

Oil markets and global economic conditions: Evidence from a general equilibrium model *

Romain Houssa[†] Jolan Mohimont[‡]

February 22, 2024

Abstract

In this paper, we study the role of the elasticities of oil supply and demand on the dynamics of oil price within a simple RBC model. Our framework is flexible to estimate both the short- and long-run price-elasticities of oil supply, and the price and income-elasticities of oil demand. We show that imposing prior restrictions on the short-run elasticity of oil supply has repercussions on the other elasticities. While the SVAR literature focuses on the short-term elasticity, we show that the long-term elasticity of supply plays a key role in explaining the disagreement found in that literature over the relative importance of oil demand and supply shocks in driving oil price.

JEL classifications: E3, E43, E52, C51, C33

Keywords: Oil market, RBC, Bayesian estimation, supply elasticities.

*Disclaimer: this paper does not necessarily reflect the views of the National Bank of Belgium

[†]DeFiPP (CRED & CeReFiM) - University of Namur; CES (University of Leuven), and CESifo, romain.houssa@unamur.be

[‡]National Bank of Belgium; DeFiPP (CRED & CeReFiM) - University of Namur, jolan.mohimont@nbb.be

1 Introduction

What are the causes and consequence of oil price fluctuations? This question is of key importance for academics, practitioners, and policy makers alike. From the oil shocks in the 1970s to the 2000s commodity boom, oil prices showed large fluctuations, with important repercussions on firms, households, and governments in importing and exporting countries.

A vast body of literature has used structural vector autoregressive (SVAR) models to identify the sources in oil price fluctuation but failed to reach a consensus on the relative contribution of different oil demand and supply shocks. As a result, assessing the driving forces of commodity prices is still a subject of an intense debate. For instance, [Kilian \(2009\)](#) and [Kilian and Murphy \(2012, 2014\)](#) attribute most of oil prices fluctuations to oil demand shocks.¹ On the contrary, [Baumeister and Hamilton \(2019\)](#) and [Caldara et al. \(2019\)](#) find a dominant role for oil supply shocks.

This disagreement finds its roots in the different identification strategies that are used in SVAR models ([Herrera and Rangaraju, 2020](#)). While [Kilian \(2009\)](#) (henceforth K09) and [Kilian and Murphy \(2012, 2014\)](#) (KM12 and KM14) assume that the short-run price-elasticity of oil supply is zero or extremely low, [Baumeister and Hamilton \(2019\)](#) (BH19) and [Caldara et al. \(2019\)](#) (CCI19) argue that this elasticity might be larger, and that one could exploit both prior knowledge of oil supply and demand elasticities to refine the identification strategy. Overall, one learns from these studies that imposing a lower short-run price-elasticity of oil supply is associated with a lower contribution of oil supply shocks to oil prices. In particular, the value of the short-run elasticities ranges from 0-0.01 in K09 and KM12 to 0.10-0.14 in CCI19 and BH19. While the exact value of this elasticity is at the center of the debate in the SVAR literature, little is known about why seemingly small changes in this elasticity – that only directly constrain the response of oil production in the first month that follows an unexpected shift in oil demand – translate into radically different conclusion regarding the sources of oil price fluctuations.

This paper contributes to this debate through the lenses of a RBC model, allowing to explain why different assumptions about the short-run price-elasticity of oil supply have different implications for the drivers of oil prices. The model explains global oil demand and supply. It comprises households, final good and oil producers, and storage firms that hold oil inventories. In this framework, four key elasticities are easily pinned down by a small number of structural parameters: the short and long-run price elasticities of oil supply and the price

¹ Their results were later confirmed by subsequent studies, see for e.g., [Juvenal and Petrella \(2015\)](#); [Antolín-Díaz and Rubio-Ramírez \(2018\)](#); [Zhou \(2020\)](#); [Cross et al. \(2022\)](#); [Braun \(2023\)](#).

and income elasticities of oil demand. We estimate our model with Bayesian methods using data from 1975M1 to 2019M12. We build our priors on the elasticities using information from the literature (see for e.g., BH19, CCI19, [Greene and Leiby, 2006](#); [Hamilton, 2009](#); [Chang et al., 2019](#); [Havranek and Kokes, 2015](#); [Golombek et al., 2018](#); [Taghizadeh Hesary and Yoshino, 2014](#), and [Kilian, 2022](#)). Our model also incorporates the four shocks traditionally considered in the SVAR literature: a global economic activity shocks and three shocks that are specific to the oil market, which capture exogenous shifts in oil-consumption demand, in oil-inventory demand, and in oil supply.

We show that varying the impact elasticity of oil supply (from 0 to 0.14) while leaving all other parameters of our model unchanged has a limited direct effect on the dynamics of oil prices. For instance, imposing a perfectly inelastic oil supply in the very short run amplifies the effect of a global economic activity shock on oil prices, but this amplification is confined to the short run and mitigated by the response of inventories that partially offsets the short-run imbalances between oil supply and demand. So, why do different values for this short-run oil supply elasticity translate into radically different dynamics in oil prices as reported in the studies discussed above?

We conjecture that imposing a lower short-run price-elasticity of oil supply affects the estimates of the other elasticities in the model. Our conjecture is inspired by BH19, CCI19 and [Herrera and Rangaraju \(2020\)](#), who noted that imposing a lower short-run price-elasticity of oil supply results in a higher estimated price-elasticity of oil demand (and vice-versa). In our model, a higher price-elasticity of demand mitigates the effect oil supply shocks as producers or consumer can substitute oil with other products. This explanation thus has the potential to shed light on the disagreement found in the SVAR literature.

Moreover, we uncover a new link between the short and long-run elasticity of oil supply. We show that imposing a lower price-elasticity of oil supply also reduces the estimate of this elasticity in the long-run. In our baseline estimation, the mode of the short and long-run elasticities of oil supply are 0.08 and 0.16, respectively. Imposing a zero elasticity on impact reduces the value of this elasticity to 0.09 in the long run, which, in turns, have a large repercussion on the forecast error variance decomposition (FEVD) of oil price at all horizons. Indeed, a joint reduction in the short and long-run elasticities implies that the effect of oil demand shocks are magnified and more persistent. So, while the SVAR debate concentrated on the value of the short-run oil elasticity of supply, our analysis shows that restrictions imposed on this elasticity may also affect its value in the long run. In fact, by reproducing the results of KM14 and BH19, one can note that their medium-run elasticities of oil supply also substantially differ. Furthermore, we show that the short-run elasticity of

oil supply is negatively correlated with the income elasticity of demand in our model.

We use our framework to revisit the question of the source of oil price fluctuations. An advantage of our model is that it combines different identifying restrictions used in the SVAR literature with prior studies of the oil market in a statistically and structurally coherent framework. While BH19 proposed this idea in a SVAR model, our approach relies on a RBC model. Moreover, we exploit prior knowledge on the long-run price-elasticity of oil supply. The ability of our model to produce a long-run elasticity of oil supply as a function of a few structural parameters is a clear advantage over SVAR, where this elasticity depends on a complex function of many parameters. The SVAR literature has also used a large variety of sign and magnitude restrictions on the impulse response functions (in addition to the priors or restrictions imposed on the elasticities of oil demand and supply).² In our model, these well-grounded sign and magnitude restrictions are all jointly satisfied.

Our results confirm that oil demand shocks are the most important drivers of oil prices, but we also find that supply shocks do matter. Global economic activity and oil-specific consumption demand shocks both explain about two-fifth of the variance in oil prices in the 1975-2019 period. We attribute the oil price boom in the early 2000s to the cumulative effect of aggregate demand shock, confirming the results of K09, KM12 and KM14. Nevertheless, we also find that oil supply shocks play a larger role than in K09, albeit not as large as in BH19 and CCI19. In our baseline experiment, oil supply shocks explain about 20% of the variance in oil prices. Our historical decomposition reveals that weaker-than-expected supply conditions contributed to maintaining oil prices high in the early 80's and after the global financial crisis.

We also demonstrate that the prior imposed on the long-run price elasticity of oil supply has a strong impact on the dynamics of oil prices. We perform two experiments where we sequentially impose extremely small or large elasticities. When we constraint the value of the long-run elasticity to be equal to its value in the short-run, we find that the contribution of supply shocks to the FEVD of oil prices fall to below 10%. This result is very close to that of KM12, who estimate a small elasticity in the medium run. In contrast, this FEVD becomes higher than 50% when oil supply is assumed to be perfectly elastic in the long run. This result bears important repercussion for the design and calibration of models of the oil markets (or commodities in general). A good model should be capable of producing different elasticities at different horizons. Since [Kose \(2002\)](#), it has become common to

² Sign restrictions are discussed in [Lippi and Nobili \(2012\)](#); [Baumeister et al. \(2010\)](#); [Baumeister and Peersman \(2013b,a\)](#); [Peersman and van Robays \(2009\)](#); [Lütkepohl and Netšunajev \(2014\)](#), while magnitude restrictions are additionally introduced in KM12.

model commodity supply with a combination of labor, capital, and land (a fixed production factor). We show how differences in the factor share and in the elasticity of substitution between these factors define the short and long-run elasticities of supply and translate into different dynamics for commodity prices.

We also show that, despite our efforts to combine the available prior knowledge of the oil market with popular identifying restrictions, substantial uncertainty remains over what are the main drivers of the oil market. Our elasticities are well-identified according to the results of the tests designed by [Andrle \(2010\)](#); [Iskrev \(2010\)](#), but as we argue in the first part of our paper, small changes in these elasticities have large repercussions on oil prices dynamics. In addition, the posterior multivariate distribution of the elasticities is such that the parameter draws displaying a lower supply elasticity in the short run tend to be associated with a lower supply elasticity in the long-run, a higher price-elasticity demand, and a higher output elasticity of demand (and vice-versa). This particular correlation in the elasticities magnifies the uncertainty in the FEVD.

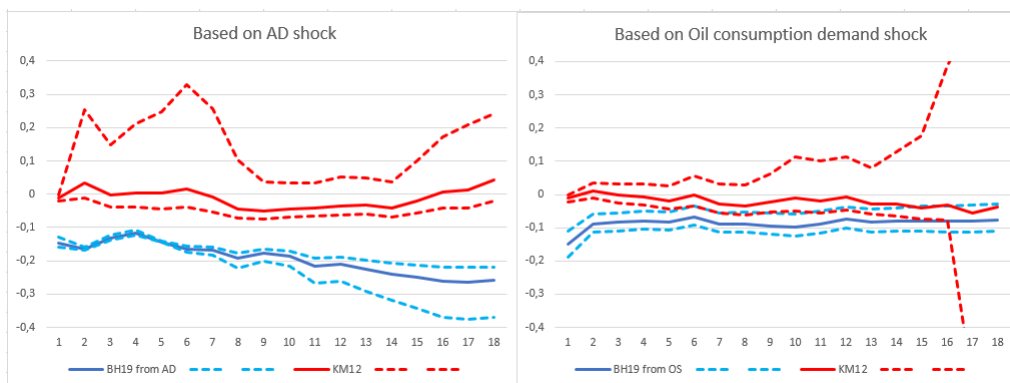
Finally, our paper also relates to a growing literature that endogenizes the price of oil (or other commodities) in DSGE models (e.g., [Balke et al., 2010](#); [Nakov and Pescatori, 2010](#); [Bodenstein et al., 2011](#); [Unalmis et al., 2012](#); [Peersman and Stevens, 2013](#); [Bergholt et al., 2019](#); [Houssa et al., 2023](#)). Our main contribution to that literature is to study the sources of oil price fluctuations in a framework where we derive analytically our four oil market elasticities to introduce priors. Moreover, DSGE models traditionally use standard CES productions or utility functions that impose a value of one on the output or income elasticity of oil demand, while we relax this assumption. Imposing this value would lead us to underestimate the contribution of global economic activity shocks.

The remainder of the paper is structured as followed. We motivate our analysis in section 2 by showing that SVAR studies may also produce substantially different oil supply elasticities in the medium-run. In Section 3, we present our RBC model of the global oil market and derive the analytical form of our demand and supply elasticities. There, we also discuss the Bayesian estimation of the model and we show that the sign and magnitude restrictions popular in the SVAR literature are all jointly satisfied a priori. Section 4 illustrate the crucial role that elasticities play in the transmission of demand and supply shocks to oil prices. Section 5 quantifies the source of oil price fluctuations. We show that our results are robust to to changes in the data in section 6. The last section concludes.

2 Motivation: supply elasticities in SVAR studies

We compare the price-elasticity of oil supply at different horizons (ranging from one to 18 months) reported in KM12 and BH19. We compute these elasticities based on the ratio of the impulse response of oil production and oil prices following an aggregate demand shock (left panel) and an oil-specific demand shock (right panel), which results in an approximation of the true elasticities (one can only recover the true value of this elasticity at the one-month horizon in BH19). It shows that the one-month price-elasticities of oil supply strongly differ between KM12 and BH19, which comes from the different bonds or priors imposed in these studies.

Figure 1: Supply elasticities in BH19 and KM12



What stands out is that the disagreement on the one-month price-elasticity of oil demand persists in the medium run. In KM12, the price elasticity of oil demand is constrained to be within the range of $[0, 0.025]$ on impact and then becomes statistically not different from zero. In contrast, the estimated elasticities of BH19 remain larger than zero. In both cases, the median estimates are rather flat, which might indicate that imposing a certain value on impact (via magnitude restrictions in KM or using a prior in BH19) also influence this elasticity in the medium run.

These estimated elasticities also violate reasonable restrictions that one might have wished to impose on the medium-run price-elasticity of oil supply. First, some draws seem to support a negative elasticity, implying that oil production falls in response to higher oil prices. Second, it is not always clear if the medium-run elasticity is larger than the short-run elasticity. However, one might want to impose that the elasticity of oil supply cannot become negative, and that the medium-run elasticity is larger than the short-run elasticity.

It is important to note that these elasticities are only approximations. For example, a favorable aggregate demand shock could raise both oil prices and the cost of inputs used in

the oil extraction process, resulting in a fall in oil production, even with a positive price-elasticity. The implication is that, while the frontier SVAR models of KM and BH could violate reasonable restrictions on the medium or long-run elasticities of oil supply, imposing such restrictions would not be easy. BH19 showed how to impose a prior on the one-month elasticity, but extending this strategy to the medium or long-run elasticities is not straightforward in a SVAR environment.

We thus propose to build an RBC model where the short and long-run elasticities of oil supply can be expressed as a simple function of a few structural parameters. This strategy offers two advantages. First, the structural equations of the models imply that the long-run elasticity of oil supply must be larger or equal than that in the short run (for reasonable values of these structural parameters). Second, one can estimate and impose priors directly on both the short and long-run elasticities and recover the value of the structural parameters as a by-product. In the next section, we describe this model.

3 RBC model with oil

We extend a standard RBC model (e.g., [King and Rebelo, 2000](#); [Francis and Ramey, 2005](#)) featuring external habits in consumption ([Abel, 1990](#)), and investment adjustment costs ([Christiano et al., 2005](#)) with storable oil (e.g., [Unalmis et al., 2012](#)) used as input in the production of the final good. Here, we focus on the optimization problem of the firms because the household's part of the model is fully standard (they provide the firms with labor and capital).

In what follows, we describe our hypothesis and their impact on the elasticities of supply and demand for oil that are key to identify the relative contribution of aggregate demand and oil-sector specific shocks (including oil supply, oil consumption demand, and oil inventory demand shocks). In [appendix A](#), we show in detail how we obtain our elasticities.

3.1 Oil supply

Oil producers combine capital (K), labor (H), and land (L) to produce oil $Y(o)$ sold at price $p(o)$ to final good producers or storage firms. They operate in a perfectly competitive environment and maximize their profits given by

$$p_t(o)Y_t(o) - w_tH_t(o) - r_t^k(o)K_t(o)$$

where w_t is the equilibrium wage (labor is perfectly mobile across sectors) and $r_t^k(o)$ is the rental rate of capital (which is sector-specific).

Their production function is given by a CES aggregation of these three production factors:

$$Y_t(o) = \varepsilon_t^{os} \left[(1 - \alpha_o - \gamma_o) \left(\frac{H_t(o)}{H(o)} \right)^{\frac{\theta_o-1}{\theta_o}} + \alpha_o \left(\frac{K_t(o)}{K(o)} \right)^{\frac{\theta_o-1}{\theta_o}} + \gamma_o \left(\frac{L_t(o)}{L(o)} \right)^{\frac{\theta_o-1}{\theta_o}} \right]^{\frac{\theta_o}{\theta_o-1}} \quad (1)$$

where ε_t^{os} is an oil-supply shock and θ_o is the elasticity of substitution between production factors. Throughout the section, CES production functions are written in their normalized form (e.g., [Temple, 2012](#); [Cantore and Levine, 2012](#)) and any variable without a subscript t represent its steady state that we use as a normalizing constant. Thus, α_o and γ_o are the income shares of capital and land in the oil sector, respectively. The supply of land is assumed to be fixed such that $L_t(o) = L(o)$.

To obtain an analytical expression of the short-run elasticity of oil supply, we first derive the linearized supply curve from the first order condition to the oil producer's problem knowing that the stock of capital is predetermined. We then evaluate the response of oil supply to a change in oil prices holding wages constant. The immediate oil supply response to a shift in prices is given by:

$$\hat{Y}_t(o) = \sigma_{\text{SR}} \hat{p}_t(o) \quad (2)$$

where the term $\sigma_{\text{SR}} = \frac{\theta_o(1-\alpha_o-\gamma_o)}{\alpha_o+\gamma_o}$ is the short-run (one month) price elasticity of oil supply. It increases with the elasticity of substitution between production factors θ_o and with the labor share $(1 - \alpha_o - \gamma_o)$. For given (or observed) labor share, one can thus calibrate θ_o to reach any short-run price-elasticity of oil supply (this elasticity will be estimated). In the limit case where θ_o or the labor share are calibrated to zero, supply becomes perfectly inelastic in the short-run, as it is not possible substitute capital or land with labor - the only variable production factor in the short run - to adjust production.

To find a closed-form solution for the long-run elasticity of oil supply, we evaluate how oil supply adjust to a permanent increase in oil prices holding wages constant. In the long-run, both labor and capital become variable production factors. The long-run oil supply response to a permanent shift in prices is given by:

$$\hat{Y}_t(o) = \sigma_{\text{LR}} \hat{p}_t(o) \quad (3)$$

where $\sigma_{\text{LR}} = \frac{\theta_o(1-\gamma_o)}{\gamma_o}$ is the long-run price elasticity of oil supply. The key parameter influenc-

ing this value is the land income share (γ_o): the higher the land share – a fixed production factor in the short and long-run – the harder it is to adjust production in response to a change in oil prices. After calibrating θ_o to any value reflecting the short-run elasticity of oil supply, one can easily adjust the land income share to reach any desired level for this long-run elasticity.

An advantage of our simple model is thus that we can easily estimate and/or introduce our prior knowledge on the short and long-run elasticities of oil supply. When we estimate the model, we estimate σ_{SR} and σ_{LR} directly and we assume that the labor share is set to 0.5. This allows us to easily recover the structural parameters of our model with $\theta_o = \sigma_{\text{SR}}$ and $\gamma_o = \frac{\sigma_{\text{SR}}}{\sigma_{\text{SR}} + \sigma_{\text{LR}}}$ and $\alpha_o = 0.5 - \gamma_o$. The only constraint is that the medium-run elasticity cannot be lower than the short-run elasticity, which is supported by the empirical evidence described in section 3.6.

3.2 Oil consumption demand

Oil is consumed by final good producers who combine labor, capital, and oil inputs to produce their goods ($Y_t(f)$) that are sold to households as consumption or investment goods. They operate in a perfectly competitive environment and maximize their profits given by

$$Y_t(f) - w_t H_t(f) - r_t^k(f) K_t(f) - p_t(o) O_t(f)$$

where $p_t(o)$ is the real price of oil and the final good is used as a numeraire to express all prices in real terms.

Production consists of two steps. First, firms combine labor and capital - their value added denoted VA - using a standard Cobb-Douglas technology:

$$VA_t(f) = \varepsilon_t(f) K_t(f)^\alpha H_t(f)^{1-\alpha} \quad (4)$$

where $\varepsilon_t(f)$ is productivity shock capturing exogenous shifts in global economic activity. Second, they combine their value added with oil using a modified CES production function to produce final goods:

$$Y_t(f) = Y(f) \left[(1 - \omega_o) \left(\frac{VA_t(f)}{VA(f)} \right)^{\frac{\eta_o - 1}{\eta_o}} + \omega_o \left(\frac{\varepsilon_t^{od}(O_t(f) - \bar{O})}{(O(f) - \bar{O})} \right)^{\frac{\eta_o - 1}{\eta_o}} \right]^{\frac{\eta_o}{\eta_o - 1}} \quad (5)$$

where the parameter η_o controls the oil demand elasticity of substitution and ω_o is the oil

share in production. The parameter \bar{O} controls the output-elasticity of oil demand. When \bar{O} is set to zero, we are back to a standard CES production with an output-elasticity of oil demand equal to one.

The first order condition (w.r.t. oil) of the final good producer's problem yields the following linearized oil demand curve:

$$\hat{O}_t(f) = \delta_y \hat{Y}_t(f) - \delta_p \hat{p}_t(o) - \hat{\varepsilon}_t^{od} \quad (6)$$

where $\delta_y = \frac{O(f) - \bar{O}}{O(f)}$ is the output elasticity of oil demand and $\delta_p = \eta_o \frac{O(f) - \bar{O}}{O(f)}$ is the price elasticity of oil demand.

When we estimate the model, we estimate δ_y and δ_p directly. This allows us to easily recover the structural parameters of our model with $\bar{O} = (1 - \delta(y))O(f)$ and $\eta_o = \frac{\delta(p)}{\delta(y)}$. An advantage of our framework is thus that we can easily calibrate or estimate the price and output elasticities of oil demand.

3.3 Storage and oil-inventory demand

Competitive risk-neutral storage firms maximize their expected profits of holding oil inventories $V_t(o)$:

$$\varepsilon_t^v(o) p_{t+1}(o) V_t(o) - p_t(o) V_t(o) \Phi(V_t(o)) \quad (7)$$

where ε_t^v is a storage-demand shock, $p_{t+1}(o)$ is the expected future price of oil, and $\Phi(V_t(o)) = \beta (V_t(o)/V(o))^{1/\varphi(o)}$ captures a storage cost that depends on the total level of oil inventories. Under these assumptions, the linearized FOC of the storage firm is:

$$\hat{V}_t(o) = \varphi(o) \left\{ \hat{p}_{t+1}(o) - \hat{p}_t(o) - \hat{\varepsilon}_t^v \right\} \quad (8)$$

The implication is that storage demand depends positively on the expected price gains of holding oil inventories and on a shock that capture speculation/precaution motives.

3.4 Market clearing

The oil markets clear when the gap between demand and supply matches the change in inventories: $V_t(o) - V_{t-1}(o) = Y_t(o) - O_t(f)$. The final good market clears when production equals the demand of households for consumption and investment goods: $Y_t(f) = C_t + I_t(o) + I_t(f)$. Appendix B describes the households' consumption and investment demand.

3.5 Bayesian estimation method and the data

We estimate our model with Bayesian methods (e.g., [DeJong et al., 2000](#); [Otrok, 2001](#); [Schorfheide, 2000](#)) in Dynare ([Adjemian et al., 2011](#)). We focus on a period ranging from 1975M1 to 2019M12. The choice of our starting date is guided by the fact that K09, K12, K14 and BH19 use a sample of monthly data starting in 1973 to estimate SVAR with 24 lag. The end date excludes the pandemic period. We experimented starting after the Great Moderation (as CCI19) and in 1995 after oil supply became less volatile to account for the potential presence of time-variation in our estimated parameters and to evaluate the drivers of oil prices in the most recent period.³

Our data set consists in four variables: the prices, inventories, and production levels of oil, and a world industrial production index. In our baseline analysis, we use the monthly growth rate of these variables while we use HP-filtered data in a robustness exercise. The price of oil comes from the world bank pink sheet. We compute real oil prices by dividing the nominal prices by the US CPI. Crude oil production at the world level directly comes from the EIA. Following [Kilian and Murphy \(2014\)](#), we use total US crude oil inventories scaled by the ratio of OECD petroleum stocks over US petroleum stocks. Data are from the EIA and this strategy is followed because crude oil inventories are not available at the world level. The world industrial production index is borrowed from [Baumeister et al. \(2022\)](#).

In a robustness exercise, we also use the first principal component extracted from the prices of a set of non-energy industrial commodities following [Alquist et al. \(2019\)](#). As argued by [Kilian and Zhou \(2018\)](#), this measure can capture anticipations about the future level of economic activity. The price of industrial commodities also has the advantage of being easily introduced in our model (by adding a second commodity used to produce final goods).

³Our starting date excludes the 1973/74 oil crisis. As argued by [Baumeister and Kilian \(2016\)](#), there are two conflicting views regarding the drivers of the oil price increase in that period. The conventional view relies on an oil supply cut in response to the war between Israel and a coalition of Arab countries. But another likely explanation for the rise in oil price is an endogenous adjustment to a strong demand for oil combined to a devaluation of the dollar. At that time, the price of oil produced by Middle Eastern countries had been fixed for a duration of five years. It is thus possible that the increase in oil prices were the reflection of demand/supply imbalances that predated the war. Exploring this question would however best be answered with a model capturing a structural break in the pricing of oil. By applying our simple model to this period, we might wrongly fully associate the oil price increase in 73/74 to an oil supply shock, because our model does not account for fixed prices. Avoiding this issue is an extra motivation for starting our estimation in 1975.

3.6 Prior’s distributions

The estimation of the short-run elasticity of oil supply has been the subject of large body of literature including the contributions and reviews of the literature of CC19, BH19 and Kilian (2022). While the exact value of this elasticity is debated, it is likely very low. On the one hand, CC19 estimate its value around 0.1, which also correspond to the average value in their search of the literature. The mode of BH’s prior is also set to 0.1. On the other hand, Kilian (2022) provides more support for lower values and the bounds of 0.025 introduced in KM12 and KM14. We build our prior as a compromise between these two views, with a mean of 0.06 and confidence bands that includes values in the range of 0 to 0.1.

We set the prior mode of the long-run elasticity of oil supply to 0.4 based on the literature reviews of Greene and Leiby (2006) and a bit larger than the estimate of Krichene (2002). The price-elasticity of oil supply is thus likely larger in the long run than in the short run.

We set the prior mean of our price-elasticity of demand to 0.13 corresponding to the mode of the values reported in the search of the literature of CCI19 and not far from their own estimates. Our prior is also similar to that of BH19 and consistent with the meta-analysis of Hamilton (2009) who reported values in the range of 0.05 to 0.17. Our prior mean is however slightly lower than the estimate of KM14 (0.26) and Inoue and Kilian (2020) (0.18) but we allow for a sufficiently large variance such that our prior does not rule out these values.

For the prior mean of the output elasticity of oil demand, we chose a value of one, which is the implicit assumption imposed by the literature that uses standard CES production or utility functions. This value stand within the bond of the values reported in the literature (Havranek and Kokes, 2015; Golombek et al., 2018; Chang et al., 2019). A value of one tends to be supported by studies that do not control for the stock of car and thus capture a total long-run elasticity. We however allow for a relatively wide prior to account for the dispersion in the results obtained in the literature. We discuss the remaining parameters in appendix B.2.

3.7 Restrictions implied by our model and priors

We now discuss the sign and magnitude restrictions implied by our modelling, calibration, and prior. For this purpose, we generate the set of admissible IRFs based on 100 random draws from our priors. The results are displayed in figure 2. We consider that a sign restriction is imposed a prior when all draws are of the same size.

We compare our restrictions with popular sign restrictions used in the SVAR literature summarized in table 1. Our restrictions are consistent with the sign restrictions of KM14 and


Table 1: Sign restrictions in KM14 and BH19

	Oil supply	Global activity	Oil-consumption	Oil-inventory
Oil production	-/-	+/+	+	+/+
Industrial prod	-/-	+/+	-	-/-
Oil price	+/+	+/+	+	+/+
Inventories				+/+

Sign restrictions imposed a priori in KM14 (black) and checked a posteriori in BH19 (blue). In BH19, their priors already imply a high probability that these restrictions are satisfied.

Figure 2: Set of admissible IRFs

Note: Variables expressed in percentage deviation from steady state. Horizon in quarters. All lines correspond to a particular draw from our priors. 100 draws in total.



BH19, which have been widely used in the SVAR literature (e.g., [Lippi and Nobili, 2012](#); [Baumeister et al., 2010](#); [Baumeister and Peersman, 2013b,a](#); [Peersman and van Robays, 2009](#); [Lütkepohl and Netšunajev, 2014](#)), Moreover, the impact of oil-specific consumption (and inventory) demand shock on global industrial production is quite small, which echoes the magnitude restriction of KM12. Thus, our model combines different identifying restrictions used in the SVAR literature with prior studies of the oil market in a statistically and structurally coherent framework.

4 Oil elasticities and the transmission of structural shocks

Since the seminal contribution of K09, a large literature studied the determinants of oil price fluctuations. Within this SVAR literature, there exist a disagreement over the relative size of the effect of oil demand and oil supply shocks on oil prices: models that allow for a larger elasticity of oil supply in the short run are associated with a larger relative contribution of oil supply shocks to oil prices (and vice-versa). We use our model to show that varying the impact elasticity of oil supply (from 0 to 0.14) while leaving all other parameters of our model unchanged has a limited direct effect on the dynamics of oil prices (section 4.1). So, the reason why different values for this short-run oil supply elasticity translate into radically different dynamics in oil prices as reported in the SVAR literature is far from obvious.

We conjecture that imposing a lower short-run price-elasticity of oil supply affects the estimates of the other elasticities in the model. We show that changes in the medium or long-run price elasticity of oil supply (section 4.2) and in the price and income elasticities of oil demand (section 4.3) can substantially alter the dynamic response of oil prices to structural shocks. A conclusion of our exercise is that the debate, which initially concentrated on the short-run price elasticity of oil supply, and later included the price elasticity of demand, should be broadened to include the medium and long-run elasticities of oil supply.

4.1 The one-month price elasticities of oil supply (and demand)

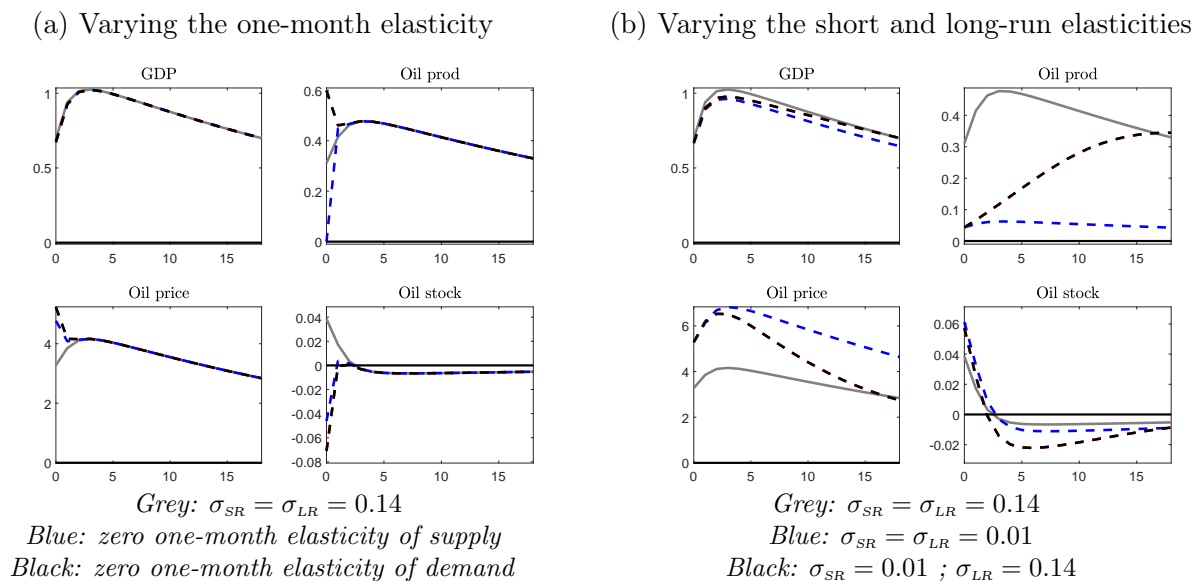
A difference in the one-month elasticity of oil supply, alone, is unlikely to explain the very different response of oil prices to a global economic activity shock. To illustrate this claim, we simulate the effect of a global economic activity shock in two economies. In the first economy, set a value of 0.14 for both the short and long-run elasticities of oil supply. These values are consistent with those obtained by BH19, so we refer to this economy as the BH19-inspired calibration. Other parameters are calibrated to our priors. In the second, we set the one-month price-elasticity of oil supply to zero while leaving all other parameters unchanged.

We implement this zero-restriction on the one-month elasticity by assuming that hours worked in the oil sector are fixed one month in advance. This assumption only restricts the response of oil supply on impact. Simply put, the one-month price-elasticity of oil supply is zero while the two-month elasticity is identical to that of the BH19-inspired calibration (for more details, see appendix A.3). The experiment allows us to isolate the effect of a change in the one-month elasticity of oil supply, which is at the centre of the debate in the SVAR literature.

Imposing a perfectly inelastic oil supply in the very short run amplifies the effect of a

global economic activity shock on oil prices but only in the short run (figure 3a). The initial rise in oil prices is larger when oil supply is not allowed to respond within a month, but the gap with the baseline response of oil prices quickly decay. Even on impact, the difference between these two scenarios is limited, because inventories can be used to partially accommodate the short-run imbalances in oil supply and demand.

Figure 3: Oil supply elasticities and the transmission of global economic activity shocks



An important implication of these results is that varying the one-month elasticity of oil supply while leaving all other elasticities and parameters unchanged only has limited direct repercussion on the effect of demand shocks on oil prices. The widely used restrictions on the one-month price-elasticity of oil supply of K09, KM12 and KM14 thus have a limited direct impact on the relative contribution of global economic activity shocks - if they do not affect other parameters of the model.

Since the literature reviewed above not only disagree on the short-run drivers of oil prices, but also how what causes their fluctuations in the medium to long run, one can conjecture that these restrictions affect some other parameters of the model.⁴ A possibility that we discuss next is that imposing a prior on the one-month price elasticity of oil supply also bears repercussion on the other three elasticities: the long-run price-elasticity of oil supply, and the price and income elasticities of oil demand.

⁴ CCI18 reports the 24-month ahead forecast error variance decomposition of oil prices; KM12 up to 18 months; BH19 and K09 describe their cumulative effects (HD)

4.2 The long-run price elasticity of oil supply

A joint disagreement over the short and long-run price-elasticities of oil supply – such as the one we described in section 2 for K12 and BH19 – has the potential to strongly affect the dynamics of oil prices. To make this point, we simulate the effect of an aggregate demand shock under two calibrations. The first calibration, inspired by the results of BH19, set a value of 0.14 for both the short and long-run elasticities of oil supply. The second, inspired by K12, has a value of 0.01 for both the short and long-run elasticities of oil supply. Note that neither K12 nor BH19 imposes a value on this elasticity in the long run. However, it turns out that their IRFs-implied elasticities are consistent with this calibration.

A more inelastic oil supply in the short- and long-run amplifies the response of oil prices to a global economic activity shock. An increase in global economic activity always stimulates the demand for oil, as captured by equation (6), which pushed up the price of oil. Higher oil prices then encourage industrial good producers to substitute oil inputs with other production factors and encourages oil producers to raise production. When oil supply is more inelastic (in blue), all the adjustment goes through the demand for oil, and it takes a larger increase in oil prices to convince industrial good producers to substitute enough oil inputs with other production factors.

Moreover, we show that changes in the value of the long-run price-elasticity of oil supply can substantially affect the dynamics of oil prices. We perform one last experiment, where we calibrate the short-run price-elasticity of oil supply to 0.01 and its long-run counterpart to 0.14. Comparing the results of this simulation with that inspired by KM12 allows us to isolate the role of this long-run elasticity. On impact, these two economies have the same price-elasticity of supply. In the KM12-inspired economy, this elasticity remains low, while it gradually increases up to 0.14 in the other experiment.

When supply becomes more elastic in the long run, the increase in prices after an aggregate demand shock is much less persistent. The gradual increase in oil supply helps mitigating the effect of a higher oil demand on oil prices. The behavior of storage firms also depends on the medium-run response of oil supply. While storage firms find it optimal to sell their inventories when oil supply gradually picks up in the medium run, they hold on to their inventories when they anticipate that the imbalance between oil demand and supply will persist. Thus, in a model that incorporate optimal storage decisions, the effect of an economic activity shock on oil price is also stronger in the short run when the medium-run price elasticity of oil supply is lower.


An important conclusion that emerges from our analysis is that, while the short-run

elasticity of oil supply has received a lot of attention, its long-run counterpart is also crucial when it comes to evaluate the sources of oil prices fluctuations. In SVARs, it is difficult to pin down the value of this elasticity. In contrast, it is easy to introduce a prior and to estimate this elasticity in our RBC model, as it depends on a few structural parameters. In the section 5, we show how different assumptions on the short- and long-run elasticities interact, and we show that imposing different assumptions on the long-run elasticity can lead to substantially different conclusions.

4.3 The income and price elasticities of oil demand

Another possibility is that imposing a lower (higher) short-run price-elasticity of oil supply affects the estimated oil demand elasticity. This point is raised in BH19, CCI19 and [Herrera and Rangaraju \(2020\)](#), who noted that in SVAR studies, a low supply elasticity in the short run tend to come with a large demand elasticity.

Figure 4: Oil demand elasticity and the transmission shocks



In figure 4, we experiment with an increase in the price-elasticity of oil demand. Specifically, we increase the value of parameter δ_p from 0.13 (the value of our prior mode) to 0.27 (as in KM14). A higher demand elasticity mitigates both the effect of global economic activity and oil supply shocks. The higher this elasticity, the easier it is for firms to substitute oil producers with other inputs, when the economy is faced with a favorable global economic activity shock, or when faced with an exogenous reduction in oil supply. A reduction in

this demand elasticity, alone, is thus unlikely to affect the relative contribution of global economic activity and oil supply shocks.

Finally, we consider potential effect on raising the income elasticity of oil demand (δ_v), from one in the baseline, to 1.5. A higher income elasticity of oil demand amplifies the effect of global economic activity shocks on oil prices but does not affect the transmission of oil supply shocks.

In this section, we thus showed that a decrease in the short and long-run price-elasticity of oil supply combined with an increase in oil demand elasticities can shift the FEVD towards a lower relative contribution of oil supply shocks. To determine which shock dominates, it is thus important to pay a close attention to the value of all these elasticities. This is what we do in the next section, paying a particular attention to the long-run price elasticity of oil supply, which has not been discussed much in the SVAR literature.

5 Empirical results


5.1 Our estimated elasticities of oil demand and supply

We now discuss the estimated elasticities of oil demand and supply. The modes of our price-elasticity of oil supply stand at 0.09 and 0.14 in the short and long-run, respectively. While the point estimates are similar, the width of the distribution differs. Indeed, our posterior almost completely rules out values that are equal or larger than 0.14 for the short-run elasticity. In contrast, the posterior distribution is wider for the long-run elasticity, with a non-zero mass associated with values ranging from about 0.05 to 0.40. There are only few estimates available for the long-run elasticity, that we tried to summarize into our prior. Compared to that prior, we find a relatively low estimate in the long run.

A question that naturally arises is that of our interpretation of the long run. From the theoretical perspective, it is a value computed based on the steady-state properties of the model. How fast our economy converges to that steady state is an empirical question that in parts depends on our estimated capital adjustment costs in the oil sector. Based on our posterior distributions, we compute the elasticity of oil supply at different horizons using the ratios of the IRFs of oil production and oil prices (as in section x) when simulating persistent oil demand shocks. The elasticity tends to have fully converged to its steady-state value in two to five years. This relatively short horizon may explain why we find a relatively low elasticities compared to prior studies.

The mode of the short-run elasticity of oil supply is close to that of BH19 and CCI19,

Figure 5: Oil demand and supply elasticities




but our distribution is tighter. We can't find draws larger than 0.15, while 16% of the draws are larger than 0.19 in BH19's baseline model. Our estimates are larger than the upper bounds implemented in K12 and K14. In our baseline, the mode of the price-elasticity of oil demand stands at 0.22, which is in between the (short-run) estimates of BH19 (0.36), KM14 (0.25), [Inoue and Kilian \(2020\)](#) (0.18), and CCI19 (0.14). The mode of the output-elasticity of oil demand is 1.7, which is larger than our prior and the mode of BH19. One way to understand such a difference is that our parameter also captures this elasticity in the long run. Unfortunately, this parameter cannot be easily compared with SVAR studies other than that of BH19.

While many of the papers reviewed in CCI19 report short-run price elasticities of oil supply in the range of our estimates, KM12 build a strong case in favor of their upper-bound. We thus also estimate our model in two alternative specifications. First, we change our prior or the short-run elasticity of oil supply for a prior truncated at 0.025 while keeping our original prior for the long-run elasticity. Second, we use the truncated prior for the short-run elasticity in combination with a uniform prior of the long-run elasticity. This last experiment best represents the assumption of KM12 regarding the elasticity of supply: very inelastic on impact, and unconstrained in the subsequent periods. As is well-known in the SVAR literature, lowering the value of the oil supply elasticity translates into higher price-elasticity of oil demand.

As we discussed in the case of SVAR in section 2, a lower short-run elasticity of oil demand is also associated with a lower elasticity in the long-run. The estimated long-run

elasticity gets very small when we impose a bound on the short-run value, especially when we additionally do not introduce a prior on the long-run value. An advantage of our model over SVAR that remains in either case is that the long-run elasticity remains non-negative and higher than that in the short run. We also uncover a relationship between the value of the elasticity of oil supply and the output-elasticity of oil demand. In our model, imposing a bound on the short-run elasticity of oil supply also results in an increase in the output-elasticity of oil demand.

Figure 6: Relation between the short-run price-elasticity of oil supply and other elasticities



Note: Histogram representing the bivariate distribution of the SRPEOS with one other elasticity. Based on 50 000 MCMC draws. SRPEOS = short-run price-elasticity of oil supply; PEOD = price-elasticity of oil demand; IEOD = income-elasticity of oil demand; LRPEOS = long-run price-elasticity of oil supply.

Finally, we show that the relationship that we described above between the short-run elasticity of oil supply in the other elasticities is also present in the joint distribution of the parameters in our baseline estimation. Figure 6 sequentially plots the MCMC draws of the short-run elasticity of oil supply with the other three elasticities. It shows that the parameter draws displaying a lower supply elasticity in the short run tend to be associated with a lower supply elasticity in the long-run, a higher price-elasticity demand, and a higher output elasticity of demand (and vice-versa). This will bear repercussion for the confidence bands of the variance decomposition we present in the next section, because we can easily find combinations of draws for which all elasticities would tend to produce a high contribution of supply shocks, as well as the exact opposite. Intuitively, the correlation in the MCMC draws means that the uncertainty around the parameters estimates compounds to produce wider confidence bands of the FEVD.

5.2 What are the main drivers of oil prices?

What are the main drivers of the oil market? In this section, we revisit this question through the lenses of our RBC model. Our methodology offers two advantages over SVARs. First, we

easily combine different well-accepted identification assumptions in a structurally simple and coherent framework. Second, we can introduce prior knowledge on the long-run elasticity of oil supply.

Table 2: Variance decomposition

	OS	AD	OCD	OID
Oil prod	62.80	16.50	18.54	0.02
IP	2.23	86.98	9.92	0.00
Oil price	24.27	36.93	35.60	0.06
<i>68%</i>	(16.42 ; 37.28)	(26.51 ; 47.10)	(27.48 ; 44.22)	(0.04 ; 0.09)
<i>95%</i>	(12.11 ; 49.94)	(18.34 ; 61.21)	(19.01 ; 57.06)	(0.03 ; 0.12)
Oil stock	0.66	1.21	1.70	96.32

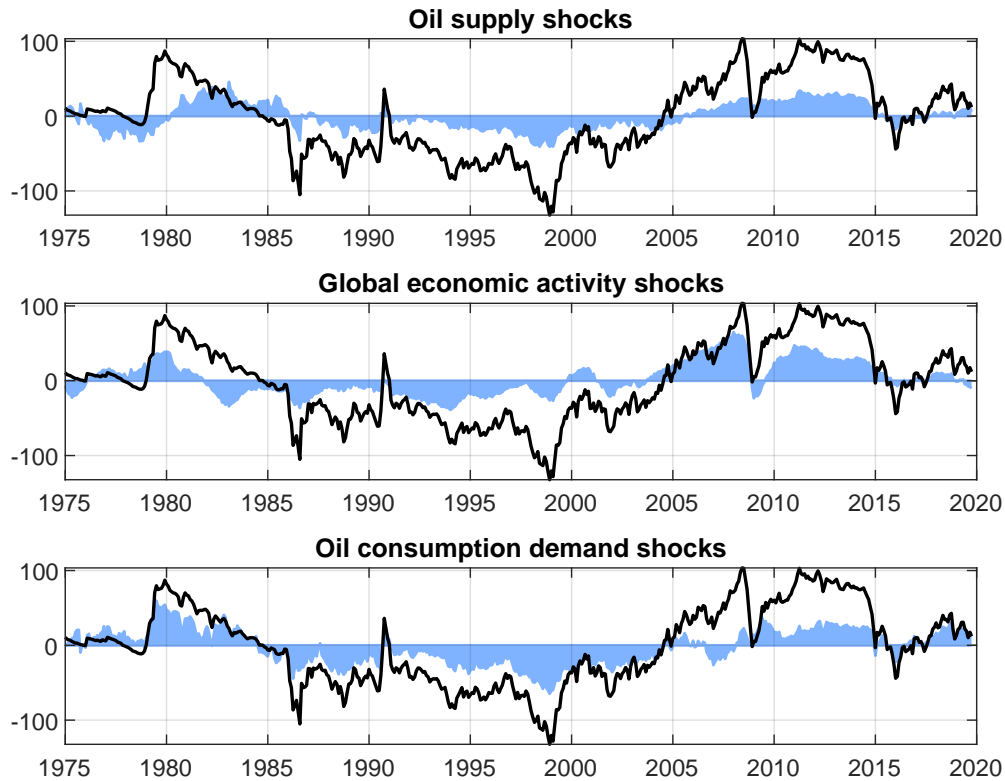
Note: OS = Oil supply shocks; AD = Global economic activity; OCD = Oil consumption demand; OID = Oil-inventory demand.

The results from our baseline analysis are summarized in table 2. We find that global economic activity and oil-demand shocks are the most important drivers of oil prices: they explain 37% and 36% of the FEVD in this variable, respectively. However, oil supply shocks are also very important as they explain about a quarter of the fluctuations in oil prices. Overall, we thus attribute a larger share of oil prices fluctuations to oil supply shocks compared to K09, KM12 and KM14. The reason is that we estimate larger short- and long-run price-elasticities of oil supply, as discussed in the previous section.

We also document a large uncertainty around the contribution of oil supply shocks. Indeed, we cannot completely rule out the possibility that oil supply shocks explain the majority of oil price fluctuations: the odds that oil supply shocks explain 50% of the variance in oil prices stand close to 5%. These results are surprising considering that we build a model that combines different identifying restrictions used in the SVAR literature with prior studies of the oil market in a single framework. This high level of uncertainty has two origins. First, small changes in the value of key elasticities have an important repercussion on the dynamics of oil prices. Second, the multivariate posterior distribution of the elasticities shows a correlation in the elasticities that magnifies the uncertainty in the FEVD.

The historical decomposition in figure 7 nevertheless confirms the results of Kilian (2009) that attributes the oil price boom in the early 2000s to the cumulative effect of aggregate demand shock. Global economic activity shocks were also a crucial driver of oil prices during the GFC and its recovery, and also explains about one-third of the fall in oil prices during the commodity price slump in 2015/16.

Figure 7: Historical decomposition of oil prices fluctuations



Oil supply shocks played a more limited role in the recent period, even though they contributed to the rise in oil prices observed after the GFC (shocks include the Libyan uprising in 2011 and the sanctions imposed on Iran in 2012) and to their fall during the commodity price slump (coinciding with the rapid growth of US shale and Iran's return to the market).

In contrast, oil supply shocks were a key driver of oil prices in the earlier sample. A succession of adverse supply shocks resulted in a strong upward pressure on oil prices in the early 80s. The Iranian revolution in 1979 was followed by the Iran-Iraq war, which began in September 1980 and lasted for eight years. In this period, the OPEC group also tried to proactively counter the fall of oil prices from their 1979 peak. Our results indicate that, together, these supply shocks mitigated the return of oil prices towards their equilibrium level, that would have been much faster without these shocks. The 1986 oil price collapse is then explained by a combination of weak demand and abundant supply, in part due by Saudi Arabi, that reversed its policy of responding to falling oil prices with lower production.

The peak in oil prices reached in 1979 is however explained by a combination of global

economic activity as well as oil consumption and inventory demand (the latter could be related to the higher uncertainty that followed the Iranian revolution), which echoes the results of Barsky and Kilian (2002) and Kilian and Murphy (2014). The contribution of global economic activity shocks then turned negative as the economy weakened after Volcker raised policy rates in the US. In that period, the contribution of oil-specific demand shocks turned from positive to negative, potentially reflecting efforts to reduce the use of oil following the two major oil crises of the 70s. Global economic activity shocks also were a key determinant of oil prices during the early-90s recession.

5.3 Does the long-run elasticity of oil supply matter?

In this section, we discuss the sensitivity of our results to changes in the long-run elasticity of oil supply. For this purpose, we re-estimate our model thrice. First, we assume that oil supply is perfectly elastic in the long run ($LRPEOS=\infty$) by calibrating the land share to zero. Second, we impose that the short and long-run elasticities of oil supply are identical ($LRPEOS=SRPEOS$), which is something that we obtain by assuming that the capital share is equal to zero. Third, we re-estimate our model with a flat prior on the long-run elasticity.

Table 3: The effects of varying the long-run elasticity of oil supply

	OS contribution to oil prices	LRPEOS	SRPEOS	PEOD	IEOD
Baseline	24.27 (16.42 ; 37.28)	0.16	0.09	0.22	1.71
LRPEOS= ∞	57.57 (50.53 ; 65.61)	∞	0.13	0.17	1.40
LRPEOS=SRPEOS	15.49 (10.19 ; 22.05)	0.09	0.09	0.23	1.76

Note: SRPEOS = short-run price-elasticity of oil supply; PEOD = price-elasticity of oil demand; IEOD = income-elasticity of oil demand; LRPEOS = long-run price-elasticity of oil supply.

The first experiment shows that imposing a perfectly elastic oil supply in the long run has a very strong impact on the relative contribution of oil supply shocks to oil prices. Indeed, the contribution of oil supply shocks increase to above 50%. The mode of the short-run price elasticity of oil supply is also bigger while demand elasticities are smaller when oil supply is assumed to be perfectly elastic in the long run.

The second experiment shows that a failure to distinguish between the short and long-run elasticities of oil supply in structural models can also bear repercussions on the dynamics of oil prices. Imposing that the short and long-run elasticities are identical results in a decline

in the long-run elasticity and marginal changes in the other elasticities. The contribution of oil supply shocks to oil prices falls to 15%.

The third experiment shows that the data support a long-run price elasticity of oil supply that is equal to its value in the short-run. The results are thus similar to those of the second experiment. The fact that we estimate a posterior elasticity that is larger in the long-run than in the short-run thus entirely comes from our prior that we build based on studies. Our prior thus supports a higher elasticity in the long run that translate into a larger contribution of oil supply shocks to oil prices.

The conclusion of this section is that disagreement regarding the value of the long-run elasticity of oil supply – or failure to account for the empirically relevant role of capital and land in the design of structural model – thus have the potential to strongly affect the dynamics of oil prices.

5.4 Does the income elasticity of oil demand matter?

We now estimate our model assuming that the income elasticity of oil demand is fixed to one. This is the implicit assumption made in most of the DSGE literature that often relies on standard CES production or utility function to model oil demand. Results are presented in table 4.

Table 4: The effects of varying the income elasticity of oil demand

	OS contribution to oil prices	LRPEOS	SRPEOS	PEOD	IEOD
Baseline	24.27 (16.42 ; 37.28)	0.16	0.09	0.22	1.71
IEOD=1	44.16 (31.46 ; 57.93)	0.19	0.13	0.17	1.00

Note: SRPEOS = short-run price-elasticity of oil supply; PEOD = price-elasticity of oil demand; IEOD = income-elasticity of oil demand; LRPEOS = long-run price-elasticity of oil supply.

We find that imposing a value of one to the income elasticity of oil demand raises the relative contribution of oil supply shocks to oil prices. A lower income elasticity of oil demand directly mitigates the effect of global economic shocks on oil prices and is associated with a higher elasticities of oil supply, which reinforces this effect. It is however associate with a lower price-elasticity of oil demand.

5.5 Sensitivity analysis

5.5.1 The short-run elasticity of oil supply

It is well known from the SVAR literature that imposing a lower short-run price-elasticity of oil supply is associated with a lower contribution of oil supply shocks to oil prices (e.g., [Herrera and Rangaraju, 2020](#)). This result is confirmed by our analysis.

Table 5: The effects of varying the short-run elasticity of oil supply

	OS contribution to oil prices	LRPEOS	SRPEOS	PEOD	IEOD
Baseline	24.27 (16.42 ; 37.28)	0.16	0.09	0.22	1.71
SRPEOS<0.025	10.26 (6.92 ; 14.63)	0.11	0.02	0.35	2.16
with LRPEOS=uniform	7.66 (5.50 ; 10.69)	0.06	0.02	0.35	2.17
SRPEOS=0	9.10 (6.29 ; 12.66)	0.10	0.00	0.39	2.25

Note: SRPEOS = short-run price-elasticity of oil supply; PEOD = price-elasticity of oil demand; IEOD = income-elasticity of oil demand; LRPEOS = long-run price-elasticity of oil supply.

A novelty of our approach is that we can describe the consequence of imposing different assumptions on the short-run elasticity of oil supply to the other elasticities. A reduction in the short-run elasticity is associated with a lower long-run oil-supply elasticity and with a higher income elasticity of demand, which both contribute to lowering the relative contribution of oil supply shocks.

5.5.2 Most recent sample

We also estimate our model starting after the great moderation (as in CCI19) or after 1995 (when oil supply became relatively less volatile). It is easy to estimate small-scale RBC models with Bayesian methods in small samples, and this experiment accounts for the potential time-variation in oil-market elasticities. [Baumeister and Peersman \(2013b,a\)](#) show that the volatility of oil price and production changed in the second-half of the 80s, and that this shift is related to a decline in the price-elasticities of oil demand and supply.

Two results stand out. First, the short-run price-elasticity of oil supply and the elasticities of oil demand are smaller in the recent sample, which corroborates the results of [Baumeister and Peersman, 2013b,a](#)). Second, oil supply shocks remain an important driver of oil prices. The reason is that the long-run elasticity of oil supply did not decreased as much as its

Table 6: Sub-period analysis

	OS contribution to oil prices	LRPEOS	SRPEOS	PEOD	IEOD
Baseline	24.27 (16.42 ; 37.28)	0.16	0.09	0.22	1.71
1985-2019	30.02 (19.19 ; 42.51)	0.17	0.04	0.17	1.42
1995-2019	23.83 (14.79 ; 37.62)	0.13	0.04	0.12	1.13

Note: SRPEOS = short-run price-elasticity of oil supply; PEOD = price-elasticity of oil demand; IEOD = income-elasticity of oil demand; LRPEOS = long-run price-elasticity of oil supply.

short-run counterpart. Also note that contribution of global economic activity increased strongly, while the contribution of oil-specific demand shocks declined substantially.

6 Robustness checks

6.1 Non-energy commodity prices prices and news shocks

The price of non-energy commodities can capture anticipations about the future level of economic activity. Several authors thus argued that the joint behavior of oil and metal prices can help to identify oil-specific shocks (e.g., [Barsky and Kilian, 2001](#); [Arezki and Blanchard, 2014](#); [Kilian and Zhou, 2018](#); [Caldara et al., 2019](#)). We thus check whether the introduction of a news shock together with the first principal component in the price of non-energy commodities ([Alquist et al., 2019](#)) helps in identifying the drivers of oil price fluctuations.

The news shock we consider is a signal about the future level of productivity and thus captures an expected global economic activity shock. When receiving favorable news shocks, agents can increase their inventory demand to sell their stocks during the boom, when prices will be high. The immediate effect of a news shocks is to raise the price and production of a commodity through a higher demand for inventories. We re-estimated our model allowing for anticipated shocks following the methodology of [Davis \(2007\)](#); [Fujiwara et al. \(2011\)](#) and [Schmitt-Grohé and Uribe \(2012\)](#). We implement a prior on the relative contribution of the anticipated and unanticipated global economic activity shocks that gives an equal weight to these two components. We also assume that the signal is received 1 months in advance (corresponding to the time it takes to ship commodities across the world). Adding news shocks in the estimation does not quantitatively affect our main results.

6.2 HP-filtered data

Our results are also robust to the use of hp-filtered data.

7 Conclusion

In this paper, we show how one can introduce priors on the short- and long-run elasticity of oil supply, as well as on the price and income elasticities of oil demand in a simple RBC model. The ability to introduce priors on both the short- and long-run elasticity of oil supply is advantage of our framework compared to SVAR, where long-run elasticities depend on a complex function of many parameters.

We use our model to shed more lights on the sources of oil prices fluctuations. We find an important role for the long-run elasticity of oil supply in explaining the disagreement found in the SVAR literature. While the SVAR literature has focused on the short-run price elasticities of oil supply and demand, we show that the long-run elasticity of oil supply is also important.

Our results support the view that demand factors are the most likely dominant factors in explaining oil prices fluctuations, although supply shocks do play a non negligible role. However, uncertainty around the oil market elasticities translate into substantial uncertainty in oil prices dynamics, implying that we find a small possibility that oil supply shocks might dominate oil demand shocks.

While we leave this issue for future research, our methodology could inspire the calibration of models that study the costs of a green transition. In these models, the long-run elasticities of commodity demand and supply are crucial and have an important effect on the cost estimates. Our framework could easily be extended to introduce multiple commodities that will be required in this transition.

A Oil market elasticities

A.1 The short-run elasticity of oil supply

The short-run price elasticity of oil supply (SRPEOS) is defined as the percentage change in the quantity supplied divided by the percentage change in the price of oil within one month, holding the equilibrium wage in the economy constant. In this appendix, we show in detail how we obtain a closed-form solution for this elasticity.

Oil producers maximize their profits given by

$$p_t(o)Y_t(o) - w_tH_t(o) - r_t^k(o)K_t(o)$$

The first order conditions w.r.t. hours worked is $p_t(o)\frac{\partial Y_t(o)}{\partial H_t(o)} = w_t$. Developing gives:

$$w_t = p_t(o)(1 - \alpha_o - \gamma_o) \left(\frac{H(o)}{Y(o)}\right)^{\frac{1-\theta_o}{\theta_o}} \left(\frac{H_t(o)}{Y_t(o)}\right)^{-\frac{1}{\theta_o}} (\varepsilon_t^{os})^{\frac{\theta_o-1}{\theta_o}} \quad (9)$$

Isolate $H_t(o)$:

$$H_t(o) = (1 - \alpha_o - \gamma_o)^{\theta_o} Y_t(o) \left(\frac{p_t(o)}{w_t}\right)^{\theta_o} \left(\frac{H(o)}{Y(o)}\right)^{1-\theta_o} (\varepsilon_t^{os})^{(\theta_o-1)} \quad (10)$$

Since the CES is normalized, we know that $(1 - \alpha_o - \gamma_o)$ is the labor income share and that $wH(o) = (1 - \alpha_o - \gamma_o)p(o)Y(o)$. Plugging $\frac{H(o)}{Y(o)} = \frac{(1-\alpha_o-\gamma_o)p(o)}{w}$ in the previous expression gives:

$$H_t(o) = (1 - \alpha_o - \gamma_o) Y_t(o) \left(\frac{p_t(o)}{w_t}\right)^{\theta_o} \left(\frac{p(o)}{w}\right)^{1-\theta_o} (\varepsilon_t^{os})^{(\theta_o-1)} \quad (11)$$

Similarly, the first order conditions w.r.t. capital is

$$K_t(o) = \alpha_o Y_t(o) \left(\frac{p_t(o)}{r_t^k(o)}\right)^{\theta_o} \left(\frac{p(o)}{r^k(o)}\right)^{1-\theta_o} (\varepsilon_t^{os})^{(\theta_o-1)} \quad (12)$$

Linearizing these last two expression gives:

$$\hat{H}_t(o) = \hat{Y}_t(o) + \theta_o (\hat{p}_t(o) - \hat{w}_t) - (1 - \theta_o) \hat{\varepsilon}_t^{os} \quad (13)$$

$$\hat{K}_t(o) = \hat{Y}_t(o) + \theta_o (\hat{p}_t(o) - \hat{r}_t^K(o)) - (1 - \theta_o) \hat{\varepsilon}_t^{os} \quad (14)$$

Linearizing the production function (1) gives:

$$\hat{Y}_t(o) = (1 - \alpha_o - \gamma_o) \hat{H}_t(o) + \alpha_o \hat{K}_t(o) \quad (15)$$

To find a closed-form solution for the short-run elasticity of oil supply, we assume that the stock of capital is fixed at its steady state ($\hat{K}_t(o) = 0$), that the labor supply is perfectly elastic ($\hat{w}_t = 0$), and that there are no supply shocks ($\hat{\varepsilon}_t^{os} = 0$). Plugging the linearized labor demand equations (13) into the production function:

$$\hat{Y}_t(o) = \frac{\theta_o(1 - \alpha_o - \gamma_o)}{\alpha_o + \gamma_o} \hat{p}_t(o) \quad (16)$$

Finally, let us define the SRPEOS as $\sigma_{\text{SR}} = \frac{\theta_o(1 - \alpha_o - \gamma_o)}{\alpha_o + \gamma_o}$. In our empirical analysis, we assume that the labor share is set to 1/2 and the SRPEOS is simply θ_o .

A.2 The long-run elasticity of oil supply

The long-run price elasticity of oil supply (LRPEOS) is measured as the change in the quantity of oil supplied in the long-run in response to a permanent one percent increase in the price of oil, holding the equilibrium wage in the economy constant. In this appendix, we show in detail how we obtain a closed-form solution for this elasticity.

To find a closed-form solution for the long-run elasticity of oil supply, we evaluate how oil supply adjust to a permanent increase in oil prices. We plug the linearized labor and capital demand equations (13) and (14) into the production function and get:

$$\hat{Y}_t(o) = \frac{1 - \gamma_o}{\gamma_o} [\theta_o \hat{p}_t(o) - (1 - \theta_o) \hat{\varepsilon}_t^{os}] - \frac{1 - \alpha_o - \gamma_o}{\gamma_o} \theta_o \hat{w}_t - \frac{\alpha_o}{\gamma_o} \theta_o \hat{r}_t^K(o) \quad (17)$$

Holding $\hat{\varepsilon}_t^{os} = 0$, $\hat{r}_t^K(o) = 0$ and $\hat{w}_t = 0$, this gives:

$$\hat{Y}_t(o) = \frac{\theta_o(1 - \gamma_o)}{\gamma_o} \hat{p}_t(o) \quad (18)$$

where γ_o is the share of land in the oil sector. Remember that when $1 - \alpha_o - \gamma_o = 0.5$, θ_o represents the SRPEOS (see equation 16), and $\gamma_o \in [0; 0.5]$. In that case, the LRPEOS must be equal or larger than the SRPEOS, and the LRPEOS is bounded by $[\theta_o; \infty]$. Let us define the LRPEOS as $\sigma_{\text{LR}} = \frac{\theta_o(1 - \gamma_o)}{\gamma_o}$.

In our empirical analysis, we directly estimate the SRPEOS and the LRPEOS - rather than estimating the structural parameters of the model - to easily introduce our prior knowl-

edge on these elasticities with Bayesian priors on σ_{SR} and σ_{LR} . It implies that the structural parameters are retrieved from these elasticities as $\theta_o = \sigma_{\text{SR}}$, $\gamma_o = \frac{\sigma_{\text{SR}}}{\sigma_{\text{SR}} + \sigma_{\text{LR}}}$ and $\alpha_o = 0.5 - \gamma_o$.

There are two points that deserve more explanations. First, the fact that $\hat{r}_t^k(o) = 0$ in the long run is not an assumption but an outcome that derives from the steady-state properties of the model. Indeed, the steady-state level of the rental rate of capital is $r^k(o) = 1/\beta - (1 - \Delta)$. Thus, it is not affected by a permanent increase in oil prices, and $\hat{r}_t^k(o) = 0$ in the long run.

Second, holding $\hat{w}_t = 0$ is a simplifying assumption that almost perfectly hold in our model and ensures that our elasticity correspond to the the idea of "holding everything else constant" that is often mentioned in other definitions of this elasticity. In our model, the steady-state level of the real wage is given by $w_t = \frac{U'(H_t)}{U'(C_t)}$. This is potentially affected by a permanent increase in oil prices. However, our simulations show that a one percent change in oil prices only has a limited effect on hours worked and consumption, and thus also has a very small impact on the equilibrium wage. The change in wages is one order of magnitude weaker than the change in oil prices. Thus, assuming that $\hat{w}_t = 0$ simplifies the computation of the long-run elasticity of oil supply at a very low cost.

A.3 Imposing a zero elasticity of oil supply on impact

We show how we can impose a zero elasticity of oil supply on impact (one month) while leaving this elasticity unchanged over all other horizons. We simply assume that hours worked are set one period in advance. Under this assumption, all production factors are fixed over a one-month horizon and the supply elasticity is thus zero. Starting in the second month after a shock, the elasticity is identical to that of the baseline model. Specifically, the production function is:

$$Y_t(o) = \varepsilon_t^{os} \left[(1 - \alpha_o - \gamma_o) \left(\frac{H_{t-1}(o)}{H(o)} \right)^{\frac{\theta_o - 1}{\theta_o}} + \alpha_o \left(\frac{K_t(o)}{K(o)} \right)^{\frac{\theta_o - 1}{\theta_o}} + \gamma_o \left(\frac{L_t(o)}{L(o)} \right)^{\frac{\theta_o - 1}{\theta_o}} \right]^{\frac{\theta_o}{\theta_o - 1}} \quad (19)$$

Oil producers maximize their profits given by

$$p_t(o)Y_t(o) - w_t H_{t-1}(o) - r_t^k(o)K_t(o)$$

The first order conditions w.r.t. hours worked is $E_t p_{t+1}(o) \frac{\partial Y_{t+1}(o)}{\partial H_t(o)} = E_t w_{t+1}$. Following the same route as in [A.1](#) but additionally assuming that wages remain fixed for two months, we

end up with:

$$\hat{Y}_t(o) = 0 \quad (20)$$

$$E_t \hat{Y}_{t+1}(o) = \frac{\theta_o(1 - \alpha_o - \gamma_o)}{\alpha_o + \gamma_o} E_t \hat{p}_{t+1}(o) \quad (21)$$

Where the impact (one month) elasticity is zero while the elasticity remains the same as that of the baseline model from the second month onwards.

A.4 The price and output elasticities of oil demand

Final good producers maximize their profits given by

$$Y_t(f) - w_t H_t(f) - r_t^k(f) K_t(f) - p_t(o) O_t(f)$$

where the final good is used as a numeraire to express all prices in real terms and $p_t(o)$ is the real price of oil (deflated with the price of the final good). The first order condition w.r.t. oil inputs is $\frac{\partial Y_t(f)}{\partial O_t(f)} = p_t(o)$. Developing gives:

$$p_t(o) = \omega_o \left(\frac{O(f) - \bar{O}}{\varepsilon_t^{od} Y(f)} \right)^{\frac{1-\eta_o}{\eta_o}} \left(\frac{O_t(f) - \bar{O}}{Y_t(f)} \right)^{-\frac{1}{\eta_o}} \quad (22)$$

We then assume that $p(o) = 1$ at steady state, which simply reflects a choice of unit. It implies that $\omega_o = \frac{O_t(f) - \bar{O}}{Y_t(f)}$. Plug and isolate $O_t(f)$:

$$O_t(f) = \bar{O} + \omega_o Y_t(f) p_t(o)^{-\eta_o} (\varepsilon_t^{od})^{\eta_o - 1} \quad (23)$$

Linearize to obtain the oil demand curve:

$$\hat{O}_t(f) = \frac{O(f) - \bar{O}}{O(f)} \left[\hat{Y}_t(f) - \eta_o \hat{p}_t(o) - (1 - \eta_o) \hat{\varepsilon}_t^{od} \right] \quad (24)$$

To properly introduce prior knowledge on the output and price elasticities of oil demand and to simplify their estimation, we express equation (24) as

$$\hat{O}_t(f) = \delta(y) \hat{Y}_t(f) - \delta(p) \hat{p}_t(o) - \hat{\varepsilon}_t^{od} \quad (25)$$

where $\delta(y) = \frac{O(f) - \bar{O}}{O(f)}$ is the output elasticity of oil demand and $\delta(p) = \eta_o \frac{O(f) - \bar{O}}{O(f)}$ is the price elasticity of oil demand. It implies that the structural parameters are retrieved from these

elasticities as $\bar{O} = (1 - \delta(y))O(f)$ and $\eta_o = \frac{\delta(p)}{\delta(y)}$. Also note the size of the shock is rescaled by a factor of $\frac{\delta_y - \delta_p}{\delta_y}$ which has no effect on the estimation.

B More details on the model

B.1 The households's problem

Households maximize their utility given by

$$\mathbf{E}_0^j \sum_{t=0}^{\infty} \beta^t \left[\frac{(C_{j,t} - \mathbf{b}C_{t-1})^{1-\sigma_c}}{1-\sigma_c} - A_h \frac{(H_{j,t})^{1+\sigma_h}}{1+\sigma_h} \right], \quad (26)$$

where \mathbf{b} is an external habit parameter, under their budget constraint

$$C_t + I_t(f) + I_t(o) + B_{t+1} = R_{t-1}B_t + w_t H_t + r_t^k(f)K_t(f) + r_t^k(o)K_t(o) \quad (27)$$

where R_t is the real gross rate of interest, $H_t = H_t(f) + H_t(o)$, and the capital accumulation rule is

$$K_{t+1} = \left[1 - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 \right] I_t + (1 - \Delta)K_t \quad (28)$$

The FOC w.r.t. C_t , H_t , B_t , K_t and I_t are:

$$(C_t - \mathbf{b}C_{t-1})^{-\sigma_c} = \lambda_t \quad (29)$$

$$A_h H_t^{\sigma_h} = \lambda_t w_t \quad (30)$$

$$\lambda_t = \beta R_t \mathbf{E}_t \lambda_{t+1} \quad (31)$$

$$\mu_t = \beta \mathbf{E}_t (\lambda_{t+1} r_{t+1}^k + (1 - \Delta)\mu_{t+1}) \quad (32)$$

$$\lambda_t = \mu_t \left(1 - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 - \phi \left(\frac{I_t}{I_{t-1}} - 1 \right) \frac{I_t}{I_{t-1}} \right) + \beta \mathbf{E}_t \mu_{t+1} \phi \left(\frac{I_{t+1}}{I_t} - 1 \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \quad (33)$$

where λ and μ are the shadow values associated with constraints (27) and (28). Let $P_t^k = \frac{\mu_t}{\lambda_t}$ and linearize:

$$\hat{w}_t = \sigma_h \hat{H}_t + \frac{\sigma_c}{1-b} \left(\hat{C}_t - b \hat{C}_{t-1} \right) \quad (34)$$

$$\hat{C}_t = \frac{1}{1+b} \hat{C}_{t+1} + \frac{b}{1+b} \hat{C}_{t-1} - \frac{1-b}{\sigma_c(1+b)} \hat{R}_t \quad (35)$$

$$\hat{p}_t^k = (1 - \beta(1 - \Delta)) \hat{r}_{t+1}^k + \beta(1 - \Delta) \hat{p}_{t+1}^k - \hat{R}_t \quad (36)$$

$$\hat{I}_t = \frac{1}{1+\beta} \hat{I}_{t-1} + \frac{\beta}{1+\beta} \hat{I}_{t+1} + \frac{1}{(1+\beta)\phi} \hat{p}_t^k \quad (37)$$

B.2 Calibration and estimation of standard parameters

Other estimated parameters For the investment adjustment cost, the variable utilization rate of capital, we use the priors of [Smets and Wouters \(2007\)](#) (SW07). For shocks' variance and autoregressive parameters, we use standard priors. Specifically, we set the mean of prior of the autoregressive parameters to $0.8^{(1/3)}$ to match a value often used when working with quarterly data. We set the prior mean of the habit parameter to 0.85, resulting in a similar persistence that that assumed by SW07 (0.7 with quarterly data). The share of oil in the production basket is calibrated to match the total oil market value to the world industrial production.

C Additional results

C.1 All other estimated parameters

References

- Abel, A.B., 1990. Asset Prices under Habit Formation and Catching Up with the Joneses. *American Economic Review* 80, 38–42.
- Adjemian, S., Bastani, H., Juillard, M., Karamé, F., Maih, J., Mihoubi, F., Perendia, G., Pfeifer, J., Ratto, M., Villemot, S., 2011. Dynare: Reference Manual Version 4. Dynare Working Papers 1. CEPREMAP.
- Alquist, R., Bhattarai, S., Coibion, O., 2019. Commodity-price comovement and global economic activity. *Journal of Monetary Economics* doi:<https://doi.org/10.1016/j.jmoneco.2019.02.004>.
- Andrle, M., 2010. A note on identification patterns in DSGE models. Working Paper Series 1235. European Central Bank.
- Antolín-Díaz, J., Rubio-Ramírez, J.F., 2018. Narrative sign restrictions for svars. *American Economic Review* 108, 2802–29. doi:[10.1257/aer.20161852](https://doi.org/10.1257/aer.20161852).
- Arezki, R., Blanchard, O., 2014. Seven Questions About The Recent Oil Price Slump. Technical Report. IMF Blog.
- Balke, N.S., Brown, S.P.A., Yücel, M.K., 2010. Oil price shocks and U.S. economic activity: an international perspective. Working Papers 1003. Federal Reserve Bank of Dallas.
- Barsky, R.B., Kilian, L., 2001. Do we really know that oil caused the great stagflation? a monetary alternative. *NBER Macroeconomics Annual* 16, 137–183. doi:[10.1086/654439](https://doi.org/10.1086/654439).
- Barsky, R.B., Kilian, L., 2002. Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative, in: *NBER Macroeconomics Annual 2001*, Volume 16. National Bureau of Economic Research, Inc. NBER Chapters, pp. 137–198.
- 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. doi:[10.1257/aer.20151569](https://doi.org/10.1257/aer.20151569).
- Baumeister, C., Kilian, L., 2016. Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives* 30, 139–60. doi:[10.1257/jep.30.1.139](https://doi.org/10.1257/jep.30.1.139).
- Baumeister, C., Korobilis, D., Lee, T.K., 2022. Energy Markets and Global Economic Conditions. *The Review of Economics and Statistics* 104, 828–844. doi:[10.1162/rest_a_00977](https://doi.org/10.1162/rest_a_00977).
- Baumeister, C., Peersman, G., 2013a. The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics* 28, 1087–1109. doi:[10.1002/jae.2283](https://doi.org/10.1002/jae.2283).
- Baumeister, C., Peersman, G., 2013b. Time-varying effects of oil supply shocks on the us economy. *American Economic Journal: Macroeconomics* 5, 1–28. doi:[10.1257/mac.5.4.1](https://doi.org/10.1257/mac.5.4.1).
- Baumeister, C., Peersman, G., Robays, I.V., 2010. The Economic Consequences of Oil Shocks: Differences across Countries and Time, in: Fry, R., Jones, C., Kent, C. (Eds.), *Inflation in an Era of Relative Price Shocks*. Reserve Bank of Australia. RBA Annual Conference Volume (Discontinued).
- Bergholt, D., Larsen, V.H., Seneca, M., 2019. Business cycles in an oil economy. *Journal of International Money and Finance* 96, 283–303. doi:[10.1016/j.jimonfin.2017.0](https://doi.org/10.1016/j.jimonfin.2017.0).
- Bodenstein, M., Erceg, C.J., Guerrieri, L., 2011. Oil shocks and external adjustment. *Journal of International Economics* 83, 168–184. doi:[10.1016/j.jinteco.2010.10.006](https://doi.org/10.1016/j.jinteco.2010.10.006).
- Braun, R., 2023. The importance of supply and demand for oil prices: evidence from non-Gaussianity. *Quantitative Economics* 14, 1163–1198. doi:<https://doi.org/10.3982/QE2091>.
- Caldara, D., Cavallo, M., Iacoviello, M., 2019. Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics* 103, 1–20. doi:<https://doi.org/10.1016/j.jmoneco.2018.08.004>.
- Cantore, C., Levine, P., 2012. Getting normalization right: Dealing with "dimensional constants" in macroeconomics. *Journal of Economic Dynamics and Control* 36, 1931–1949. doi:[10.1016/j.jedc.2012.05.009](https://doi.org/10.1016/j.jedc.2012.05.009).

- Chang, B., Kang, S.J., Jung, T.Y., 2019. Price and Output Elasticities of Energy Demand for Industrial Sectors in OECD Countries. *Sustainability* 11, 1–17. doi:doi.org/10.3390/su11061786.
- Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy* 113, 1–45. doi:[10.1086/426038](https://doi.org/10.1086/426038).
- Cross, J.L., Nguyen, B.H., Tran, T.D., 2022. The role of precautionary and speculative demand in the global market for crude oil. *Journal of Applied Econometrics* 37, 882–895. doi:<https://doi.org/10.1002/jae.2905>.
- Davis, J.M., 2007. News and the Term Structure in General Equilibrium.
- DeJong, D., Ingram, B., Whiteman, C., 2000. A Bayesian approach to dynamic macroeconomics. *Journal of Econometrics* 98, 203–223. doi:[10.1016/S0304-4076\(00\)00019-1](https://doi.org/10.1016/S0304-4076(00)00019-1).
- Francis, N., Ramey, V.A., 2005. Is the technology-driven real business cycle hypothesis dead? shocks and aggregate fluctuations revisited. *Journal of Monetary Economics* 52, 1379–1399. doi:<https://doi.org/10.1016/j.jmoneco.2004.08.009>.
- Fujiwara, I., Hirose, Y., Shintani, M., 2011. Can news be a major source of aggregate fluctuations? a bayesian dsge approach. *Journal of Money, Credit and Banking* 43, 1–29. doi:<https://doi.org/10.1111/j.1538-4616.2010.00363.x>.
- Golombek, R., Irarrazabal, A.A., Ma, L., 2018. Opec’s market power: An empirical dominant firm model for the oil market. *Energy Economics* 70, 98–115. doi:<https://doi.org/10.1016/j.eneco.2017.11.009>.
- Greene, D.L., Leiby, P.N., 2006. The oil security metrics model. Technical Report. Oak Ridge National Laboratory.
- Hamilton, J.D., 2009. Causes and Consequences of the Oil Shock of 2007–08. *Brookings Papers on Economic Activity* 2009, 215–261.
- Havranek, T., Kokes, O., 2015. Income elasticity of gasoline demand: A meta-analysis. *Energy Economics* 47, 77–86. doi:<https://doi.org/10.1016/j.eneco.2014.11.004>.
- Herrera, A.M., Rangaraju, S.K., 2020. The effect of oil supply shocks on us economic activity: What have we learned? *Journal of Applied Econometrics* 35, 141–159. doi:<https://doi.org/10.1002/jae.2735>.
- Houssa, R., Mohimont, J., Otrok, C., 2023. Commodity exports, financial frictions, and international spillovers. *European Economic Review* 158, 104465. doi:<https://doi.org/10.1016/j.euroecorev.2023.104465>.
- Inoue, A., Kilian, L., 2020. The Role of the Prior in Estimating VAR Models with Sign Restrictions. Working Papers 2030. Federal Reserve Bank of Dallas. doi:[10.24149/wp2030](https://doi.org/10.24149/wp2030).
- Iskrev, N., 2010. Local identification in DSGE models. *Journal of Monetary Economics* 57, 189–202. doi:[10.1016/j.jmoneco.2009.12.007](https://doi.org/10.1016/j.jmoneco.2009.12.007).
- Juvenal, L., Petrella, I., 2015. Speculation in the oil market. *Journal of Applied Econometrics* 30, 621–649. doi:<https://doi.org/10.1002/jae.2388>.
- Kilian, L., 2009. Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review* 99, 1053–1069. doi:[10.1257/aer.99.3.1053](https://doi.org/10.1257/aer.99.3.1053).
- Kilian, L., 2022. Understanding the estimation of oil demand and oil supply elasticities. *Energy Economics* 107, 105844. doi:<https://doi.org/10.1016/j.eneco.2022.105844>.
- Kilian, L., Murphy, D., 2012. WHY AGNOSTIC SIGN RESTRICTIONS ARE NOT ENOUGH: UNDERSTANDING THE DYNAMICS OF OIL MARKET VAR MODELS. *Journal of the European Economic Association* 10, 1166–1188.
- 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. doi:<https://doi.org/10.1002/jae.2322>.
- Kilian, L., Zhou, X., 2018. Modeling fluctuations in the global demand for commodities. *Journal of International Money and Finance* 88, 54–78. doi:<https://doi.org/10.1016/j.jimonfin.2018.07.001>.

- King, R.G., Rebelo, S.T., 2000. Resuscitating Real Business Cycles. Working Paper 7534. National Bureau of Economic Research. doi:[10.3386/w7534](https://doi.org/10.3386/w7534).
- Kose, M.A., 2002. Explaining business cycles in small open economies: 'How much do world prices matter?'. *Journal of International Economics* 56, 299–327. doi:[10.1016/S0022-1996\(01\)00120-9](https://doi.org/10.1016/S0022-1996(01)00120-9).
- Krichene, N., 2002. World crude oil and natural gas: a demand and supply model. *Energy Economics* 24, 557–576. doi:[https://doi.org/10.1016/S0140-9883\(02\)00061-0](https://doi.org/10.1016/S0140-9883(02)00061-0).
- Lippi, F., Nobili, A., 2012. Oil and the macroeconomy: A quantitative structural analysis. *Journal of the European Economic Association* 10, 1059–1083. doi:<https://doi.org/10.1111/j.1542-4774.2012.01079.x>.
- Lütkepohl, H., Netšunajev, A., 2014. Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural vars. *Journal of Applied Econometrics* 29, 479–496. doi:<https://doi.org/10.1002/jae.2330>.
- Nakov, A., Pescatori, A., 2010. Monetary Policy Trade-Offs with a Dominant Oil Producer. *Journal of Money, Credit and Banking* 42, 1–32.
- Otrok, C., 2001. On measuring the welfare cost of business cycles. *Journal of Monetary Economics* 47, 61–92. doi:[10.1016/S0304-3932\(00\)00052-0](https://doi.org/10.1016/S0304-3932(00)00052-0).
- Peersman, G., van Robays, I., 2009. Oil and the Euro area economy [Labour market implications of EU product market integration]. *Economic Policy* 24, 603–651.
- Peersman, G., Stevens, A., 2013. Analyzing Oil Demand and Supply Shocks in an Estimated DSGE-Model. Ghent University.
- Schmitt-Grohé, S., Uribe, M., 2012. What's news in business cycles. *Econometrica* 80, 2733–2764. doi:<https://doi.org/10.3982/ECTA8050>.
- Schorfheide, F., 2000. Loss function-based evaluation of DSGE models. *Journal of Applied Econometrics* 15, 645–670. doi:[10.1002/jae.582](https://doi.org/10.1002/jae.582).
- Smets, F., Wouters, R., 2007. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review* 97, 586–606. doi:[10.1257/aer.97.3.586](https://doi.org/10.1257/aer.97.3.586).
- Taghizadeh Hesary, F., Yoshino, N., 2014. Monetary policies and oil price determination: an empirical analysis. *OPEC Energy Review* 38, 1–20. doi:<https://doi.org/10.1111/opec.12021>.
- Temple, J., 2012. The calibration of CES production functions. *Journal of Macroeconomics* 34, 294–303. doi:[10.1016/j.jmacro.2011.12.006](https://doi.org/10.1016/j.jmacro.2011.12.006).
- Unalmis, D., Unalmis, I., Unsal, D.F., 2012. On the Sources and Consequences of Oil Price Shocks : The Role of Storage. Working Papers 1230. Research and Monetary Policy Department, Central Bank of the Republic of Turkey.
- Zhou, X., 2020. Refining the workhorse oil market model. *Journal of Applied Econometrics* 35, 130–140. doi:<https://doi.org/10.1002/jae.2743>.