

3.3.2. Further discussion of the identified shocks

In summary and similar to results of Herwartz and Plödt (2016), the ICA-based identification approach yields unique shocks with sound economic underpinnings. The structural outcomes are fairly comparable with insights offered by several benchmark studies relying on more restrictive identification techniques such as recursive schemes, sign restrictions combined with elasticity bounds or informative priors (e.g., Kilian, 2009; Kilian and Murphy, 2012, 2014; Baumeister and Hamilton, 2019). Providing further justification for the credibility of the agnostically identified shocks, a few additional points are worth remarking.

First, all shocks identified in this study exhibit strong positive and significant correlations with their counterparts obtained from the Bayesian approach of Baumeister and Hamilton (2019), which relies on several informative priors and a shorter sample period ending in 2021M2. Specifically, exogenous shocks affecting oil supply, economic activity (or income), oil consumption and oil inventory as identified in this study exhibit linear correlations of 0.775, 0.176, 0.665 and 0.385, respectively, with posterior median shocks in Baumeister and Hamilton (2019).[4]

Second, structural outcomes comply with consensual narratives that have been suggested for the refinement of identified sets of oil market models that fulfill (weak) sign restrictions. Antolín-Díaz and Rubio-Ramírez (2018) have convincingly argued that aggregate demand shocks should be considered the weakest contributors to oil price movements during September and October 1980 (outbreak of the Iran–Iraq war) and August 1990 (Persian Gulf crisis). Based on historical decompositions for the real price of oil, we find that price surges during the outbreak of the Iran–Iraq war have been muted by a relaxation of oil inventories, while all shocks contributed positively to price changes associated with the second Gulf crisis. Among those shocks that contributed positively to oil price increases, the historical effects of aggregate demand shocks are the smallest throughout.[5]

Third, while KML estimation of structural parameters and shocks offers economically plausible insights into the functioning of oil markets, it is fair to say that – as discussed above – these insights are not necessarily novel. What renders statistical identification approaches especially useful in the present context is the provision of over-identifying information, such that otherwise just-identifying restrictions can be subjected to testing. For instance, with 10% significance, bootstrap-based Wald tests provided in Online Appendix A cast doubt on the imposition of a recursive structural model as suggested for a three-dimensional system by Kilian (2009).

Fourth, to address the robustness of core model implications, Online Appendix B documents complementary analysis that consist of four variants of the benchmark specification. Specifically, we consider (i) a model conditioned on pre-COVID data, (ii) a model utilizing the full sample data but adjusted for extreme realizations during the pandemic periods (Lenza and Primiceri, 2022), (iii) a model using the world industrial production index constructed by Baumeister and Hamilton (2019) as an alternative proxy for the global business cycle, and (iv) a model using the oil futures-spot spread as a financial forward-looking variable to replace inventory changes. Throughout, model-implied structural impulse response functions, and the structural decompositions of oil production confirm the key insights derived from the benchmark specifications.

Finally, the suggested approach fully identifies the four-dimensional system. Using high-frequency data on changes in oil price futures surrounding OPEC announcements, Känzig (2021) identifies an *oil supply news shock* that attributes the need for oil inventories to expected future oil shortfall. While the partially identified shock in Känzig (2021) naturally relates to expectations of future oil supply (and, consequently, the inventory demand shock), the oil supply shock examined in our fully identified system pertains to innovations in the supply curve, i.e., shocks to current oil production. Unreported results indicate that all identified shocks show significant impacts on oil futures-spot spreads, pointing to the fact that partially identified oil supply news shocks could be understood to capture mixed source signals which complicates their interpretation in a structural sense. Interestingly, all demand shocks identified by means of KML (and the lagged oil supply) show a moderate but significant correlation (501 observations) with the partially identified oil supply news shock. Positive correlations are between 0.099 and 0.1330, while the correlation with the lagged supply shock is −0.0837.

3.4. Oil supply and demand price responsiveness

The literature on the responsiveness of oil production with regard to changes in oil demand or supply comprises both micro- and macroeconomic approaches. Based on the recent review by Kilian (2022), one might conclude that the majority of structural assessments agrees on the limited responsiveness of the actual oil supply to oil price changes (see, e.g., Kilian and Murphy, 2014; Hafner et al., 2023; Herwartz and Plödt, 2016; Herrera and Rangaraju, 2020; Braun, 2021,

[4] We have used updated shocks provided by Christiane Baumeister for the period from 1976M2 until 2021M2 covering 553 observations, which were downloaded from https://sites.google.com/site/cjsbaumeister/datasets in June 2023. Very similar results are obtained for the original sample period of Baumeister and Hamilton (2019) ending in 2016M12. The smaller correlations documented for aggregate demand shocks are partly due to the fact that Baumeister and Hamilton (2019) condition their analysis on world industrial production as an alternative indicator of economic activity (see Online Appendix B for the respective robustness analysis). While the documented significant correlations point to the relevance of the insinuated approximation, cross correlations among distinct shocks identified by means of ICA and their Bayesian counterparts of Baumeister and Hamilton (2019) lack significance throughout.

[5] Historical decompositions of real oil prices yield the following contributions:

Sep. 1980				Oct. 1980				Aug. 1990			
ξ_{1t}	ξ_{2t}	ξ_{3t}	ξ_{4t}	ξ_{1t}	ξ_{2t}	ξ_{3t}	ξ_{4t}	ξ_{1t}	ξ_{2t}	ξ_{3t}	ξ_{4t}
1.54	0.86	1.02	−3.91	5.81	1.47	−1.37	−4.81	5.43	2.89	24.42	3.64

.

as exemplary studies concluding in favor of minor oil supply elasticities).[6] Contrary to the prevailing notion of, at most, a mild impact responsiveness of oil supply to exogenous price changes, Baumeister and Hamilton (2019) argue in favor of a considerably stronger sensitivity summarized in a posterior median price elasticity estimate of 0.15. Similarly, the weekly model of Valenti et al. (2023) points to a notable impact sensitivity of US producers to changes in the West Texas Intermediate crude oil price. Although Baumeister and Hamilton (2019) conduct numerous robustness exercises to support their key conclusion on the price sensitivity of oil supply, their Bayesian analysis throughout relies on the assumption of a Gaussian likelihood. In this regard, the analysis in Braun (2021) provides an insightful additional robustness analysis. While the benchmark results of Baumeister and Hamilton (2019) hold under the assumption of a Gaussian likelihood, substituting the Gaussian density with a flexible non-Gaussian model yields price elasticities of oil supply close to zero. To elucidate this outcome, Braun (2021) highlights that the specification of informative priors strongly shapes posterior outcomes, when sample information indicates only a weak correlation between changes in oil prices and oil production (which can also be implicitly inferred from time series displayed in Fig. 2). Given strong indications of non-Gaussian structural shocks that has been documented in the literature (e.g., Herwartz and Plödt, 2016; Hafner et al., 2023) and detected in the present sample, we recall that the majority of empirical evidence points to a minor price sensitivity of oil supply. Moreover, we view the ongoing debate in the literature as a crucial motivation to consider outcomes of fully agnostic identification schemes as valuable complements to more informed approaches in global oil market studies.

Among alternative representations, the structural model in (3) holds particular merit, especially when the focus is on understanding causal relationships among the observed variables. In Online Appendix A, we provide detailed estimates for the matrix Π_0, along with bootstrap means and t-ratios. Notably, the first and third equation of the system in (3) represent the oil supply and oil demand equations, respectively, from which short- and long-run price elasticities can be calculated (see, e.g., Baumeister and Hamilton, 2019). With $\pi_{ij}^{(l)}$ denoting the ijth element of Π_l, the short-run price elasticity of oil supply and oil demand can be computed as $-\pi_{13}^{(0)}/\pi_{11}^{(0)}$ and $-\pi_{33}^{(0)}/\pi_{31}^{(0)}$, respectively.[7] Due to the stationarity condition, the structural VAR model enables us to estimate not only the short-run (one-month) demand and supply curves, but also the slopes of the long-run stationary distribution. As a result, the unconditional mean of y_t exists and can be considered as the long-run predictor with minimum mean-squared error loss, i.e., $\lim_{h\to\infty} y_{T,h} \equiv \lim_{h\to\infty} \mathbb{E}[y_{T+h}|\Omega_T]$, where Ω_T denotes the filtration associated with y_t. The first equation of the system in (3) allows us to derive the long-run price elasticity of oil supply as the limit $\lim_{h\to\infty} \frac{\partial \mathbb{E}[Q_{T+h}-Q_T|\Omega_T]}{\partial \mathbb{E}[P_{T+h}-P_T|\Omega_T]}$, with the changes in quantity and price being caused by a unit oil supply shock. Corresponding arrangements of the third equation allow for the construction of the long-run price elasticity of oil demand. To derive explicit representations of these long-run elasticities, it is useful to consult the so-called autoregressive distributed lag model implied by the system in (3) and replace all variables by their unconditional mean. Then, the long-run price elasticities of oil supply and oil demand are given by $(-\pi_{13}^{(0)}+\sum_{l=1}^{p}\pi_{13}^{(l)})/(\pi_{11}^{(0)}-\sum_{l=1}^{p}\pi_{11}^{(l)})$ and $(-\pi_{33}^{(0)}+\sum_{l=1}^{p}\pi_{33}^{(l)})/(\pi_{31}^{(0)}-\sum_{l=1}^{p}\pi_{31}^{(l)})$, respectively.

With regard to the marginal responsiveness of oil production to demand and supply shifts caused by the aforementioned shocks, the SVAR model yields the following conclusions. In response to a 10% decrease in oil prices, oil producers will cut oil production by 0.2% on impact. In the long-run, we do not observe an increase in the supply elasticity as our estimate remains approximately the same. In contrast, in response to the same 10% price drop, oil demand will increase by 13.6% on impact and by 4.7% in the long-term. According to corresponding confidence intervals with 90% coverage, all these elasticity estimates are characterized by sufficient precision.[8]

Our findings align with the existing literature, which suggests that oil supply is relatively inelastic in the short-run, while oil demand is much more elastic (Kilian, 2022). As a result, aggregate and oil-specific demand shocks play a dominant role in driving the real price of oil, whereas oil supply shocks are crucial in determining the level of oil production.

3.5. Forecast error variance decompositions

As a complement to structural IRFs, forecast error variance decompositions (FEVDs) have become a common tool to assess the relative importance of structural shocks in explaining system variation for a specific variable. In the context of the global oil market model, Fig. 4 shows implied decompositions of the uncertainty associated with h-step ahead forecasts of changes in oil production. The alternative forecast horizons considered include impact effects, as well as mid-term (6 months to 1 year) and long-term (5 to 10 years) perspectives. The results show that oil supply shocks are the major determinant of production, explaining up to 96% and about 80% of model-implied forecast uncertainty in the short- and long-run, respectively. Other shocks contribute only marginally to predictive uncertainty, with the aggregate demand shock having the weakest contribution at all horizons. These FEVDs suggest that oil supply shocks are the main causal factor driving oil production.

3.6. Historical decompositions of oil production

The structural implications of the model discussed so far, including IRFs, elasticity estimates, and FEVDs, provide a characterization of the structural model in an unconditional manner, allowing for stylized insights to be considered to hold ‘on average’. However, to describe the structural model implications with finer time resolution, historical decompositions can be used to trace actual data quotes back to the underlying structural shocks. While historical decompositions have been very beneficial in the literature on global oil markets to understand time-varying structural patterns of oil pricing, to the best of our knowledge, historical decompositions of crude oil *production* have not been discussed yet.

Fig. 5 presents historical decompositions of the stochastic component of global oil production. Overall, the results suggest that the observed oil production can be largely attributed to the procession of underlying supply shocks. In particular, we observe that the contribution of aggregate demand shocks to shaping oil production patterns is very limited throughout. Starting after the early millennium economic slowdown until the outbreak of the Great Financial Crisis, there were mild positive contributions of aggregate demand shocks to oil production. Subsequent to the European sovereign debt crisis and from 2013 to 2017 mildly negative effects imply that aggregate demand effects have muted oil production to a minor degree. For an extended period

[6] Kilian (2022) concludes that ‘the one-month oil supply elasticity is low, which implies that oil demand shocks are the dominant driver of the real price of oil’, while ‘recent findings of rather large one-month oil supply elasticities are misleading’. In fact, many studies document oil supply elasticity estimates that are well in line with the upper bound suggested in Kilian and Murphy (2014) of 0.0258.

[7] Unlike stylized regressions, the set of equations in (3) lacks normalization, i.e., there is no (left-hand-side) variable with a coefficient of unity. In Online Appendix A, we also document a normalized version of the matrix Π_0 that allows for a more direct recovery of the defined short-run elasticities.

[8] Specifically, the marginal price elasticities of oil supply and demand are with 90% probability covered by the interval estimates of $[0.0179; 0.0218]$ (short-run oil supply elasticity), $[-1.9721; -1.0614]$ (short-run oil demand), $[0.0217; 0.0284]$ (long-run oil supply) and $[-0.8727; -0.2708]$ (long-run oil demand).

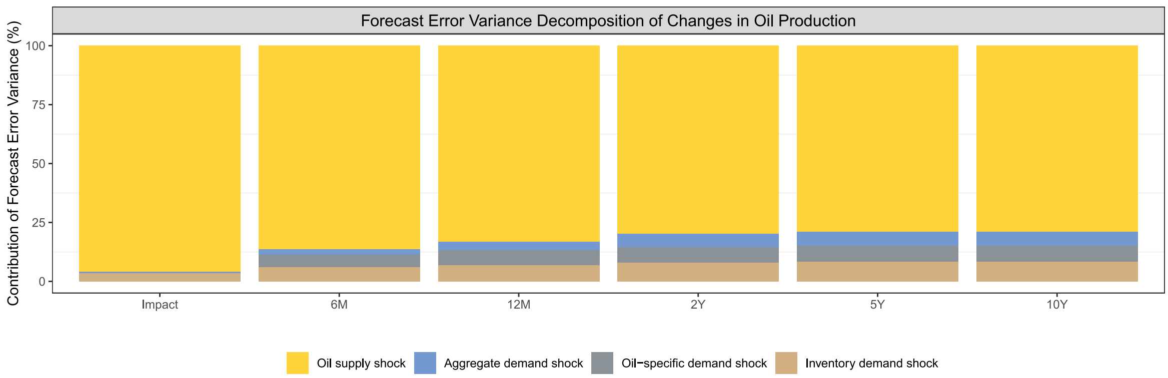

Fig. 4. FEDVs for changes in oil production on impact and at short, medium and long horizons.

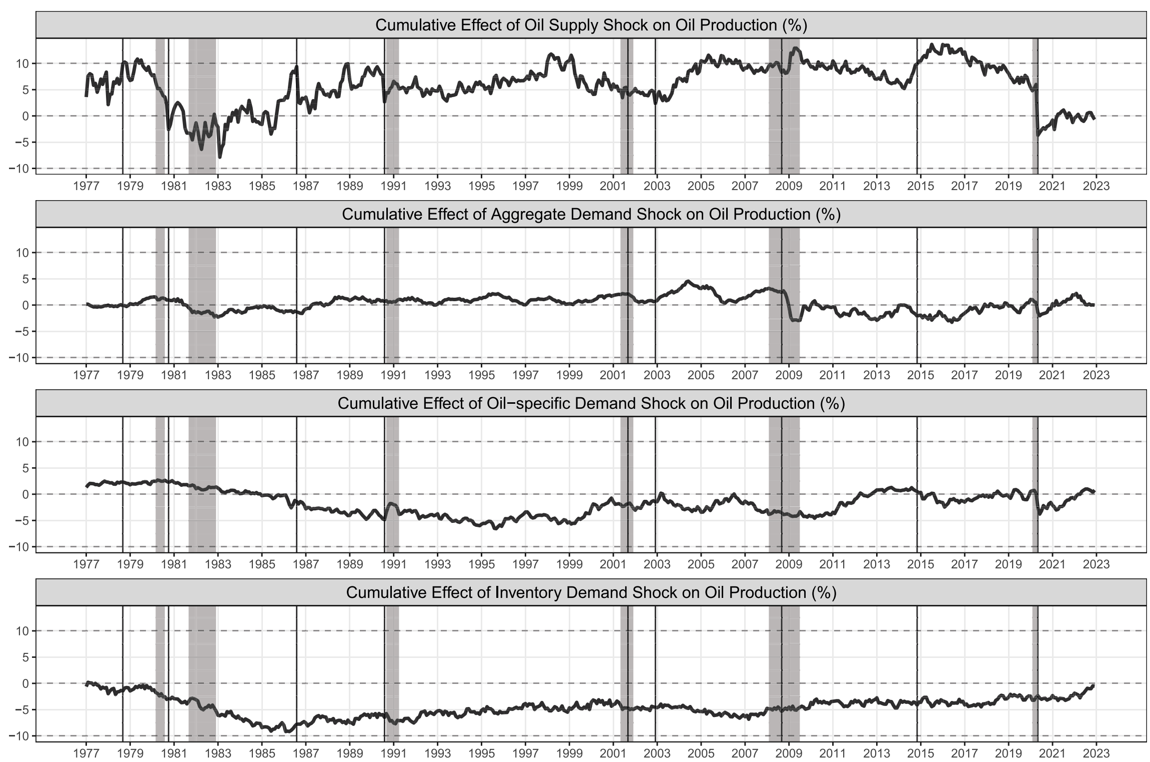

Fig. 5. Historical decompositions of oil production. Vertical lines indicate major events in the global oil market, including the outbreak of the Iranian Revolution (1978M9), the Iran–Iraq War (1980M9), the agreement of OPEC in Geneva to restrict production (1986M8), the Persian Gulf War (1990M8), the 9/11 terrorist attack (2001M9), the Venezuelan crisis (2002M12), the collapse of Lehman Brothers (2008M9), the unexpected block of OPEC's agreement on production cut by Saudi Arabia (2014M11), and the unprecedented production cut by OPEC+ (2020M5). For further notes see Fig. 1.

starting after the Geneva agreement of OPEC and ending in 2012, oil-specific demand shocks have slightly but consistently contributed to a moderation of oil production. Similarly, oil-specific demand shocks have contributed to the reduction of oil production in the context of the measures taken to counter the outbreak of the SARS-CoV-2 virus early in 2020. A noteworthy finding is that the contribution of inventory demand shocks to oil production has been negative throughout the sample period.

4. Structural oil market shocks, CO2 emissions and climate policy implications

The previous analysis indicates that: (i) aggregate, oil-specific, or inventory demand shocks have only mild (long-run) effects on oil production; (ii) the short- and long-run oil supply curve is rather inelastic such that demand-shifts cause oil price fluctuations but hardly impact on oil production; and (iii) the historical decomposition of oil production reveals a limited role of aggregate demand shocks for oil production in the past. These findings suggest that current policies aiming to address climate change by a reduction in the demand of fossil fuel-based energy might only have a limited impact on fossil fuel production and, thus, ultimately on climate change. To further assess this issue in the context of the global oil market, it is necessary to unravel the relationship between GHG emissions and oil production, and, finally, the effects of the shocks identified above on GHG emissions.

From eyeballing the time series of oil production and CO2 emissions as displayed in Fig. 1, it is evident – and confirmed by explicit cointegration testing (see notes to Fig. 1) – that the former can be considered a key determinant of the latter. However, it is important to notice that oil production is actually the result of underlying shocks, which are specific to the supply side of the oil market on the one hand, and heterogeneous motives for demanding oil on the other hand. With this in mind and to enable more focused policy strategies to achieve reductions in GHG emissions, it is instructive to trace CO2 emissions,

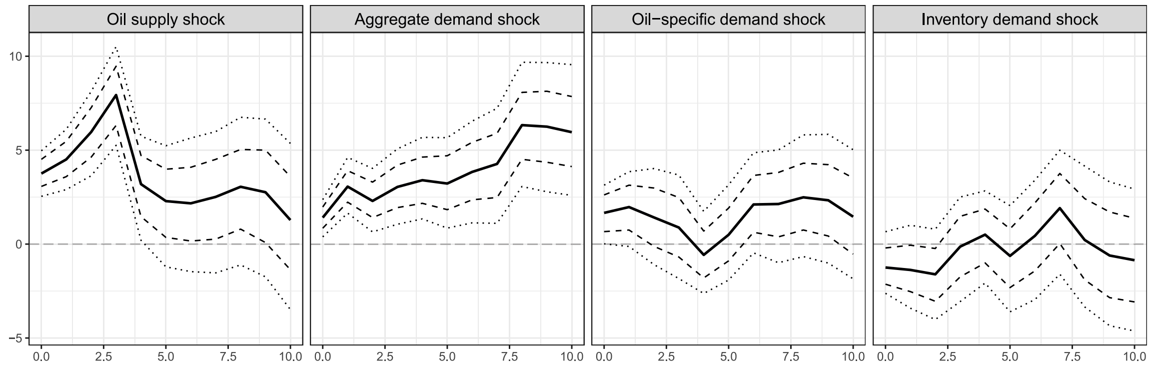

Fig. 6. Local projection results. Responses of world CO2 emissions (in percent, i.e. cumulated estimates $\hat{\beta}_l$ from (8)) to structural shocks. Panels show medians (solid curves) and confidence bonds with 68% (dashed) and 90% pointwise coverage (dotted) from 1000 replications of a moving block bootstrap scheme.

as displayed in Fig. 1, back to exogenous shocks hitting the crude oil market.

To unravel the role of the identified structural oil market shocks for the actual GHG emissions, it appears straightforward to augment the four-dimensional VAR system with time series information on changes of (log) CO2 emissions. Unfortunately, global emission data are only available at the annual frequency. As a feasible alternative to a full-information VAR system, we employ local projection IRFs in the vein of Jorda (2005) to unravel structural sources of GHG emissions. In comparison with a full-information analysis utilizing a five-dimensional SVAR model, limited-information methods like LPs have been shown robust to misspecification. Specifically, define $x_\tau = (g_\tau, y_\tau^{*'})'$, where g_τ is the annual log change of the world CO2 emissions ($\times 100$) in year τ and y_τ^* is the vector consisting of VAR variables, which are aggregated to annual data by averaging realizations over all months during year τ. Similarly, let $\xi_{i\tau}^*$ denote averaged structural shocks of year τ. Then, local projection impulse responses of g_τ to shocks $\xi_{i\tau}^*$ after l years with $l \in \mathbb{N}$ and $i \in \{1, \ldots, K\}$ are estimated from the model

$$g_{\tau+l} = \beta_0 + \beta_l \xi_{i\tau}^* + \sum_{j=1}^{p^*} \Gamma_j' x_{\tau-j} + \zeta_\tau, \tag{8}$$

where the lag order is chosen in accordance with the monthly VAR, i.e., $p^* = 2$ (years). We consider in total eleven such regressions with $l = 0, 1, \ldots, 10$. The resulting local projection IRFs of CO2 emissions are displayed in Fig. 6.

LPs reveal that in comparison with all other structural determinants of oil production, oil supply shocks have a predominant impact on CO2 emissions, particularly in the short-run (i.e., within the first two years after their occurrence). Among the demand shocks, the oil-specific and inventory demand shocks lack significant impacts on CO2 emissions, in the short-, mid- and long-run. Instead, an aggregate demand shock manifests in CO2 emissions, exhibiting particular strength in the medium- to longer-term (i.e, from five to ten years).

Regarding climate policy implications, these results emphasize the critical significance of the scope of demand-based policies as exemplified in Fig. 7. For instance, if a cap on emissions implemented through an Emissions Trading System (ETS) in a group of countries is understood as a negative oil-specific demand shock, this policy would have only negligible impacts on oil production and CO2 emissions (see Fig. 7(b)). In contrast, a global cap on emissions, corresponding to an aggregate demand shock, would prove effective in reducing CO2 emissions in the mid- and long-run (see Fig. 7(c)). With these arguments in mind, it becomes evident why the European Union's considerable endeavors to transition from fossil fuels to cleaner energy sources have resulted in a mere 1% decrease in oil consumption from 2011 to 2021 (BP, 2021).

To investigate how far-reaching climate demand-side policies must be in order to be effective in the global oil market is beyond the scope of this study, and will depend on the extent to which oil supply remains inelastic (for an exemplary case of eventually changing transmission patterns of oil market shocks to the US economy see Bruns and Lütkepohl, 2023). In particular, it could be argued that the oil supply will be more elastic in the case of a demand shift that exceeds the current reduction invoked by the substitution of fossil fuels with renewable energy sources in developed economies. However, recent evidence puts this argument into question. According to Masnadi et al. (2021), the national volume-weighted average marginal production costs for crude oil per barrel ranged from as low as US$2.8 in Iraq to as high as US$21.5 in Colombia. Even during the COVID-19 pandemic demand shock in April 2020, when the oil price dropped to around US$40 per barrel, it still remained significantly above the marginal cost of production. Therefore, it seems unlikely that demand-driven reductions in the oil price will make crude oil production unprofitable on a large scale in the near future. Moreover, it is also unlikely that a reduction in oil supply will occur in the near future for two reasons. First, the world's total proven oil reserves at the end of 2020 were 1732.4 thousand million barrels, which, at current production rates, would last for 53.5 years (BP, 2021). Second, proven oil reserves have been rising by 33.4% over the last 20 years, making it unlikely to achieve the climate goals set by the UNFCCC for 2030 and 2050 as a result of the exhaustion of oil resources.

In summary, our results provide an explanation for why current policies, which aim to reduce oil demand by promoting the substitution of fossil fuels with renewable energy sources, have been largely ineffective in reducing global oil production and, consequently, GHG (in particular CO2) emissions. Specifically, demand-side policies, supported by only a limited number of countries, may not be sufficient to achieve the 1.5 °C objective set for 2050 by the Paris Agreement of 2015. Indeed, it was recognized soon after its ratification that all major industrialized countries were failing to meet the pledges to cut GHG emissions made in the Paris Agreement (Victor et al., 2017). In the same vein, Boehm et al. (2022) point out that in 2021 'global GHG emissions are higher than they were when more than 190 parties adopted the Paris Agreement in 2015'. Accordingly, these authors conclude that 'much greater ambition and action is urgently needed' to meet the agreement's objective of limiting global warming to 1.5 °C for 2050. This failure can be attributed to the inherent limitations of unilateral demand-based policies in fossil fuel markets due to their demand and supply characteristics. In this context, the adoption of a global cap-and-trade system appears to be the most viable policy alternative for reducing oil production and GHG emissions worldwide.

5. Conclusions

This study is the first to present an explicit view of the structural determinants of oil production across the short-, medium-, and long-term. The analysis relies on a structural vector autoregressive model utilizing a purely agnostic and ICA-based identification approach. The

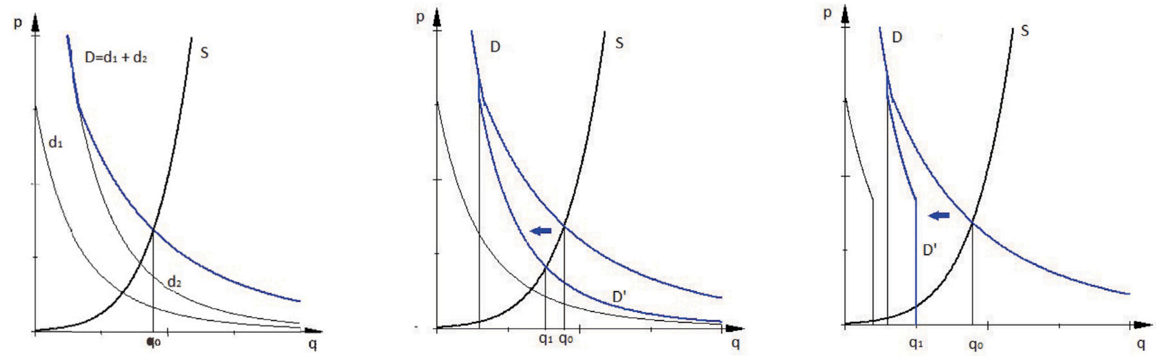

Fig. 7. The effect of jurisdiction-specific and global caps on consumption with an inelastic supply curve. Panel (a) *Market equilibrium with aggregated demand from two jurisdictions* displays the market equilibrium quantity (q_0) resulting from the interaction between the supply (S) and aggregated demand (D) from two jurisdictions (d_1 and d_2). Panel (b) *Market equilibrium with a jurisdiction-specific cap on consumption* indicates how the market equilibrium values change after imposing a strict consumption cap in jurisdiction 1, shifting aggregated demand from D to D'. Panel (c) *Market equilibrium with a global cap on consumption* shows the change in equilibrium quantity after imposing a cap in both jurisdictions.

results, obtained from impulse response functions, estimated oil supply and demand elasticities, the decomposition of forecast error variances, and the historical decomposition of oil production, all indicate that, over the past five decades, changes in crude oil demand have had only minor impacts on the actual level of oil production. Local projections of global CO2 emissions on annualized oil market shocks reveal that only supply and aggregate demand shocks have impacted on CO2 emissions, while oil-specific demand shocks did not exert any significant influence.

An important implication of this result is that unilateral demand-side strategies in energy and climate protection policies will experience limited effectiveness, as observed in past policies adopted by a limited number of countries. This underscores the need for a global approach, involving the implementation of a worldwide cap on GHG emissions, as outlined in international climate summits such as the 2015 Paris Agreement or the COP28 UN Climate Change Conference in Dubai. The necessity of establishing a global cap is further underscored by the fact that other alternatives may take too much time to be effective in reaching the climate goals for 2050. One such alternative involves focusing on technological innovation to make alternative green fuels more cost-competitive than oil in the global market. However, this approach requires considerable investments in research and development, and the process may take time. Another promising alternative consists of implementing carbon capture technology, which involves capturing CO2 emissions from power plants and other sources and storing them underground. Nonetheless, this technology requires further research and development to be scaled up and become cost-effective (Hepburn et al., 2019).

While the focus of this study has been on the crude oil market, for future research it would be interesting to determine whether these insights also apply to other fossil fuel markets, such as coal and natural gas.

CRediT authorship contribution statement

Helmut Herwartz: Conceptualization, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Bernd Theilen:** Investigation, Writing – original draft, Writing – review & editing. **Shu Wang:** Data curation, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial and non-financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft, Germany (HE 2188/17-1), the Ministerio de Ciencia e Innovación, Spain (PID2019-105982GB-I00 and PID2022-137382NB-I00) and the Generalitat de Catalunya, Spain (SGR2021-00729). We are grateful to Eugenio Figueroa, Helmut Lütkepohl, Werner Roeger and Hans-Werner Sinn for insightful discussions. Excellent research assistance from Huynh Tuan Duy Bui is gratefully acknowledged. We also acknowledge constructive comments that we received from two anonymous reviewers.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.eneco.2024.107488.

References

Antolín-Díaz, J., Rubio-Ramírez, J.F., 2018. Narrative sign restrictions for SVARs. Amer. Econ. Rev. 108 (10), 2802–2829.

Baumeister, C., Hamilton, J.D., 2019. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. Amer. Econ. Rev. 109 (5), 1873–1910.

Bhattacharya, P.K., 1967. Estimation of a probability density function and its derivatives. Sankhyā: Indian J. Statist. A (1961-2002) 29 (4), 373–382.

Blanchard, O.J., Quah, D., 1989. The dynamic effects of aggregate demand and supply disturbances. Amer. Econ. Rev. 79 (4), 655–673.

Boehm, S., Jeffery, L., Levin, K., Hecke, J., Schumer, C., Fyson, C., Majid, A., Jaeger, J., Nilsson, A., Naimoli, S., Thwaites, J., Cassidy, E., Lebling, K., Sims, M., Waite, R., Wilson, R., Castellanos, S., Singh, N., Lee, A., Geiges, A., 2022. State of Climate Action 2022. Tech. rep, Bezos Earth Fund, Climate Action Tracker, Climate Analytics, ClimateWorks Foundation, NewClimate Institute, the United Nations Climate Change High-Level Champions, and World Resources Institute, Berlin and Cologne, Germany, San Francisco, CA, and Washington, DC.

BP, 2021. Statistical review of world energy. Available at:http://www.bp.com/statisticalreview, London: BP Statistical Review of World Energy.

Braun, R., 2021. The importance of supply and demand for oil prices: evidence from non-Gaussianity. Bank of England working papers 957, Bank of England.

Brüggemann, R., Jentsch, C., Trenkler, C., 2016. Inference in VARs with conditional heteroskedasticity of unknown form. J. Econometrics 191 (1), 69–85.

Bruns, M., Lütkepohl, H., 2023. Have the Effects of Shocks to Oil Price Expectations Changed?: Evidence from Heteroskedastic Proxy Vector Autoregressions. Tech. rep.

Comon, P., 1994. Independent component analysis, a new concept? Signal Process. 36 (3), 287–314.

Enerdata, 2023. World energy & climate statistics - yearbook 2022: CO2 emissions from fuel combustion. URL https://yearbook.enerdata.net/co2/emissions-co2-data-from-fuel-combustion.html.

Faust, J., 1998. The robustness of identified VAR conclusions about money. Carnegie-Rochester Conf. Series Public Policy 49 (1), 207–244.

Fiorentini, G., Sentana, E., 2023. Discrete mixtures of normals pseudo maximum likelihood estimators of structural vector autoregressions. J. Econometrics 235 (2), 643–665.
Genest, C., Quessy, J.-F., Rémillard, B., 2007. Asymptotic local efficiency of cramér von mises tests for multivariate independence. Ann. Statist. 35 (1), 166–191.
Gouriéroux, C., Monfort, A., Renne, J., 2017. Statistical inference for independent component analysis: Application to structural VAR models. J. Econometrics 196, 111–126.
Hafner, C.M., Herwartz, H., 2023. Dynamic score-driven independent component analysis. J. Bus. Econom. Statist. 41, 298–308.
Hafner, C.M., Herwartz, H., Wang, S., 2023. Causal inference with (partially) independent shocks and structural signals on the global crude oil market. Discussion Paper ISBA 2023/4, LIDAM.
Hepburn, C., Adlen, E., Beddington, J., Carter, E.A., Fuss, S., Mac Dowell, N., Minx, J.C., Smith, P., Williams, C., Yang, J., et al., 2019. The technological and economic prospects for CO2 utilization and removal. Nature 575 (7728), 87–97.
Herrera, A.M., Rangaraju, S.K., 2020. The effect of oil supply shocks on US economic activity: What have we learned? J. Appl. Econometrics 35 (2), 141–159.
Herwartz, H., 2018. Hodges lehmann detection of structural shocks - an analysis of macroeconomic dynamics in the euro area. Oxford Bull. Econ. Statist. 80 (4), 736–754.
Herwartz, H., Plödt, M., 2016. The macroeconomic effects of oil price shocks: Evidence from a statistical identification approach. J. Int. Money Finance 61 (C), 30–44.
Herwartz, H., Wang, S., 2023. Point estimation in sign-restricted SVARs based on independence criteria with an application to rational bubbles. J. Econom. Dynam. Control 151, 104630.
Jorda, O., 2005. Estimation and inference of impulse responses by local projections. Amer. Econ. Rev. 95 (1), 161–182.
Känzig, D.R., 2021. The macroeconomic effects of oil supply news: Evidence from opec announcements. Amer. Econ. Rev. 111 (4), 1092–1125.
Keweloh, S.A., 2021. A generalized method of moments estimator for structural vector autoregressions based on higher moments. J. Bus. Econom. Statist. 39 (3), 772–782.
Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. Amer. Econ. Rev. 99 (3), 1053–1069.
Kilian, L., 2022. Understanding the estimation of oil demand and oil supply elasticities. Energy Econ. 107, 105844.
Kilian, L., Lütkepohl, H., 2017. Structural Vector Autoregressive Analysis. Cambridge University Press.
Kilian, L., Murphy, D.P., 2012. Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models. J. Eur. Econom. Assoc. 10 (5), 1166–1188.
Kilian, L., Murphy, D.P., 2014. The role of inventories and speculative trading in the global market for crude oil. J. Appl. Econometrics 29 (3), 454–478.
Kilian, L., Zhou, X., 2018. Modeling fluctuations in the global demand for commodities. J. Int. Money Finance 88, 54–78.
Kilian, L., Zhou, X., 2023. The econometrics of oil market VAR models. Adv. Econom. 45B, 65–95.
Lanne, M., Luoto, J., 2021. GMM estimation of non-Gaussian structural vector autoregression. J. Bus. Econom. Statist. 39 (1), 69–81.
Lanne, M., Meitz, M., Saikkonen, P., 2017. Identification and estimation of non-Gaussian structural vector autoregressions. J. Econometrics 196, 288–304.
Lenza, M., Primiceri, G.E., 2022. How to estimate a vector autoregression after march 2020. J. Appl. Econometrics 37 (4), 688–699.
Lütkepohl, H., Netšunajev, A., 2014. Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural VARs. J. Appl. Econometrics 29 (3), 479–496.
Lütkepohl, H., Netšunajev, A., 2017. Structural vector autoregressions with heteroskedasticity: A review of different volatility models. Econometr. Statist. 1, 2–18.
Magnusson, L.M., Mavroeidis, S., 2014. Identification using stability restrictions. Econometrica 82 (5), 1799–1851.
Masnadi, M., Benini, G., El-Houjeiri, H., Milivinti, A., Anderson, J., Wallington, T., De Kleine, R., Dotti, V., Jochem, P., Brandt, A., 2021. Carbon implications of marginal oils from market-derived demand shocks. Nature 599, 80–84.
Silverman, B.W., 1998. Density Estimation for Statistics and Data Analysis. Chapman and Hall/CRC.
Sims, C.A., 1980. Macroeconomics and reality. Econometrica 48 (1), 1–48.
Székely, G.J., Rizzo, M.L., Bakirov, N.K., 2007. Measuring and testing dependence by correlation of distances. Ann. Statist. 35 (6), 2769–2794.
Uhlig, H., 2005. What are the effects of monetary policy on output? Results from an agnostic identification procedure. J. Monetary Econ. 52 (2), 381–419.
Valenti, D., Bastianin, A., Manera, M., 2023. A weekly structural VAR model of the US crude oil market. Energy Econ. 121, 106656.
Victor, D.G., Akimoto, K., Kaya, Y., Yamaguchi, M., Cullenward, D., Hepburn, C., Taschetti, R., 2017. Prove Paris was more than paper promises. Nature 548 (7665), 25–27.