



EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Claudia Foroni, Livio Stracca **Much ado about nothing?
The shale oil revolution and the
global supply curve**

Revised December 2019

No 2309 / August 2019

Abstract

We focus on the implications of the shale oil boom for the global supply of oil. We begin with a stylized model with two producers, one facing low production costs and one higher production costs but potentially lower adjustment costs, competing *à la Stackelberg*. We find that the supply function is flatter for the high cost producer, and that the supply function for shale oil producers becomes more responsive to demand shocks when adjustment costs decline. On the empirical side, we apply an instrumental variable approach using estimates of demand-driven oil price changes derived from a standard structural VAR of the oil market. A main finding is that global oil supply is rather vertical, practically all the time. Moreover, for the global oil market as a whole, we do not find evidence of a major shift to a more price elastic supply as a result of the shale oil boom.

Keywords: Oil supply, shale oil, oil shocks, structural VAR, instrumental variables, sign restrictions.

JEL Classification Code: Q33, Q41, Q43, C32

Non-technical summary

This paper focuses on the question of whether the world oil supply function has changed in a fundamental way with the advent of the shale oil on the market. A change in the elasticity of the oil supply function can have, in fact, potential consequences for oil prices and broader macroeconomic consequences given oil's still central role in the world economy.

Our contribution is twofold. On the conceptual side, we propose a very stylized model for the oil market with two producers, one representing conventional oil producers and one shale oil producers, we find that the supply function is flatter for the high cost producer (shale). Further, we find that the supply function for shale oil producers becomes more responsive to demand shocks when adjustment costs decline. On the other hand, more efficient production by the shale oil producer actually reduces its production elasticity.

Second, we conduct an empirical analysis based on a combination of structural VAR analysis and instrumental variables (IV). We first identify demand shocks influencing the price of oil using a more parsimonious Bayesian version of the standard structural VAR approach in Kilian and Murphy (2014). We then use these shocks as instruments in an IV estimate of the oil supply function, for the world as well as the main producers (OPEC and non-OPEC). Finally, we run time-varying IV estimates to capture possible structural breaks in the supply function, in particular before and after the boom in shale oil production around 2010.

We find that the oil supply curve is typically relatively vertical. At the same time, we find a flattening of the supply curve for the U.S., in particular for oil demand shocks around 2010. Taking our model as the guide to analyze the empirical evidence, we find that this is consistent with the idea that a fall in adjustment costs has been a major factor for shale oil production, rather than an outright reduction in marginal production costs. For the global oil market as a whole, however, we do not find evidence of a major shift to a flatter supply curve around 2010. This may be explained by the fact, its impressive growth notwithstanding, US shale oil production only represents a relatively small fraction of world oil production.

1 Introduction

Since the summer of 2014, the oil industry has been on the roller-coaster, with the oil price falling to 30 dollars per barrel from above 100, before recovering to the level of 50 dollars per barrel since mid 2016, and increasing again to 80 dollars in mid 2018. In the meanwhile, we have seen the OPEC first implicitly trying to fight the shale oil by dragging the shale oil companies into unprofitable territories, then agreeing in cutting the production of oil to sustain the price.

While the financial press and policy institutions are evoking the shale oil as a major change in the oil market, and the reason behind the plunge in the oil price, there is surprisingly still few studies which look at the shale phenomenon in the economic literature. Kilian (2017b) aims at quantifying the impact of the U.S. shale oil boom on U.S. imports and exports of crude and refined oil. Further, he aims at estimating the effects of a change in the balance trade of the U.S. on the Arab countries, and in particular on Saudi Arabia. Based on a structural VAR model, he shows that the U.S. shale oil revolution is not the cause of the 2014/15 decline in the Brent oil price. He furthermore quantifies the cumulative effects of the decline in the Brent price caused by the shale oil revolution on the foreign exchange revenue of Saudi Arabia. Kilian (2016) examines how the increased availability of shale oil has shaped the U.S. oil and gasoline prices. A broader review on how the shale oil affected the U.S. economy is provided by Kilian (2017a). Mohaddes and Pesaran (2016) take a different methodological perspective and look at the effects of lower oil prices to the global economy, with the use a multi-country GVAR model. With a similar GVAR approach, Mohaddes and Raissi (2015) investigate more specifically the macroeconomic consequences of the U.S. oil revolution for the global economy in general and the Middle East and North Africa in particular. Kilian (2009) was the first paper to attribute the sluggishness of the supply response to the costs of adjusting oil production and the uncertainty about the state of the crude oil market (p. 1059). This observation was formalized by Anderson et al. (2018), who showed that in equilibrium the price elasticity of oil supply should be zero, given the high costs of shutting down and reopening conventional oil wells. In this case, the optimal response of oil producers to an oil price change induced by oil demand shifts is to adjust investment in future oil production rather than the level of oil production from existing wells. This theoretical prediction is supported by extraneous micro evidence. Although there are no micro-economic estimates of the global one-month price elasticity of oil supply, recent micro-economic estimates based on regional data from the United States are

all close to zero and statistically insignificant, consistent with economic theory (see, e.g., Anderson et al. (2018)). Recently, Bjornland et al. (2017) analyse the extent to which shale oil producers respond to price incentives by changing either the completion of new wells or the production from existing ones. Using data from the single oil wells in North Dakota, they find large differences in the supply function of conventional and unconventional oil technology (although their results are not strongly statistically significant, as highlighted by Kilian and Zhou (2018)).

The goal of our paper is to understand whether there has been a change in the oil market, induced by the shale technology. In particular, we want to check whether with the advent of the shale oil, the oil supply has become more responsive¹. In other words, while a few studies in the literature typically found oil production to be inelastic in the short run, we wonder whether this is still true, especially after the shale oil boom. If shale oil production is more elastic than conventional oil production, this can have potential consequences for the evolution of the price of oil.

To conduct our analysis, we start with a stylized model for the oil market, in which we have two producers, one facing low costs but with very high adjustment costs (representing conventional oil producers), and one facing a higher cost, but lower adjustment costs (representing shale oil producers). The two players compete *à la Stackelberg*: they produce a homogeneous good, they do not collude, but only Saudi Arabia has market power and internalises the demand constraint in its choice. They compete in quantities and choose them simultaneously. Further, they act rationally and strategically, aiming at maximizing their own profits. This is in keeping with most of the literature, which considers typically a dominant firm with a competitive fringe model (see, for example, Nakov and Nuño (2013) and Golombek et al. (2018)).

Based on simple but realistic calibrations of our model, we find that the supply function is flatter for the high cost producer. This reflects the fact that a positive demand shock pushes production in the area that makes it profitable for the high cost producer, whereas the low cost producer finds it profitable even when demand is scarce. Further, we find that the supply function for shale oil producers becomes more responsive to demand shocks when adjustment costs decline, while we find the opposite for the case of a lower marginal production cost. The intuition for the latter result is that the decline in production cost

¹For simplicity, we will talk throughout the paper about oil supply curve. More precisely, we refer to the short-run oil supply curve conditional on past data. The estimated model is shown in equation (7).

makes the high cost producer more similar to the low cost producer (Shale closer to Saudi), and hence flattens the supply curve. Note that our model is stylized and aims at capturing the role of adjustment and marginal costs, more than modelling realistically a complex market as Saudi Arabia or the OPEC. Therefore, our model is much more simplified than the one by Nakov and Nuño (2013), who look at the behaviour of Saudi Arabia as dominant producer with a competitive fringe in a general equilibrium model.

We conduct then an empirical analysis, in which we try to analyze a potential structural change in the oil production, as the advent of the shale oil on the market, which can be thought as a major technology change. In order to do that, we estimate a VAR model similar to the one of Kilian and Murphy (2014), and derive estimates of oil demand shocks. Then we use these shocks as proxies of movements of oil prices around the oil supply curve, and as instruments to recover the underlying oil supply function. A recent paper by Caldara et al. (2016) is focusing also on estimating supply and demand elasticities in the oil market for a selected number of countries, combining narrative analysis with a panel of observations on country-specific oil production and consumption to estimate oil supply and demand elasticities.

Contrary to Caldara et al. (2016), who look only on the full sample period for their analysis, we focus on whether there is *time variation* or not in the supply elasticities, investigating whether any technological break or other major changes have occurred in the oil market after around 2010, the time when US shale oil production started to soar. In addition, we do not limit ourselves to consider the global supply, but we will look also at the supply of the OPEC, and of Saudi Arabia alone, as representative of the conventional oil producers, as well as at the supply on non-OPEC and the two main players in that group, mainly, Russia and the US.

Overall, we find that the oil supply curve is relatively price inelastic, in line with most previous findings in the literature. At the same time, we find a flattening (increase in price responsiveness) of the supply curve for the U.S. in the post 2010 sample (although not in the short run elasticity). In the language of our stylised model, this is consistent with the idea that a fall in adjustment costs has been a major factor for Shale oil production, rather than an outright reduction in marginal production costs. For the global oil market as a whole, however, we do not find evidence of a major shift to a flatter supply curve around 2010 as it is sometimes argued. This reflects the fact that, despite its impressive growth, US shale oil production only represents a relatively small fraction of world oil production, and time will tell if this will change over time as technological advances continue.

The paper is organized as follows. In Section 2 we present our stylized model for the oil market, and discuss the consequences for the oil supply function when different calibrations are chosen. In Section 3 we describe the data used in the analysis. In Section 4 we describe the identification strategy for the oil supply curve. In Section 5 we present our empirical findings. Section 6 concludes.

2 A model with two producers

We introduce a model of the oil market which features two producers, one facing lower production costs (henceforth Saudi Arabia, or SA), and one facing a higher cost, but potentially lower adjustment costs (henceforth Shale, or SH). The model is in real terms, i.e. all the variables are real.

Our intent is to construct a small model to highlight mainly the role of marginal and adjustment shocks in the optimal supply of oil and guide the empirical analysis. It is therefore far less rich than the model introduced by Nakov and Nuño (2013) to describe the behaviour of Saudi Arabia as a dominant producer with competitive fringe. Further, that model is constructed into a dynamic stochastic general equilibrium setting, while ours is a simplified, partial equilibrium one. Similar to Nakov and Nuño (2013) and other papers in the literature we choose to model the market with only one strategic player, Saudi Arabia.²

In addition, note that we do not consider the fact that oil is an exhaustible resource (Hotelling rule) as most of the existing literature considers this dimension not to be relevant for producers decisions (see, e.g., (Anderson et al. (2018) and Cairns and Calfucura (2012)). Likewise, we do not consider the question of how OPEC coordinates its actions and how it behaves like a cartel or not; on that question see for example Almoguera et al. (2011) and Alkhatlan et al. (2014). We consider one block of low cost producers as acting as price takers. Unlike Behar and Ritz (2016) we do not focus on the OPEC or Saudi decision to accommodate or squeeze out shale oil out of the market. Moreover, our modelling strategy is different, and notably the marginal cost of production is an endogenous variable in our setting, while it is a parameter in Behar and Ritz (2016).

²A previous version of this paper has a Cournot model and results are similar. For a comparison between the Cournot and Stackelberg models, see Anderson and Engers (1992). While in a Stackelberg model firms choose output sequentially, in the Cournot model they choose it simultaneously, without knowing the competitors' choice.

2.1 Setting

We consider a standard linear demand function:

$$P_t = a - dQ_t + \epsilon_t, \quad (1)$$

where Q_t is the total quantity of oil produced in the market, and P_t is its price.

The two producers produce a homogeneous good, do not collude, but only Saudi Arabia has market power (i.e., internalises the demand constraint in its choices and does not take the oil price as given). They compete in quantities and choose them simultaneously. Further, they act rationally and strategically, with the only objective to maximize their profits.

Each producer faces its own cost function $C_{i,t}$:

$$C_{i,t} = \alpha_i Q_{i,t}^{\delta_i} P_{i,t}^{\eta_i} + \frac{\gamma_i}{2} (Q_{i,t} - Q_{i,t-1})^2 \quad (2)$$

where $i = SA, SH$. Costs follow a power function and adjustment costs to production are assumed to be quadratic.³ The term $P_{i,t}^{\eta_i}$ captures the fact that production costs may be a function of the oil price, which is particularly plausible for Shale producers due to the cost of rigs.⁴ In the calibration, we set $\eta_i > 0$ for Shale but $\eta_i = 0$ for Saudi Arabia.

Further, $\alpha_{SA} < \alpha_{SH}$ and $\delta_{SA} < \delta_{SH}$ but $\gamma_{SA} = \gamma_{SH}$. Hence, we assume that Saudi Arabia has lower variable and fixed production costs and, in the baseline, the same adjustment costs as Shale. In the calibration, we also consider a version where Shale has lower adjustment costs than Saudi Arabia, and also where it has lower marginal production costs than in the baseline, while still higher than Saudi Arabia. In both cases, oil production hits diminishing marginal returns (that is, there are increasing production costs); this implies that $\delta_i > 1$.

³We are therefore assuming that the adjustment costs are symmetric, which may not be the case in reality (in particular, reducing production may be easier than increasing it). We are working at an extension which allows for asymmetry.

⁴We thank Axel Pierru for this suggestion.

2.2 Equilibrium

Each producer chooses the quantity to produce Q_i to maximize its profits. That is, each of the two producers solve the following maximization problem:

$$\max \sum_{t=0}^{\infty} \beta^t \left\{ P_t Q_{i,t} - \alpha_i Q_{i,t}^{\delta_i} P_{i,t}^{\eta_i} - \frac{\gamma_i}{2} (Q_{i,t} - Q_{i,t-1})^2 \right\} \quad (3)$$

with $i = SA, SH$, subject to the demand function in equation (1) for the Stackelberg leader (Saudi Arabia) and the market clearing condition:

$$Q_t = Q_{SA,t} + Q_{SH,t}. \quad (4)$$

This leads to the following first order condition for Saudi Arabia:

$$P_t - dQ_{SA,t} - \alpha_{SA} \delta_{SA} Q_{SA,t}^{\delta_{SA}-1} P_t^{\eta_{SA}} - \gamma_{SA} (Q_{SA,t} - Q_{SA,t-1}) + \beta \gamma_{SA} (E_t(Q_{SA,t+1}) - Q_{SA,t}) = 0, \quad (5)$$

For Shale, the same FOC applies apart from the term $-dQ$, because the demand constraint is not internalised by this producer, who acts as a price-taker.

Equation (5) for $i = SA, SH$, together with equations (1) and (4), determine the equilibrium of the model.

2.3 Model calibration

In the calibration of the slope of the demand function we follow Behar and Ritz (2016) and impose a slope parameter of -8.⁵ Production is expressed in million barrels per day.

Our cost calibration, when setting adjustment costs at zero, ensures that the low cost player (Saudi Arabia) has a marginal production cost of 15 in USD per barrel at the steady state, while the high cost producer (Shale) faces a marginal cost of 50 (again USD per barrel). Further, adjustment costs are set at 1000 for both Saudi Arabia and Shale, as this value gives a reasonable degree of persistence in the impulse responses and a realistic value of the short term supply elasticity.

⁵Observe that they, unlike us, define the demand function as $Q = f(P)$.

Table 1: Calibrated parameters

Parameter	Baseline	Calibration 1: Lower adjustment cost in Shale	Calibration 2: Lower marginal cost in Shale
d	-8	-8	-8
α_{SA}	10	10	10
α_{SH}	15	15	12
δ_{SA}	1.15	1.15	1.15
δ_{SH}	1.25	1.25	1.20
γ_{SA}	1000	1000	1000
γ_{SH}	1000	500	1000
η_{SA}	0	0	0
η_{SH}	0.1	0.1	0.1

Given the parameters chosen for the calibration, we obtain a steady state of the model with the oil price at 50 and total production at 15 million barrels per day, which is approximately the sum of Saudi and Shale production in the data. What the model gets wrong, however, is the distribution of production between Saudi and Shale, where in the steady state it is Shale the dominant producer (10.5 million barrel vs. 4.4 million barrel for Saudi), owing to the fact that only the Saudi producer internalizes the demand constraint, which leads it to produce less than it otherwise would.⁶ We also propose two alternative calibrations aimed at understanding the changing role of Shale producers. In a first alternative calibration, we consider a lower adjustment cost for Shale, capturing the idea of Shale production becoming more flexible. In the second alternative, we make Shale production more efficient and closer to the Saudi producer. All the other parameters remain as in the baseline. The calibration of our baseline model and of our two alternatives is summarized in Table 1.


In Figure 1 we plot the supply elasticity of Saudi Arabia (left panel), Shale (middle panel) and world (sum of Shale and Saudi; right panel) to a demand shock in our baseline calibration (red solid line), when the adjustment cost for shale oil decreases (dashed blue line), and when marginal costs for Shale decrease (magenta lines with circles).⁷

In general, we find that the supply elasticity is (as expected) positively sloped for both

⁶Indeed, the relative production shares change dramatically if we impose that Saudi is also a price-taker producer like Shale.

⁷Note that the supply elasticity is computed as $\frac{\Delta q}{q^{SS}} / \frac{\Delta P}{P^{SS}}$, where SS denotes the steady state of the variables.

Figure 1: Stackelberg competition model simulations: Supply elasticity following an oil demand shocks



The **red solid line** refers to the baseline calibration, the **blue dashed line** to a calibration with lower adjustment costs for Shale ($\gamma_2 = 500$) and the **magenta line with circles** to a calibration with lower marginal costs for Shale ($\alpha_{SH} = 12, \delta_{SH} = 1.20$). The calibrated parameters are summarized in Table 1.

producers and flatter for the high cost producer. This reflects the fact that a positive demand shock pushes production in the area that makes it profitable for the high cost producer, whereas the low cost producer finds it profitable to produce even when demand is scarce. With this general principle in mind, we find that the supply function for the high cost producer (Shale) becomes flatter (more responsive to demand shocks) when adjustment costs decline. For the case of a lower marginal production cost we find the opposite to be the case, namely that the high cost producer becomes less responsive (the supply curve steepens); see Figure 1. The intuition for the latter result is that the decline in production cost makes the high cost producer more similar to the low cost producer (Shale closer to Saudi), and hence steepens the supply curve.

All in all, our simple model suggests two testable implications for the Shale (high cost) producer:

1. *More flexible production, as captured by lower adjustment costs in the model, should lead to a steeper supply curve, i.e. to a larger reaction of supply to demand-induced price increases;*
2. *More efficient production, as measured by less decreasing returns to scale in oil production, leads to a less steep supply curve.*

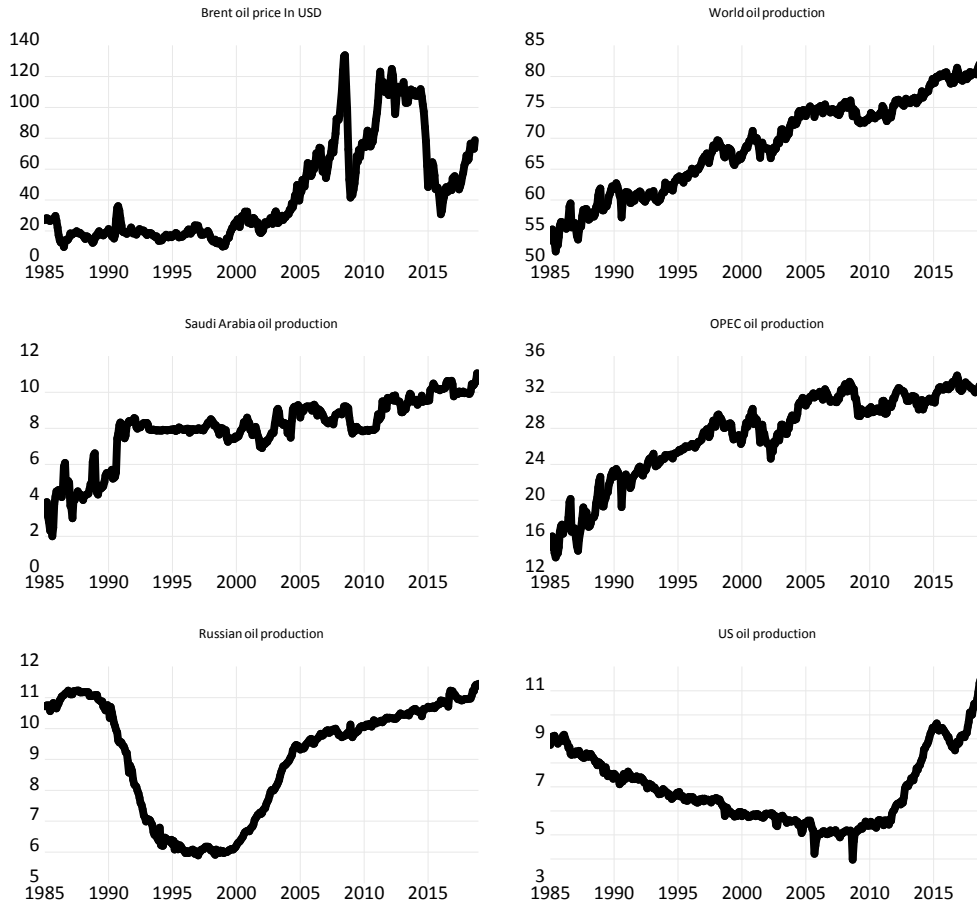
For the Saudi (low cost) producer, the implications are the opposite to those of Shale, but more muted, at least in the calibration that we use.

3 The oil market data

In this section we introduce our data, which are monthly and cover the period from 1985 to 2018. Figure 2 offers an overview of all the variables that we will be using in this paper. A summary on the data sources are provided in Table 2.

Our goal is to try to analyze a potential structural change in the oil production with the advent of the shale oil on the market, which can be thought as a major technology change. Global crude oil production has increased steadily since the oil counter-shock in the mid- to late 1980s, in parallel with the increase of the OPEC production. However, OPEC production has levelled off (or at least it grew much less) starting from 2005. Other

Figure 2: Oil market data.



Note: Data for the Brent oil price are in USD, oil production data are in million barrels per day. See Table 3 on the sources of the data.

Table 2: Data description

Data	Source	Transformation
Global real economic activity	Lutz Kilian's webpage	level
Crude oil production (global and per country)	Energy Information Administration Monthly Energy Review	percentage changes
Brent oil price	Datastream	log levels (deflated by US CPI)
Inventories	Energy Information Administration	changes (obtained as in Kilian and Murphy (2014))

countries were responsible of the oil production increase, in particular the U.S. and Russia experienced the biggest increase in their production. The U.S. crude oil production actually declined until the first decade of the 2000s, and only from 2009 it shows a marked increase, which can be linked to the rise of shale oil production and the arrival of the fracking technology. Russian production declined markedly in the aftermath of the collapse of the Soviet Union, but recovered strongly in the early 2000s and it is now one of the major oil producers.

Whereas oil production has tended to be reasonably stable over time, the oil price (shown in USD in the figure) has fluctuated wildly especially since the early 2000s. In particular, oil prices increased strongly from 2003, driven by the strong growth in economic activity in China and worldwide. This demand driven increase was interrupted in 2008 by the global financial crisis, which caused a drop in the oil price of about 70% in just a few months. This sharp drop was followed by a strong rebound in 2009, and prices remained at record-high levels until 2013, when the rise of shale oil production and the slowdown of the Chinese economy made the price plunge from 130 dollars/barrel to roughly 30 in 2016. At the time of writing, prices have seen a sustained rise and stand at around 75 US dollars per barrel.

4 Empirical identification of the oil supply

In order to identify the oil supply function we need to obtain a source of variability *along* rather than *of* the oil supply due to demand factors. This implies the need of instrumental variables to deal with simultaneity problems. For each country in our analysis, we instrument the price of oil to clean for demand factors. An innovation compared with previous literature is our use of the estimated demand shocks coming from a SVAR that describes the global oil market as instruments to identify the supply curve. In particular, we estimate a VAR model similar to the Kilian and Murphy (2014) model of the oil market in order to derive estimates of oil demand shocks, which we use to identify movements in the real price of oil driven by exogenous demand shifts. In a second stage, we use the estimated shocks as instruments to recover the underlying oil supply function.

The idea to use instrumental variables to identify the oil supply curve is similar in spirit to Caldara et al. (2016), but note that we use a different instrument. In fact, in their paper Caldara et al. (2016) combine narrative analysis with a panel of observations on country-specific oil production and consumption to estimate oil supply and demand elasticities. In

addition, their focus is different: while our objective is to estimate a time-varying oil supply curve, their goal is to recalculate the contributions of oil demand and supply shocks to the fluctuations in oil prices and quantities.

Other papers in the literature have used a series of events to identify shocks to demand and supply, which is close to an instrumental variable approach as in Caldara et al. (2016) and in this paper. For example, Hamilton (2003) suggests measures of exogenous oil production shortfall as instrument, and Kilian (2008) builds on that and proposes a new measure of exogenous supply shocks.

4.1 Deriving demand shocks in the oil market

A standard model in the literature to analyze the global oil market is the one proposed by Kilian and Murphy (2014), which in a sign-restricted VAR framework, allows to identify three types of shocks in the global oil market: shocks to the storage demand for oil, to flow demand and to supply.

Therefore, following Kilian and Murphy (2014) we estimate a monthly model, and derive estimates of oil demand shocks. The model is a VAR estimated on monthly data, including a vector z with the percent change in crude oil production, the Kilian index of real economic activity (see Kilian (2009) for a description of the index), the real price of oil in log-levels (Brent price deflated with the US CPI) and the change of global crude oil inventories above the ground. The representation of the structural VAR model is the following:

$$A_0 z_t = \alpha + \sum_{i=1}^{24} A_i z_{t-i} + \epsilon_t. \quad (6)$$

where ϵ_t is a white noise process with mean zero and covariance matrix Σ_t . The identification of the oil shocks is achieved by imposing the sign restrictions as in Kilian and Murphy (2014) (see Table 3).

Further, as in Kilian and Murphy (2014) we impose a bound on the impact price elasticity of oil demand. In particular, we impose that the impact elasticity of oil demand must be weakly negative on average over the sample, included between -0.8 and 0.⁸ We also impose

⁸These bounds are standard in the literature, see the discussion in Kilian and Murphy (2014) and their references to Sweeney (1984) and Hausman and Newey (1995).

Table 3: Identifying structural shocks in the oil market

	Oil supply	Flow demand	Storage demand
Oil production	-	+	+
Kilian index of global activity	-	+	-
Real oil price	+	+	+
Inventories			+

the additional restriction that the response of the real price of oil to a negative flow supply shock must be positive not only on impact, but on the first 12 periods (to be consistent with the conventional views on unanticipated oil supply disruptions). However, given that the goal of our paper is exactly to look at potential changes in the supply elasticity, we impose only a loose bound on the impact price elasticity of oil supply (corresponding to the loosest one in the Kilian and Murphy (2014)), equal to 0.1.

Further, in order to take into account potential time-variation in the oil supply elasticity around the time of the shale oil revolution, we cannot estimate our SVAR on the full sample, which would imply no change in the structure of the economy. We therefore allow for one-time break in the data in 2010, the year when US shale oil production started to rise markedly. In practice, we estimate our model first on the sample 1985-2009, and then on the sample 2010-2018. Given that the second sample is short, we deviate from the estimation method implemented by Kilian and Murphy (2014) and we use a Bayesian method. A relatively tight Minnesota prior allows us in this case to maintain unchanged the number of lags in the estimation (equal to 24 months). In fact, this type of prior implies that the further the lag, the more confident we should be that coefficients linked to this lag have a value of zero. This means that the variance of the lagged parameters should be smaller on further lags. The same is true for coefficients relating variables to past values of other variables. The variance (i.e. the tightness) of these parameters is driven by the choice of only 4 hyper-parameters. We optimize the hyper-parameters of the prior by using a grid search, and retain the combination of hyper-parameters that maximizes the value of the marginal likelihood for the model. In practice, we estimate the sign-restricted VAR with standard Bayesian algorithms (implemented in the ECB BEAR toolbox), by constructing a set of admissible models that satisfy our restrictions. We then use the median flow demand and median storage demand shocks from this estimation procedure, which we will denote as $\hat{\epsilon}_t^{FlowDemand}$ and $\hat{\epsilon}_t^{StorageDemand}$. Later on, we will verify if using the Bayesian version

of the Kilian-Murphy model results in any loss of information from the point of view of estimating a supply curve.

4.2 Estimating the oil supply curve

We estimate the oil supply curve using local projections as in Jorda (2005). The empirical model is the following:

$$\Delta Prod_{i,t+h} = \alpha_i + \beta_h \Delta Oilp_t + \epsilon_{i,t+h} \quad (7)$$

where $\Delta Prod_{i,t+h}$ is the change in oil production for country i between t and $t+h$, $h = 1, \dots, 12$, $Oilp$ is the log oil price in USD deflated with the US CPI and the parameter β_h indicates the slope of the curve. We consider $h > 0$ in order to allow for lags in producers' reaction to new information, consistent with, e.g., Anderson et al. (2018). One would normally expect this parameter to be positive in a supply function, as it is more convenient to produce more when the price is high. In order to study whether the boom in the shale oil production has influenced oil supply by other producers and globally, we estimate the supply equation in the 1985-2009 and 2010-2018 period separately, and compare them.

The estimation of this equation by OLS may lead to inconsistent results, because supply shocks to oil production (say, a disruption in supply due to a natural disaster or political event) may well influence the oil price, creating a simultaneity bias. For this reason, we instrument $\Delta Oilp_t$ with the two demand shocks, $[\hat{\epsilon}_t^{FlowDemand}, \hat{\epsilon}_t^{StorageDemand}]$.

Caveats and possible questions on the identification. As any identification strategy, also ours is subject to a number of caveats and qualifications. One first concern is that the instrument should be both valid and strong for the instrumentation strategy to be successful. In our case, the instruments are very strong (more evidence on this below). Second, the identification of the demand shocks hinges on the specification of the structural VAR of Kilian and Murphy (2014). Although this model is very well established in the literature, it should be observed that the stability of the model over time is particularly important for us, because we are investigating possible changes in the slope of the oil supply curve over time.⁹

⁹Of course, the strategy of using extraneous information for estimating supply or demand elasticity is not new in the literature; see, e.g., Kilian and Lee (2014).

First stage regressions

	1985-2009	2010-2018
Aggregate demand shock	0.07*** (0.00)	0.04*** (0.00)
Precautionary demand shock	0.05*** (0.00)	0.03*** (0.00)
Observations	300	105
F statistic	748.1	212.0

Table 4: OLS regression. Dependent variable: monthly change in the oil price in USD.

5 Results

Before describing the results in detail, it is useful to provide an overview of the main findings. First, we look at the first stage regressions, and conclude that our choice of instruments appears appropriate as our estimated demand shocks are strong instruments and affect oil prices with the expected sign. Second, we focus on the full sample and confirm previous results in the literature that the supply curve in the short run is relatively vertical, with short term elasticity between 0.02 and 0.04. Finally, has the oil supply curve has changed over time? We find that the shale oil boom around 2010 does not appear to have fundamentally affected the slope of the global oil supply curve, which remains rather vertical. Moreover, the boom appears to have increased the medium term responsiveness of US production to demand-driven changes in oil prices, which in our theoretical model is consistent with lower adjustment costs but not with lower marginal costs of production.

5.1 First stage regressions

We begin by reporting the results of the first stage regressions in Table 4, where we regress the growth rate of the oil price in USD on the estimated aggregate demand and precautionary demand shocks (estimates from the Bayesian version), in the 1985-2009 and 2010-2018 sample periods. In both samples the instruments are correctly signed and strong, as suggested by the F statistic much above the typical benchmark of 10.

5.2 Estimates for the whole sample 1985-2018: a relatively vertical supply

Armed with the instruments, we estimate equation (7) and derive impulse responses from the β_h coefficients. In Figure 3 we compare the estimates based on instruments from

the original Kilian-Murphy VAR model (black lines) and from the Bayesian version (blue lines) on the pre-shale sample, 1985-2009. Two results are noteworthy.

First, the impact elasticity is positive but small, between .02 and 0.04, implying a relatively vertical curve. This implies that a 10% increase in the oil price leads to a rise in oil production by 0.2-0.4%. The effect is relatively persistent but loses statistical significance after one month. Observe that this number is well within the range of the previous findings in the literature. For example, it is well in line with the 0.025 estimated in Kilian and Murphy (2012). Bjornland et al. (2017) find low and insignificant elasticity for conventional production, although they find a high elasticity (between 0.6 and 1) for shale oil wells. Anderson et al. (2018) find that production from existing wells in Texas does not respond to price signals (implying a vertical short term supply curve), although drilling activity does. Caldara et al. (2016) report larger elasticities between 0.05 and 0.1, as do Bornstein et al. (2017), who however use annual data.¹⁰ Finally, it accords well with our stylised theoretical model, where the supply elasticity turned out to be in the neighborhood of 0.02, although this reflected, of course, the chosen calibration in particular for the adjustment costs.

Second, it is noteworthy that the two lines are practically overlapping, signalling that, at least for the purpose of this exercise, the full Kilian-Murphy and our more parsimonious Bayesian version are largely interchangeable at least for the purpose of this analysis.


Overall, based on these two results, we can turn to compare the pre- and post-shale supply curve estimates.

5.3 Did the shale oil boom bring a structural break in oil supply?

Figure 4 compares estimates of the oil supply for the world, Saudi Arabia, OPEC, Russia and the United States in the pre-shale sample 1985-2009 (black lines) and in the post shale one 2010-2018 (blue lines), based on the Bayesian version of the Kilian-Murphy model. The impulse responses are almost always overlapping, implying that the data do not allow us to draw firm conclusions about the impact of the boom in shale oil production on global oil supply. Statistical significance aside, however, one conclusion that seems to emerge from the figure is that the Saudi Arabian and OPEC has become more vertical

¹⁰They also find that OPEC production is less price responsive, which is consistent with our results.

Figure 3: Supply curve estimates 1985-2009: Original Kilian-Murphy model (black lines) vs. Bayesian version (blue lines).



The figure reports impulse responses derived from local projections estimates as in equation (7) estimated with instrumental variables on the sample 1985-2009, where $\Delta Oilp_t$ is instrumented with $[\hat{\epsilon}_t^{Demand}, \hat{\epsilon}_t^{Oildemand}]$. The black lines refers to the Kilian-Murphy original VAR model, the blue lines to the Bayesian version with one lag. The dashed lines indicate confidence bands at the 90% level.

(i.e. production responds less to prices), whereas the US production has flattened (production responds more to prices, especially over time; consistent with, e.g., Bjornland et al. (2017)).¹¹ Because OPEC production is still much larger than US production, overall the world supply appears to have become if anything more price inelastic.

Seeing these results through the lenses of our simple theoretical model, we would reach two main conclusions: (i) the shale oil boom does not appear to have fundamentally changed the contours of global oil production, where the supply function remains relatively vertical; (ii) tentative evidence, although not statistically significant, points to the oil supply curve becoming more vertical in Saudi Arabia, and more price responsive in the US. This in turn is consistent with the change being brought about by a decline in production adjustment costs in Shale, and not by a fall in marginal production costs compared with Saudi.


6 Conclusions

With this paper, we contribute to the understanding of the oil supply function, and in particular to whether it has changed with the advent of the shale oil on the market. A change in the elasticity of the oil supply function can have, in fact, potential consequences for oil prices and broader macroeconomic consequences given oil's still central role in the world economy.

With a very stylized model for the oil market with two producers, one representing conventional oil producers and one shale oil producers, we find that the supply function is steeper for the high cost producer (Shale). Further, we find that the supply function for Shale producers becomes more responsive to demand shocks when adjustment costs decline. These results find confirmation in our empirical analysis. Conducting a structural VAR exercise, we find that the oil supply curve is difficult to estimate, generally stable and not very responsive to fluctuations in oil demand - i.e., the curve is typically vertical. At the same time, we find a flattening of the supply curve for the U.S. after 2010 (albeit not in the short run elasticity). This is consistent with the idea that a fall in adjustment costs has been a major factor for Shale oil production, rather than an outright reduction in marginal production costs. For the global oil market as a whole, however, we do not find evidence of a major shift to a flatter supply curve around 2010. Most likely this finding

¹¹Observe that our data on US production cover both conventional and shale oil.

Figure 4: Supply curve estimates, 1985-2009 (black lines) vs. 2010-2018 (blue lines).



The figure reports impulse responses derived from local projections estimates as in equation (7) estimated with instrumental variables on the sample 1985-2009, where $\Delta Oilp_t$ is instrumented with $[\hat{\epsilon}_t^{Demand}, \hat{\epsilon}_t^{Oildemand}]$. The black lines refer to the VAR model (Bayesian version) estimated on 1985-2009, the blue lines to 2010-2018. Dashed lines indicate confidence bands at the 90% level.

reflects the fact that, despite its impressive growth, US shale oil production only represents a relatively small fraction of world oil production.

The rise in shale oil production is a relatively recent phenomenon and it is still early to make predictions about its effect on the global oil market. This paper should be seen as a first step, but its conclusions should be revisited at a later point in time as technological advances in fracking technology continue and further evidence accumulates.

References

- Alkathlan, K., Gately, D., and Javid, M. (2014). Analysis of saudi arabia's behavior within opec and the world oil market. *Energy Policy*, 64:209 – 225.
- Almoguera, P. A., Douglas, C. C., and Herrera, A. M. (2011). Testing for the cartel in opec: non-cooperative collusion or just non-cooperative? *Oxford Review of Economic Policy*, 27(1):144–168.
- Anderson, S. P. and Engers, M. (1992). Stackelberg versus cournot oligopoly equilibrium. *International Journal of Industrial Organization*, 10:127–135.
- Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3):984–1026.
- Behar, A. and Ritz, R. A. (2016). An Analysis of OPECs Strategic Actions, US Shale Growth and the 2014 Oil Price Crash. IMF Working Papers 16/131, IMF.
- Bjornland, H. C., Nordvik, F. M., and Rohrer, M. (2017). Supply Flexibility in the Shale Patch: Evidence from North Dakota. Working Papers No 2/2017, Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School.
- Bornstein, G., Krusell, P., and Rebelo, S. (2017). Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry. NBER Working Papers 23423, National Bureau of Economic Research, Inc.
- Cairns, R. D. and Calfucura, E. (2012). OPEC: Market failure or power failure? *Energy Policy*, 50(C):570–580.
- Caldara, D., Cavallo, M., and Iacoviello, M. (2016). Oil Price Elasticities and Oil Price Fluctuations. International Finance Discussion Papers 1173, Board of Governors of the Federal Reserve System (U.S.).
- Golombek, R., Irarrazabal, A. A., and Ma, L. (2018). Opec's market power: An empirical dominant firm model for the oil market. *Energy Economics*, 70:98 – 115.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2):363–398.
- Hausman, J. and Newey, W. (1995). Nonparametric estimation of exact consumers surplus and deadweight loss. *Econometrica*, 63(6):1445–76.

- Jorda, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.
- Kilian, L. (2008). Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *The Review of Economics and Statistics*, 90(2):216–240.
- Kilian, L. (2009). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review*, 99(3):1053–69.
- Kilian, L. (2016). The Impact of the Shale Oil Revolution on U.S. Oil and Gasoline Prices. *Review of Environmental Economics and Policy*, 10(2):185–205.
- Kilian, L. (2017a). How the Tight Oil Boom Has Changed Oil and Gasoline Markets. *Papeles de Energia*, 3:79–102.
- Kilian, L. (2017b). The Impact of the Fracking Boom on Arab Oil Producers.
- Kilian, L. and Lee, T. K. (2014). Quantifying the speculative component in the real price of oil: The role of global oil inventories. *Journal of International Money and Finance*, 42(C):71–87.
- Kilian, L. and 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(5):1166–1188.
- Kilian, L. and Murphy, D. P. (2014). The Role Of Inventories And Speculative Trading In The Global Market For Crude Oil. *Journal of Applied Econometrics*, 29(3):454–478.
- Kilian, L. and Zhou, X. (2018). Structural Interpretation of Vector Autoregressions with Incomplete Information: Revisiting the Role of Oil Supply and Demand Shocks: Comment. CEPR Discussion Papers 13068, C.E.P.R. Discussion Papers.
- Mohaddes, K. and Pesaran, M. H. (2016). Oil Prices and the Global Economy; Is It Different This Time Around? IMF Working Papers 16/210, International Monetary Fund.
- Mohaddes, K. and Raissi, M. (2015). The U.S. Oil Supply Revolution and the Global Economy. IMF Working Papers 15/259, International Monetary Fund.

Nakov, A. and Nuño, G. (2013). Saudi arabia and the oil market. *The Economic Journal*, 123(573):1333–1362.

Sweeney, J. L. (1984). The response of energy demand to higher prices: What have we learned? *The American Economic Review*, 74(2):31–37.

Acknowledgements

The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the European Central Bank. We thank Matteo Ciccarelli, Lutz Kilian, Cedric Tille, Virginia di Nino, Fabrizio Venditti, Alex Pierru, Thomas Covert, and Felipe Saffie for useful comments and suggestions.

Claudia Foroni

European Central Bank, Frankfurt am Main, Germany; email: claudia.faroni@ecb.europa.eu

Livio Stracca

European Central Bank, Frankfurt am Main, Germany; email: livio.stracca@ecb.europa.eu

© European Central Bank, 2019

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-3571-5

ISSN 1725-2806

doi:10.2866/039247

QB-AR-19-090-EN-N