



Norwegian
Business School

This file was downloaded from BI Open Archive, the institutional repository (open access) at BI Norwegian Business School <https://biopen.bi.no>.

It contains the accepted and peer reviewed manuscript to the article cited below. It may contain minor differences from the journal's pdf version.

Gundersen, T.A. (2020). The Impact of U.S. Supply Shocks on the Global Oil Price. *Energy Journal*, 41(1). <https://doi.org/10.5547/01956574.41.1.tgun>

Copyright policy of IAEE, the publisher of this journal.

An early version of the paper (frequently the author's manuscript) is made available at no cost to users solely through non-commercial and institutional networks with which the author is affiliated, and is made available to the general public through the IAEE at a reasonable cost to those seeking access to the article, during a 3-year embargo period, after which the processed version will be available to all users at no charge, subject to a Creative Commons license



The Impact of U.S. Supply Shocks on the Global Oil Price*

Thomas S. Gundersen[†]

March 10, 2019

Abstract

I examine the role of the U.S. shale oil boom in driving global oil prices. Using a structural vector autoregressive (SVAR) model that identifies separate oil supply shocks for the U.S. and OPEC, I find that U.S. supply shocks can account for up to 13% of the oil price variation over the 2003–2015 period. This is considerably more than what has been found in other studies. Moreover, while U.S. oil production has increased substantially since 2010, U.S. oil supply shocks first started to contribute negatively to oil prices beginning in late 2013. This mismatch implies a temporary friction in the transmission of U.S. supply shocks to the rest of the world likely caused by logistical and technological challenges observed in the downstream supply chain.

JEL-codes: C32, Q33, Q35

Keywords: structural VARs, oil prices, demand and supply shocks, shale oil.

This paper has been accepted for publication in *the Energy Journal*.
doi.org/10.5547/01956574.41.1.tgun

*I would like to thank the Editor, four anonymous referees, Knut Are Aastveit, Hilde C. Bjørnland, Yoosoon Chang, Jamie Cross, Ana María Herrera, Nils B. L. Plosnić, Francesco Ravazzolo, Leif Anders Thorsrud and Ragnar Torvik, as well as participants at the 2017 IAAE conference in Sapporo, the 2017 MEG meeting at Texas A&M University in College Station and the 2017 CAMA workshop at the Australian National University in Canberra for helpful comments.

[†]Centre for Applied Macroeconomics and Commodity Prices (CAMP), Department of Economics, BI Norwegian Business School, Nydalsveien 37, N-0484 Oslo, Norway. Email: thomas.s.gundersen@bi.no.

1 Introduction

Few other commodities receive the same level of attention among economists as crude oil. It serves as an important input for a large share of production and is actively traded in the financial markets. It has been at the centre of wars and conflicts and can be a contributing factor to political turmoil and geopolitical tensions as well as leading to inflationary pressure and recessions (Hamilton 1983; 1985). Thus it is not surprising that a significant change in the price of oil spurs interest and debate.

During the summer of 2014, the oil price practically collapsed. From fluctuating around \$120 per barrel, the price hit the \$28 mark in January 2016, a decrease of more than 75 per cent. In the decade preceding the collapse, the United States saw an unprecedented surge in crude oil output after more than two decades of production decline. This sudden spurt was the result of innovations in shale oil extraction technology, as well as record high oil prices that made commercial shale production viable. Still, the role of U.S. shale oil in the subsequent collapse of the oil price is debated. In particular, since the seminal paper by Kilian (2009), demand shocks have been commonly viewed as the main driver of oil price fluctuations (see e.g. Kilian and Murphy (2012) and Aastveit, Bjørnland and Thorsrud (2015)). While Baumeister and Kilian (2016) argue that slow growth in emerging markets may also have played a key role in the large oil price drop of 2014, it is hard to rule out oil supply shocks considering the unprecedented surge in oil production over the last decade, especially in the United States. This is also supported by the recent findings of Caldara, Cavallo and Iacoviello (2016) who attribute a larger role for supply. Moreover, Behar and Ritz (2017) show that a switch from an accommodative strategy to a market-share strategy by OPEC can be rationalised when facing high-cost shale producers and predict that such switch happened in 2014. The paper most closely related to my study is Kilian (2017), who shows that the Brent crude price would have been up to \$10 higher without the U.S. shale oil boom by applying counterfactual analysis. Yet none of the recent papers examine the role of the United States, the shale oil boom and U.S. supply shocks explicitly.

In this paper, I aim to rectify this shortcoming and explore the implications of increased U.S. self-sufficiency in crude oil for oil prices over the 2003:M01–2015:M12 period. The hypothesis is that additional oil production produced by the U.S. shale oil fields has put a downward pressure on prices. To analyse this hypothesis, I estimate a structural vector autoregression (SVAR) model which

includes a measure of U.S. crude oil supply, OPEC production, a measure of global economic activity and the real price of oil. The model builds on Kilian (2009) in that oil supply and demand shocks are identified separately. Novel to my model, however, is the explicit distinction between U.S. and OPEC supply shocks. To quantify the role of changes in U.S. supply, I rely on the U.S. ban on crude oil exports and use a constructed U.S. imports variable. This restricts the sample to end in December 2015 with the lifting of the export ban in January 2016. The imports variable is constructed as the residual from a regression of U.S. crude oil imports on measures of domestic and foreign demand for U.S. crude oil. The choice of using U.S. imports differs from earlier studies that have used crude oil production to measure oil supply. I argue that using U.S. imports allows me to directly assess the price effect of the displacement of foreign crude oil. The identified supply shock from the U.S. is thus interpreted as a sudden negative shift in U.S. demand for foreign crude oil due to higher domestic availability. To motivate and frame the results, the analysis relies on evidence from a strand of literature that studies the insufficient capacity in the crude oil refining and transportation sectors following the U.S. shale oil boom (see among others Kilian 2016; McRae 2017; Wilkerson and Melek 2014).

I find strong support for the hypothesis that the U.S. shale industry has put a downward pressure on global oil prices. The analysis shows that the cumulative effect of U.S. supply shocks on the oil price has been negative leading up to the mid-2014 fall in prices. Further, following a positive U.S. supply shock that lowers U.S. imports by 1%, the real price of oil falls by almost 2%. U.S. supply shocks explain up to 13% of the variation in the real price of oil over the sample period. Taken together, U.S. and OPEC supply shocks account for a third of the variation in the real price of oil. This is considerably higher than what has been found by earlier studies in the literature and reintroduces supply as an important driver of oil prices.

The remainder of the paper is structured as follows. Section 2 gives a brief narrative of the U.S. shale oil boom and the plunge in the oil price during 2014 and 2015. Section 3 presents the SVAR model which includes an adjusted measure of U.S. crude oil imports to identify U.S. supply shocks. I present the results in section 4 while section 5 discusses robustness with alternative identifying restrictions and model specifications.

2 Data environment

2.1 The U.S. shale oil boom

Shale oil is petroleum found in rock formations of low permeability.¹ Primary recovery from conventional oil wells requires only drilling because the pressure differential brings the oil to the surface (Bret-Rouzaut and Favennec 2011). Shale oil, on the other hand, cannot be extracted by traditional methods as the sediment in which it is enclosed bars it from flowing freely. A combination of two technologies makes this extraction commercially viable: hydraulic fracturing (*fracking*) and horizontal drilling. The former allows the oil to escape the rocks and the latter lets more rock be fracked at the same time. The development of these technologies was fuelled by the period of high oil prices in the run-up to the financial crisis and subsequent years (Alquist and Guénette 2014; Kilian 2016; Maugeri 2013). Unconventional oil thus became competitive against conventional techniques, and investments in shale oil gained traction.

2.1.1 The Cushing glut

It is easy to justify an assertion that the fall in global oil prices, at least to some extent, was due to the U.S. shale oil boom. However, there are two caveats to consider relating to the timings of the boom and the fall in prices. The first is related to the chemical properties of shale oil and the second to the lack of appropriate transportation infrastructure in shale oil rich regions of the United States.

The chemical make-up of shale oil is different from that of conventional oil. In general, the chemical properties of shale oil are characterised as light and sweet, measured by API gravity and sulphur content respectively. These are similar to the properties of crude oil that has been produced in the U.S. traditionally. However, because of declining crude oil production from conventional wells since the 1970s, refineries along the Gulf Coast have been relying on imports from the Middle East to cover petroleum product demand. This oil, however, is characterised as heavy and sour. These refineries were fitted accordingly and have been mostly unable to process lighter oil without further investments in new or up-

¹Shale oil should not be confused with *oil shale*, sedimentary rocks with high kerogen content. Liquid petroleum can be extracted from these rocks, but it is a costly and capital intensive process (Bret-Rouzaut and Favennec 2011). Because of this confusion, shale oil is often referred to as *tight oil*.

graded equipment.² Several obstacles related to the processing of shale oil have also been reported by the industry, namely that shale oil differs chemically, not only from conventional oil, but also from sample to sample extracted from the same shale play (Baker Hughes 2013; Benoit and Zurlo 2014).

Second, North Dakota has become a prominent shale producer, but has not been an important producer of conventional oil historically. For this reason, pipelines and rail capacity did not exist to accommodate the rapid expansion of oil production in the state, making it more costly for the producers to get the oil to the market (McRae 2017; Wilkerson and Melek 2014). Traditionally, crude oil has been transported from the Gulf Coast to Cushing, Oklahoma where there has been more residual demand. The swelling supply of shale oil inland and no southbound pipelines from Cushing created an excess supply of oil in the U.S. interior (Kilian 2016).

These developments spawned a glut of light sweet crude oil in Cushing, Oklahoma. The emergence of this glut can be seen in Figure 1 as a price spread widening between the Brent Crude and the West Texas Intermediate (WTI) benchmarks.³ The persistent spread in prices created incentives for refineries to adapt to the new opportunities in shale oil refining. Over time and in addition to the use of trucks, rail and river barges, pipelines were constructed from the storage facilities in Cushing to the refineries along the Gulf Coast (see Kilian 2016, 2017; Wilkerson and Melek 2014). The closing of the Brent–WTI spread later on is an indication that the bottlenecks in the supply chain became less severe and thus U.S. oil to a larger extent was used as feedstock by the domestic refining industry.

The implication of this glut in Cushing was a temporary friction in the transmission of booming shale oil supply in the U.S. to the oil prices globally. While the WTI had to be sold at a considerable discount and with inventories of shale oil to staying high in Cushing, the refineries had to continue importing crude from abroad in order to satisfy demand for petroleum products. In other words, as the U.S. oil supply was booming, the impact on global prices was cushioned by the glut until downstream buyers were able to adapt their refining processes and

²While refining plants on the East Coast which typically import light and sweet North Sea oil could use the shale oil as feedstock, the necessary infrastructure to transport the oil is not in place. Some refineries on the Gulf Coast had the necessary equipment to process shale oil, but the lack of southbound pipelines from the Midwest hindered the adoption (Kilian 2016).

³It should be noted that decouplings between the WTI and the Brent is not a new phenomenon. There has previously been insufficient pipeline infrastructure to transport oil into Cushing causing the WTI to be sold at a premium (Fattouh 2007). As discussed by Büyüksahin et al. (2013), there has been a lack of storage capacity as well as pipeline infrastructure out of Cushing leading to discounts on the WTI as early as 2008.




Figure 1: West Texas Intermediate and Brent Crude benchmark prices (left axis) and U.S. crude oil imports and exports (right axis), 2000–2016, at monthly frequency. Source: Federal Reserve Bank of St. Louis Economic Data (FRED) and Energy Information Administration (EIA)

utilise domestic shale oil to a greater extent, thereby reducing the need for foreign imports. This mechanism will be modelled explicitly in the next section.

3 A Structural VAR with U.S. imports of crude oil

In this section, I present the empirical model that I use in my analysis. Novel to my approach is that in order to capture U.S. oil supply, I construct a measure of U.S. imports of crude oil. I then analyse the effects of U.S. oil supply shocks, along with the other shocks, on the global oil market using a structural VAR model.

3.1 U.S. imports of crude oil

Since I am interested in how developments in the U.S. oil industry have affected the global oil market, a U.S. centric model design with an appropriate measure of U.S. oil supply is needed. Ideally, one would want to capture shifts in U.S. demand for foreign crude oil that occur due to a higher availability of domestic supply. As such, U.S. crude oil production is an inappropriate measure because changes in flow supply does not translate directly into shifts in the demand schedule for imports that affect non-U.S. oil prices. After the enactment of the Energy Policy and Conservation Act of 1975, the U.S. government banned exports of crude oil and natural gas. While, given the appropriate permissions, some export could take

place, the extent of this flow was negligible relative to total U.S. oil production.⁴ A key point is then that any increase in U.S. supply of crude oil only transmits to prices globally if it displaces foreign sources of oil. However, as the export ban was effectively lifted in January 2016 with the passing of the Consolidated Appropriations Act, the sample used in the analysis must end in December 2015 for the argument to remain valid.

Equation 1 explains the relationship between U.S. crude oil self-sufficiency, net exports and changes in inventories.

$$\text{U.S. Production} - \text{U.S. Consumption} = \text{Exports} - \text{Imports} + \Delta\text{Inventories} \quad (1)$$

In this context, consumption refers to the number of barrels of crude oil U.S. refineries use as input to produce petroleum products. Exports can be set to zero due to the export ban. Moreover, holding inventories are likely only to be done to smooth out gluts and shortfalls in the flow of oil. Therefore, build-ups of inventories are likely to be transitory and even out over time.⁵ The widening of the Brent–WTI spread that began in early 2011 coincided with a build-up of stocks at tank farms and pipelines which lasted till May the same year. The International Energy Agency (IEA) and the U.S. EIA announced in June that 30 million barrels of crude oil would be released from the U.S. Strategic Petroleum Reserve to offset negative effects of the political unrest in Libya (EIA 2011). This prompted the futures prices to fall and a lowering of inventories that did not end until futures prices had rebounded in November that same year. From December to May 2012, total inventories increased by 55.5 million barrels (roughly 7 days worth of total U.S. crude oil output) and reached its pre-2014 peak. The monthly growth rate however, never exceeded $\pm 2.5\%$ during this period. By setting exports and the change in inventories equal to zero and re-arranging, equation 1 can be simplified to:

$$\text{Imports} = \text{U.S. Consumption} - \text{U.S. Production} \quad (2)$$

The demand for oil that is not satisfied by domestic production thus has to be

⁴These exports have mainly come from production sites in California and Alaska (EIA 2014; 2015). On average, exports relative to U.S. production have been around 1% from 2003 to 2015 with a peak at 6% in April 2015. The vast majority of U.S.-sourced crude goes to Canada (EIA 2014b), but starting in 2014, some crude leaving the Gulf was shipped to Europe and Asia. However, this is Canadian oil re-exported by the United States (EIA 2014a). This might help explain the growth in U.S. exports relative to U.S. production starting in 2014.

⁵The monthly changes in U.S. inventories resemble a covariance stationary process with mean close to zero, see Appendix B for details.




Figure 2: U.S. crude oil production and imports, 1970–2015. Source: EIA

covered by changes in imports.

Figure 2 shows a clear negative correlation between U.S. oil production and imports over the long run. Tropical storms and hurricanes are temporary shocks to both variables causing the series to co-move. Most of these shocks hit during the August–October hurricane season. Hurricane Katrina (2005) and Hurricane Ike (2008) are by far the most devastating in terms of volumes and clearly visible in Figure 2. Further, the decline in imports after 2005 precedes the boom in domestic supply, indicating that the initial fall in imports was due to lower U.S. consumption. Crude oil demand can be controlled for by including a real economic activity variable in the VAR model. However, it will enter the U.S. equation with a lag and not account for the contemporaneous correlation. This poses a challenge for representing U.S. oil supply with U.S. imports as variations in imports would need to be uncorrelated with shifts in domestic demand and the business cycle while at the same time being correlated with observed domestic supply. To ensure that this condition is satisfied, I regress U.S. imports on variables that reflect demand for oil to clean out this variation. I then use the residual from this regression as a proxy for U.S. oil supply. This proxy will be uncorrelated with the demand factors by construction.

Two variables with the appropriate monthly frequency and sample availability

were chosen for the regression.⁶ *Vehicle miles travelled* captures domestic demand for petroleum products through the use of road vehicles. It includes cars and larger diesel vehicles used in freight transportation.⁷ More traffic on U.S. roads implies a higher demand for petroleum products which induces demand for crude oil from refineries. The refineries can choose to import the oil or use what is produced domestically. However, the mileage on the U.S. vehicles each month does not affect the amount of oil extracted from the ground directly. *Petroleum product exports* captures demand for American crude oil abroad through exports of refined products from U.S. refineries.⁸ Again, the refineries have to make use of imports or domestic supply. How much is refined and exported does not directly determine how much crude oil is taken up from the ground. See Appendix A for more information about the dataset.

$$\Delta usimp_t = \underset{(0.28^{***})}{0.843} \Delta vmt_t + \underset{(0.03^{***})}{0.12} \Delta petrolexp_t + \hat{e}_t \quad \bar{R}^2 = 0.17 \quad (3)$$

Equation 3 describes the regression that I estimate on the sample 2003:M01–2015:M12. All variables are in log-differences and the standard errors are shown in parentheses.⁹ The residual is the measure of U.S. oil supply that I will later include in the structural VAR model (explained in the next section). The variable (residual) is orthogonal to contemporaneous demand innovations and therefore captures the supply effects on oil imports. In other words, a negative shock to the modified U.S. imports variable can be interpreted as a decision by refineries to import less crude because of a sudden abundance of domestic supply. A positive shock will then reflect the need for more imports because of less domestic production.

⁶Sales of heating oil or other 'distillate fuel oils' used in industry could have been other fitting choices. However, the set of candidate variables to include in the regression is constrained by the availability of data on a monthly frequency.

⁷The variable is compiled by the U.S. Federal Highway Administration and based on vehicle counts. Each data-point is calculated by taking the total number of vehicles travelling on a given stretch of public road that month and multiplying it by the road's length in miles and the number of lanes. The final step is to sum over all monitored roads.

⁸The U.S. became a net exporter of petroleum products in 2011 (EIA 2015)

⁹To check for robustness, lagged values of the regressors have been included as well as having the estimation of the coefficients in Equation 3 being based on pre-shale samples. These alternative specifications do not change the baseline results reported below. Finally, tests for parameter stability (Bai-Perron) have been carried out and specified with a variety of trimming percentages. The tests do not suggest any breaks in the post-shale period. Of the suggested breakpoints, none are statistically significant except one at 2009:M04 which is so at the 10% level.

3.2 The Structural VAR model

I will now include my constructed measure of U.S. imports in a Structural VAR model. Having a U.S.-specific variable in the model necessitates the use of OPEC production as an alternative measure of foreign supply. The reason for this is twofold. The first is due to a possible simultaneity issue, as U.S. oil production is a component of global production. The second reason has to do with the data itself. Aggregate global production exhibits low variation relative to more disaggregated measures, possibly reflecting that a shortfall of production in one location is met by an increase somewhere else thereby neutralising fluctuations. While the same can be said about the producers within OPEC, the member countries as a group account for most of the short-run fluctuations in the global output (see e.g. Almoguera, Douglas and Herrera 2011). OPEC production is an interesting candidate as it represents a large bulk of global production and possibly captures some interesting dynamics between itself and the U.S. Behar and Ritz (2017) predicts that OPEC will, facing increased production by shale oil producers, change behaviour and increase production as well. Adding separate supply-equations for different oil producers has been done by Kang, Ratti and Vespignani (2016; 2017) with U.S. and non-U.S. production, Ratti and Vespignani (2015) and Kolodziej and Kaufmann (2014) with OPEC and non-OPEC production. Common to these papers is the argument that aggregate global production leads to underestimation of the influence of supply shocks on oil prices.

This model is an augmentation of the Kilian (2009) 3-variable model that includes aggregate global oil production, the measure of global activity and the real price of oil.¹⁰

Consider the following reduced form VAR model

$$\mathbf{Y}_t = \boldsymbol{\mu} + \sum_{p=1}^P A_p \mathbf{Y}_{t-p} + \mathbf{e}_t \quad (4)$$

where $\mathbf{Y}_t = [\Delta opecprod_t, \Delta usimp_t^S, rea_t, lrpo_t]'$ is the vector of variables, percentage change in OPEC crude oil production, the adjusted U.S. crude oil imports constructed in section 3.1, an index of real economic activity (Kilian 2009) and the

¹⁰In Appendix D.2, I show that the results reported below are broadly robust to a specification that includes crude oil inventories as in Kilian and Murphy (2014). However, this brings the number of equations up to five and is challenging to estimate with the limited sample used here.

real price of oil respectively. rea_t and $lrpo_t$ are in logs.¹¹ $\boldsymbol{\mu}$ is a vector of intercept terms and $\mathbf{e}_t \stackrel{iid}{\sim} N(0, \Sigma_e)$ where Σ_e is positive semi-definite and symmetric. All data is on monthly frequency and the model is specified with 18 lags. Hamilton and Herrera (2004) demonstrated with their replication of the Bernanke, Gertler and Watson (1997) model that a rich lag structure is needed to capture oil price shocks. An *ex ante* choice of 1.5 years worth of lags rather than the use of information criteria is in line with the recommendations of Kilian and Lütkepohl (2017).¹²

To identify the structural shocks, let the reduced form errors be decomposed in the following way $\mathbf{e}_t = S\boldsymbol{\varepsilon}_t$, or

$$\begin{bmatrix} e^{\Delta opecprod} \\ e^{\Delta usimp^S} \\ e^{rea} \\ e^{lrpo} \end{bmatrix}_t = \begin{bmatrix} s_{11} & 0 & 0 & 0 \\ s_{21} & s_{22} & 0 & 0 \\ s_{31} & s_{32} & s_{33} & 0 \\ s_{41} & s_{42} & s_{43} & s_{44} \end{bmatrix} \begin{bmatrix} \varepsilon^{\Delta opecprod} \\ \varepsilon^{\Delta usimp^S} \\ \varepsilon^{rea} \\ \varepsilon^{lrpo} \end{bmatrix}_t \quad (5)$$

where matrix S is the lower triangular component of the Cholesky decomposition of Σ_e and $\boldsymbol{\varepsilon}_t$ is the structural shocks with the property that $\mathbb{E}[\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'] = I$. The way S is identified implies that a recursive structure is imposed where the responses of the variables ordered at the top in \mathbf{Y}_t will be restricted to zero contemporaneously.

Supply variables are ordered at the top, followed by global demand and, lastly, the oil price. OPEC supply shocks are defined as unexpected changes in oil production in OPEC member countries. A U.S. import shock is a sudden change in the importing decision of U.S. refineries reflecting the availability of domestically produced crude oil. By placing OPEC on top, a short-run vertical supply curve is imposed. Hence, OPEC cannot adjust their production within a month after shocks to aggregate demand, nor after shifts in beliefs about the state of the future oil markets (oil-specific demand shocks). Taking into consideration the adjustment costs of changing their production schedules, necessary cartel coordination among OPEC members, but also lack of information regarding business cycle movements in real time, oil producers are likely to respond to these innovations with a lag. Additionally, OPEC cannot respond to U.S. import shocks contemporaneously, reflecting OPEC's inability to observe what the United States imports from abroad

¹¹ $\Delta opecprod_t$ and $\Delta usimp_t^S$ are $I(0)$ while rea_t and $lrpo_t$ are $I(1)$. See Appendix A for more information about the dataset. Importantly, the VAR model in Equation 4 is stable and therefore has a moving average representation.

¹²Beginning with Kilian (2009), it has been customary to include 24 lags in these models. However, due to the number of parameters to be estimated relative to the sample size, only 18 lags are included in this model.

in real time. This information is published by the EIA later on.¹³ Finally, the refineries in the U.S are assumed not to react to aggregate demand and oil-specific demand shocks instantly. Although oil prices are observed in the market daily, the American suppliers are slow to ramp up their production due to adjustment costs. Hence, the effect of higher oil prices on imports when the U.S. supply situation is taken into account is delayed. Still, the United States is ordered beneath OPEC production as there is evidence which suggests that shale producers are more flexible than conventional producers (see e.g. Bjørnland, Nordvik and Rohrer 2017).¹⁴

An abrupt change in global real activity is here represented by a shock to the demand of industrial commodities, henceforth called an aggregate demand shock (see Kilian 2009). Innovations to the real price of oil that are not explained by either supply or demand are called oil-specific demand shocks and reflect primarily precautionary demand for crude oil related to expectations of future supply shortfalls (see Kilian 2009).¹⁵ The exclusion restriction implies that global real activity takes one month to adjust to oil-specific demand shocks. While oil prices are observable daily, economic agents are slow to change their behaviour, and the effect on the level of real activity is therefore delayed. This is consistent with the historical relationship between oil prices and business cycle movements (see e.g. Hamilton (1985)). The real price of oil equation is left unrestricted. These identifying restrictions are similar to those first imposed by Kilian (2009). In section 5, I show that the results are robust to alternative restrictions.

4 Empirical results

The sample period used for the estimation is 2003:M01–2015:M12. As shown by Baumeister and Peersman (2013), parameter instability is a prevalent feature of oil market models over the commonly estimated sample beginning in the early 1970s. The choice of sample period is ultimately motivated by the research question. Shale oil production in the U.S. had just begun in 2003 and the U.S. export ban on exports was lifted in December 2015. Extending the sample backwards, however, does not affect the main results until 1997 when statistical significance is lost. Estimated impulse responses are shown in Figure 3. Since the U.S. imports

¹³Since the structural shock to U.S. imports reflects domestic supply conditions, it is even tougher for OPEC to monitor.

¹⁴In Appendix D, I show that the results are robust to an alternative ordering of the variables.

¹⁵The interpretation of this shock should not be taken too literally however, as it will reflect all residual variation in the oil price not explained by the other endogenous variables.

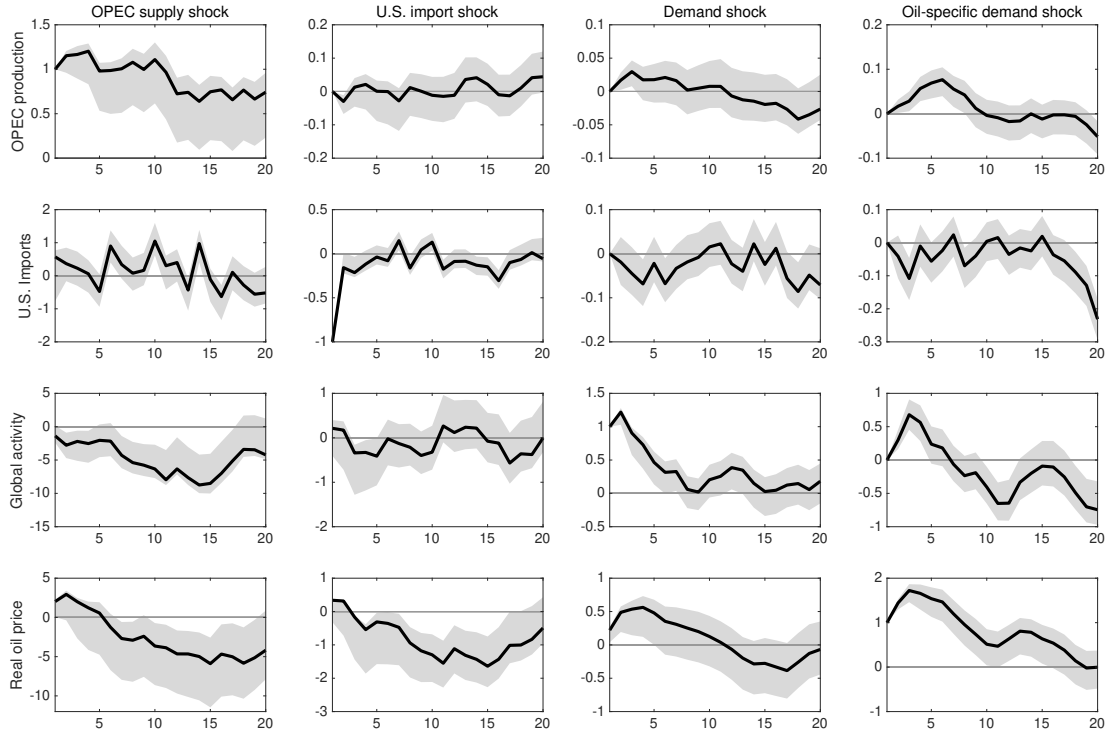


Figure 3: Impulse responses generated from the model described in equations 4 and 5. The sample is 2003:M01–2015:M12. They are all in levels of the variables. Shocks are normalised so that the response of the variables is 1 on impact, i.e. 1% for the OPEC supply shock, U.S. supply shock and aggregate demand shock while one log unit for the oil price. The shaded areas represent 68% confidence bands calculated using a bootstrap with 10,000 draws.

variable is a generated regressor, the reported confidence bands are computed using on a two-stage residual bootstrap procedure.¹⁶

Starting with the supply shocks (left columns), a sudden innovation to OPEC supply growth leads to a persistent increase in their level of production. The United States begins to import more on impact and periodically so over the next year. It is however, hard to interpret the nature of this response. The response of global activity is clearly negative and statistically significant over time. The real price of oil initially increases but is not significant and turns negative within 6 months. The persistently negative response of global activity is puzzling given

¹⁶Each iteration of the procedure follows these steps: first a residual bootstrap sample is generated from equation 3. Second, a residual bootstrap of the VAR model is carried out where the U.S. imports variable is replaced by the bootstrap sample from the first step. Finally, the bootstrap sample from the previous step is used to estimate a VAR model and impulse response functions are computed and stored.

that the OPEC supply shock causes oil prices to fall.¹⁷

The second column shows the responses to the U.S. import shock. A negative shock to U.S. imports, reflecting a sudden abundance of domestically produced crude oil, does not change OPEC production nor global activity. Interestingly, U.S. imports exhibit a very low degree of persistence as it returns to pre-shocks levels within two months. Still, the oil price exhibits a persistent but gradual decrease. When the U.S. reduces its imports by 1%, the oil price falls by almost 2% after ten months and is significantly negative after eight months. As a comparison, for the OPEC supply shock to decrease the oil price by the same amount, the OPEC supply shock must increase OPEC production by approximately 0.3%.

Turning to the demand side, the third column shows the responses to the aggregate demand shock that increases global activity. The shock leads OPEC and the United States to produce more (imports less) crude oil, but only temporarily. The oil price responds by increasing on impact and follows a hump-shaped trajectory back to zero as expected.

Following an oil-specific demand shock (right column), OPEC starts to produce more while the Americans import less, implying that their domestic supply is higher. For OPEC, the response is slow, but it is much more persistent than that of the U.S. and lasts for almost eight periods. A puzzling result is that of global activity, which initially increases following the shock to oil prices. This is a similar result as that seen in Kilian (2009), later attributed to not allowing emerging and developed economies to respond differently to oil market shocks (see Aastveit et al. 2015).

The historical decomposition of the real price of oil is presented in Figure 4. The cumulative effect of the U.S. import shock has since late-2013 contributed to pushing oil prices down. OPEC, on the other hand, seems to have been working to increase prices following the 2014 fall. Oil-specific demand has also contributed, possibly reflecting the expectations of an oversupply in the oil markets. Caldara et al. (2016) find similar results for the 2014–15 episode, but do not identify separate U.S. and OPEC supply shocks. Baumeister and Kilian (2016) find evidence that prior to the slump in oil prices, movements in oil supply could predict parts of the decline that would occur later, which is consistent with the findings here. In contrast to the U.S. import shocks, aggregate demand shocks did not influence

¹⁷When the OECD+6 industrial production index constructed by Baumeister and Hamilton (2017) and detrended according to the recommendations of Hamilton (2018) is used in place of the Kilian index, global activity responds positively (but not statistically significant) as one would expect.




Figure 4: Historical decomposition of the real price of oil derived from the model described in equations 4 and 5. The sample is 2003:M01–2015:M12 The shaded areas correspond to the cumulative effects of the different shocks on the oil price.

prices negatively until 2014 and only did so for the first of the two dips in prices that occurred between 2014 and 2016. Overall, the influence of demand factors in the historical decomposition is reduced compared to e.g. Kilian (2009).

The variance decomposition of the real price of oil is presented in Table 1. It shows the relative contributions of each shock to variations in the oil price. U.S. supply-side innovations explain up to 13% of the variation in the oil price. OPEC and the United States together account for 30% of the fluctuations in the oil price at the 18-months forecast horizon according to the model. In recent years, the literature has been giving supply-side explanations of oil price fluctuations an increasingly smaller role. The current results, together with Caldara et al. (2016), however, provide evidence of the importance of supply. To put this into dollar amounts, I do a back-of-the-envelope calculation of the implied path of the real price of oil in the absence of U.S. import shocks after September 2013. Specifically, I recompute the solid line in Figure 4 by setting the cumulative effect of U.S. import shocks to zero after September 2013, add back the estimated deterministic term and then reflate it to nominal terms. As can be seen in Figure 5, in the absence of U.S. import shocks, the oil price would have been roughly \$10 higher if it were not for the U.S. shale oil boom. This is a similar estimate to that of Kilian (2017).

The results presented here suggest that the U.S. shale oil boom has contributed significantly to lowering oil prices during 2014 and 2015. While this result might seem surprising considering the rapid growth in U.S. shale oil output that began as early as 2011, it is consistent with the accounts of frictions in the supply-

Variance decomposition of the real price of oil

Shocks	Horizons			
	1 month	6 months	12 months	18 months
OPEC supply	7.73 (0.26; 8.48)	2.83 (1.54; 11.95)	8.77 (4.27; 29.46)	17.58 (6.90; 40.71)
U.S. imports	2.12 (0.13; 5.84)	1.05 (1.99; 19.92)	8.10 (6.32; 32.00)	13.04 (8.00; 34.02)
Aggregate demand	7.15 (0.46; 10.14)	13.63 (2.30; 19.59)	10.94 (3.35; 19.43)	10.73 (4.85; 23.19)
Oil-specific demand	83.00 (79.86; 95.29)	82.48 (58.16; 85.84)	72.19 (36.61; 70.12)	58.65 (26.20; 58.44)

Table 1: Variance decomposition (in percentages) of the real price of oil for different time horizons, generated from the imports model described in equations 4 and 5. The sample is 2003:M01–2015:M12. The confidence intervals (in parenthesis) are at the 68% level and computed using a bootstrapping method with 10,000 draws.

chain and lack of pipeline infrastructure delaying the adoption of shale oil by U.S. refineries (see section 2.1.1). The closing of the Brent–WTI spread in late 2013 and early 2014 lines up with the emergence of the negative cumulative effect U.S. import shocks had on the oil price as can be seen by comparing Figure 1 and 4. This suggests that the adoption of shale oil in the domestic refining sector finally displaced foreign crude oil imports and the price spread narrowed as a result.¹⁸

The choice of representing the U.S. supply side of the oil market with (adjusted) U.S. imports was motivated by the research question. However, it has been customary in this literature to have a supply equation where the endogenous variable is the quantity of crude oil produced. The conclusion commonly drawn from these models is that oil supply shocks cannot explain oil-price fluctuations. The results from a model identical to the one presented above, but where U.S. imports are replaced by U.S. crude oil production, are presented in Appendix C. The gains from identifying the U.S. supply shocks by exploiting the relationship between domestic production and imports rather than with the quantity produced directly are evident. The response of OPEC to a positive U.S. supply shock, as seen in Figure 7, is sensible in that they increase output, consistent with attempting to keep their market share. The model does predict that the oil price temporarily

¹⁸As has been pointed out earlier, U.S. imports and U.S. crude oil production have been sensitive to hurricanes and tropical storms in the Gulf Region. However, the 2013 Atlantic Hurricane season was the least active in two decades (NOAA 2014) so these disruptions are not driving the results in this period.




Figure 5: The implied path of the nominal price of oil in the absence of U.S. import shocks.

increases which might also drive OPEC upwards. This, together with the response of global activity – which both get a temporary boost – are however, puzzling. In particular, one would not expect oil prices to increase when both the U.S. and OPEC expand output. The estimated response eventually turns negative, but is not significantly different from zero.¹⁹ In addition, the model fails at explaining the 2014 fall in oil prices as illustrated by the historical decomposition in Figure 8. Specifically, the oil-specific demand shock (a residual shock) explains the lion’s share of the movements in the real price of oil past mid-2014 suggesting that the supply (and demand) shocks are not well identified. Hence, a model that includes U.S. crude oil production rather than U.S. imports does not shed light upon the research question or reaffirm results from previous studies.

5 Robustness and sensitivity checks

To check the robustness of the results I impose different identifying restrictions as well estimate a model where inventories are included. The details are given in Appendix D, so a brief summary will suffice here.

First, I rearrange the ordering of the variables so that the U.S. is placed on top and OPEC second. The results do not change compared to those of the main baseline ordering. Second, I re-estimate the baseline model using Bayesian methods with flat priors. For identification, I impose a mix of sign and zero restrictions on the contemporaneous impact matrix. In particular, the zero-restriction on the s_{12}

¹⁹Kilian (2009) also found that the response of global activity moves in the same direction as oil production following a supply shock. However, his results also showed that the real price of oil moves in the opposite direction, contrary to the results here.

parameter is relaxed to be negative on impact. This implies that following a shock that lowers U.S. imports, OPEC will respond by increasing their output. The main results from the baseline model remain robust to the chosen identification strategy. Finally, I estimate an augmented version of the Kilian and Murphy (2014) model in order to account for changes in inventories. The structural shocks are identified with sign-restrictions. Again, the main conclusions drawn from the baseline model remains unchanged. However, the response of the oil price to a U.S. import shock is less persistent than in the baseline model and returns to pre-shock levels after 6 months compared to 16 in the baseline model.

6 Conclusion

In this paper I analyse the impact of the U.S. shale oil boom on global oil prices. In doing so, I estimated a structural VAR based on Kilian (2009) with OPEC production, a modified U.S. crude oil import variable, a measure of real economic activity and the real price of oil. The use of crude oil import data in a structural VAR to model the case of the United States directly is to my knowledge new to the literature. It is modified so as only to capture U.S. supply innovations. This approach is the most sensible given the institutional framework in place in the United States up until December 2015 as it gives a clear transmission mechanism of domestic supply shocks to oil prices abroad.

Firstly, the findings show that a 1% reduction in U.S. imports causes the oil price to decrease by almost 2% after ten months. The U.S. import shock, reflecting the domestic supply environment, explains up to 13% of the variation in the oil price over the sample period 2003–2015. The U.S. and OPEC together account for a third of the variation in the oil price. This is significantly more than what has been found in earlier studies. Secondly, the results show that the developments in the U.S. oil industry had no significant effect on global prices until the end of 2013.

These results suggest that the U.S. shale oil boom *has* in fact been able to affect global oil prices negatively. However, the analysis shows that oil prices were not affected until the end of 2013. The cause of the delay is puzzling considering the length of time U.S. production figures had been on the rise. A possible explanation for the lagged transmission of U.S. supply shocks to the rest of the world is the oil glut in Cushing, Oklahoma caused by insufficient pipeline capacity and incompatible refining equipment that postponed adoption of shale oil by the domestic refining industry, indirectly observable in the WTI–Brent price spread.

The results put forward in this paper add to the discussion of the role of the U.S. in the oil price fall of 2014/2015. Contrary to earlier studies, I find an increased importance of supply side factors in explaining oil price fluctuations. Further, due to the emergence of the U.S. shale oil industry, the role of the United States in the market has fundamentally changed and will have implications for oil prices globally going forward depending on the state of transportation and refining infrastructure.

References

- Aastveit, Knut Are, Hilde C. Bjørnland and Leif Anders Thorsrud (2015). ‘What drives oil prices? Emerging versus developed economies’. In: *Journal of Applied Econometrics* 30.7, pp. 1013–1028.
- Almoguera, Pedro A., Christopher C. Douglas and Ana María Herrera (2011). ‘Testing for the cartel in OPEC: non-cooperative collusion or just non-cooperative?’ In: *Oxford Review of Economic Policy* 27.1, pp. 144–168.
- Alquist, Ron and Justin-Damien Guénette (2014). ‘A blessing in disguise: The implications of high global oil prices for the North American market’. In: *Energy Policy* 64, pp. 49–57.
- Alquist, Ron, Lutz Kilian and Robert J Vigfusson (2013). ‘Forecasting the Price of Oil’. In: *Handbook of Economic Forecasting* 2, pp. 427–507.
- Baker Hughes (2013). *Overcoming Shale Oil Processing Challenges*. White paper. Houston, TX.
- Baumeister, Christiane J.S. and James D Hamilton (2017). ‘Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks’. In: Working Paper Series 24167.
- Baumeister, Christiane and Lutz Kilian (2016). ‘Understanding the Decline in the Price of Oil since June 2014’. In: *CESifo Working Paper Series* 5755.
- Baumeister, Christiane and Gert Peersman (2013). ‘The Role of Time-varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market’. In: *Journal of Applied Econometrics* 28.7, pp. 1087–1109.
- Behar, Alberto and Robert A Ritz (2017). ‘OPEC vs US shale: Analyzing the shift to a market-share strategy’. In: *Energy Economics* 63, pp. 185–198.
- Benoit, Brian and Jeffrey Zurlo (2014). ‘Overcoming the challenges of tight/shale oil refining’. In: *Processing Shale Feedstocks*, pp. 37–44.
- Bernanke, Ben S, Mark Gertler and Mark Watson (1997). ‘Systematic Monetary Policy and the Effects of Oil Price Shocks’. In: *Brookings Papers on Economic Activity* 1997.1, pp. 91–142.
- Bjørnland, Hilde C., Frode Martin Nordvik and Maximilian Rohrer (2017). *Supply Flexibility in the Shale Patch: Evidence from North Dakota*. CAMP Working Paper series 2/2017. Centre for Applied Macro- and Petroleum Economics, BI Norwegian Business School.
- Bret-Rouzaut, Nadine and Jean-Pierre Favennec (2011). *Oil and Gas Exploration and Production: Reserves, Costs, Contracts*. Paris: Editions TECHNIP. ISBN: 9782710809753.
- Büyüksahin, Bahattin et al. (2013). ‘Physical markets, paper markets and the WTI-Brent spread’. In: *The Energy Journal*, pp. 129–151.
- Caldara, Dario, Michele Cavallo and Matteo Iacoviello (2016). *Oil Price Elasticities and Oil Price Fluctuations*. Mimeo, Federal Reserve Board.
- Canova, Fabio (2007). *Methods for Applied Macroeconomic Research*. Princeton University Press.
- Canova, Fabio and Gianni De Nicolò (2003). ‘On the Sources of Business Cycles in the G-7’. In: *Journal of international economics* 59.1, pp. 77–100.
- Canova, Fabio and Matthias Paustian (2011). ‘Business cycle measurement with some theory’. In: *Journal of Monetary Economics* 58.4, pp. 345–361.

- EIA – U.S. Energy Information Administration (June 2011). *International Energy Agency members release strategic petroleum stocks*. URL: <https://www.eia.gov/todayinenergy/detail.php?id=1950>.
- (Oct. 2014a). *Crude exports and re-exports continue to rise; some volumes sent to Europe and Asia*. URL: <http://www.eia.gov/todayinenergy/detail.cfm?id=18631> (visited on 16/06/2016).
- (June 2014b). *U.S. crude exports in April rise to highest level in 15 years*. URL: <http://www.eia.gov/todayinenergy/detail.cfm?id=16711> (visited on 15/06/2016).
- (Sept. 2015). *Effects of Removing Restrictions on U.S. Crude Oil Exports*. Washington D.C. URL: <http://www.eia.gov/analysis/requests/crude-exports/pdf/fullreport.pdf> (visited on 29/06/2016).
- Fattouh, Bassam (Apr. 2007). ‘WTI Benchmark Temporarily Breaks Down: Is It Really a Big Deal?’ In: *Oxford Energy Comment*.
- Faust, Jon (1998). ‘The robustness of identified VAR conclusions about money’. In: *Carnegie-Rochester Conference Series on Public Policy*. Vol. 49. Elsevier, pp. 207–244.
- Hamilton, James D (2018). ‘Why you should never use the Hodrick-Prescott filter’. In: *Review of Economics and Statistics* 100.5, pp. 831–843.
- (1983). ‘Oil and the macroeconomy since World War II’. In: *The Journal of Political Economy* 91.2, pp. 228–248.
- (1985). ‘Historical causes of postwar oil shocks and recessions’. In: *The Energy Journal* 6.1, pp. 97–116.
- Hamilton, James D. and Ana María Herrera (Apr. 2004). ‘Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy. A comment by James D. Hamilton and Ana María Herrera’. In: *Journal of Money, Credit and Banking* 36.2, pp. 265–286.
- Kadiyala, K Rao and Sune Karlsson (1997). ‘Numerical methods for estimation and inference in Bayesian VAR-models’. In: *Journal of Applied Econometrics*, pp. 99–132.
- Kang, Wensheng, Ronald A Ratti and Joaquin L Vespignani (2016). ‘The impact of oil price shocks on the US stock market: A note on the roles of US and non-US oil production’. In: *Economics Letters* 145, pp. 176–181.
- (2017). ‘Oil price shocks and policy uncertainty: New evidence on the effects of US and non-US oil production’. In: *Energy Economics* 66, pp. 536–546.
- Kilian, Lutz (2009). ‘Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market’. In: *American Economic Review* 99.3, pp. 1053–69.
- (2016). ‘The impact of the shale oil revolution on US oil and gasoline prices’. In: *Review of Environmental Economics and Policy* 10.2, pp. 185–205.
- (2017). ‘The Impact of the Fracking Boom on Arab Oil Producers.’ In: *Energy Journal* 38.6.
- Kilian, Lutz and Helmut Lütkepohl (2017). *Structural Vector Autoregressive Analysis*. Cambridge University Press. ISBN: 9781108164818.
- Kilian, Lutz and Daniel P Murphy (2014). ‘The Role of Inventories and Speculative Trading in the Global Market for Crude Oil’. In: *Journal of Applied Econometrics* 29.3, pp. 454–478.

- Kilian, Lutz and Daniel P. Murphy (2012). ‘Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models’. In: *Journal of the European Economic Association* 10.5, pp. 1166–1188.
- Kilian, Lutz and Robert J Vigfusson (2011). ‘Nonlinearities in the Oil Price–Output Relationship’. In: *Macroeconomic Dynamics* 15.S3, pp. 337–363.
- Kolodziej, Marek and Robert K Kaufmann (2014). ‘Oil demand shocks reconsidered: A cointegrated vector autoregression’. In: *Energy Economics* 41, pp. 33–40.
- Maugeri, Leonardo (2013). ‘The shale oil boom: a US phenomenon’. In: *Belfer Center Discussion Paper* 2013-05.
- McRae, Shaun (2017). *Crude Oil Price Differentials and Pipeline Infrastructure*. Working Paper 24170. National Bureau of Economic Research.
- NOAA (May 2014). *National Hurricane Center Annual Summary: 2013 Atlantic Hurricane Season*. Report. National Oceanic and Atmospheric Administration, U.S. Department of Commerce.
- Ratti, Ronald A and Joaquin L Vespignani (2015). ‘OPEC and non-OPEC oil production and the global economy’. In: *Energy Economics* 50, pp. 364–378.
- Rubio-Ramirez, Juan F, Daniel F Waggoner and Tao Zha (2010). ‘Structural vector autoregressions: Theory of identification and algorithms for inference’. In: *The Review of Economic Studies* 77.2, pp. 665–696.
- Uhlig, Harald (2005). ‘What are the effects of monetary policy on output? Results from an agnostic identification procedure’. In: *Journal of Monetary Economics* 52.2, pp. 381–419.
- United States Congress (1975). *Energy Policy and Conservation Act*. 42 U.S.C. 6212. Chap. 103, (b)(1), p. 877.
- (2016). *Consolidated Appropriations Act*. H. R. 2029. Chap. 101, (a), p. 746.
- Wilkerson, Chad R and Nida Cakir Melek (2014). ‘Getting crude to market: central US oil transportation challenges’. In: *Main Street Economist* 1, pp. 1–7.

A Data

Variable	Description	Source
$\Delta opecprod$	Log-change in total OPEC crude oil production in thousands of barrels per day. Datastream identifier: OPPCOBD.P	Thomson Reuters Datastream – EIA
$\Delta usimp^S$	Log-change in U.S. crude oil imports in millions of barrels per day. Adjusted for demand factors (see text for details).	U.S. Energy Information Administration (EIA)
Δvmt	Log-change in the Vehicle Miles Travelled index compiled from automatic roadside traffic monitors. Seasonally adjusted.	Federal Highway Administration. Retrieved from FRED database, St. Louis Fed. https://fred.stlouisfed.org/series/TRFVOLUSM227SFWA
$\Delta petrolexp$	Log-change in U.S. petroleum products exports in million barrels per day.	U.S. Energy Information Administration (EIA).
$\Delta usprod$	U.S. field production of crude oil in thousands of barrels per day.	U.S. Energy Information Administration (EIA)
rea	Measure of global real economic activity based on dry cargo bulk freight rates. Monthly deviations from trend. Introduced in Kilian (2009).	http://www-personal.umich.edu/~lkilian/paperlinks.html
$lrpo$	Log of refiner’s acquisition cost of crude oil imports deflated by the U.S. CPI.	U.S. Energy Information Administration (EIA). U.S. CPI retrieved from the FRED database, St. Louis Fed.
$\Delta gprod$	Global crude oil production in thousands of barrels per day. Datastream identifier: WDPCOBD.P	Thomson Reuters Datastream – EIA
$\Delta usinv$	Percentage change in U.S. crude oil inventories. Including the United States Strategic Petroleum Reserve (SPR).	U.S. Energy Information Administration (EIA)

Commonly quoted oil prices such as Brent or WTI are not used in the analysis the reason being that these prices reflect market outcomes on particular exchanges for particular types of oil. While they serve as benchmarks for the pricing of oil produced elsewhere, they do not reflect the cost refineries actually pay. For this reason, the U.S. refiner’s acquisition cost of imported crude oil is the closest proxy to a true global oil price. It is a volume-weighted price series based on the crude oil imported to the United States. Using this price in oil market VAR models is not uncommon in the literature (see e.g. Aastveit, Bjørnland and Thorsrud 2015; C. J. Baumeister and James D Hamilton 2017; Kilian 2009). For a discussion on the different oil prices and their uses, see Alquist, Kilian and Vigfusson (2013) and Kilian and Vigfusson (2011).

B Unit Root and Stationarity Tests of U.S. Inventories




Figure 6: Plot of the log change in U.S. inventories over the sample 2003:M01–2015:M12. Data is retrieved from U.S. Energy Information Administration (EIA).

Unit root and stationarity tests

Variable	Lags	ADF	PP	KPSS
Δ U.S. inventories	N/A		-7.65***	0.12
	2	-7.98***		
	4	-6.07***		
	6	-5.63***		
	8	-5.48***		
	10	-3.60***		
Critical values				
		-3.47	-3.47	0.74
		-2.88	-2.88	0.46
		-2.58	-2.58	0.35

Table 2: Test results from an Augmented Dickey-Fuller test (ADF), a Phillips-Perron test (PP) and a Kwiatkowski-Phillips-Schmidt-Shin test (KPSS) to check for unit root and stationarity in U.S. inventory data. The series is tested with a constant term and in first difference after taking the natural logarithm. The sample range is 2003:M01–2015:M12 with monthly observations. For the ADF and PP tests, the null hypothesis is that the series has a unit root while the null hypothesis for the KPSS test is that the series is stationary. The test results suggest stationarity and no unit root.

C Alternative SVAR model with U.S. crude oil production

In section 3.2, a new way of identifying supply shocks was implemented with U.S. crude oil imports. A relevant question is what the results would be if an alternative model that included U.S. crude oil production rather than imports was estimated. The only difference from earlier is the inclusion of U.S. crude oil production instead of oil imports.

The variable ordering is similar to the previous model with supply variables on top and United States production also with a short-run vertical supply curve.

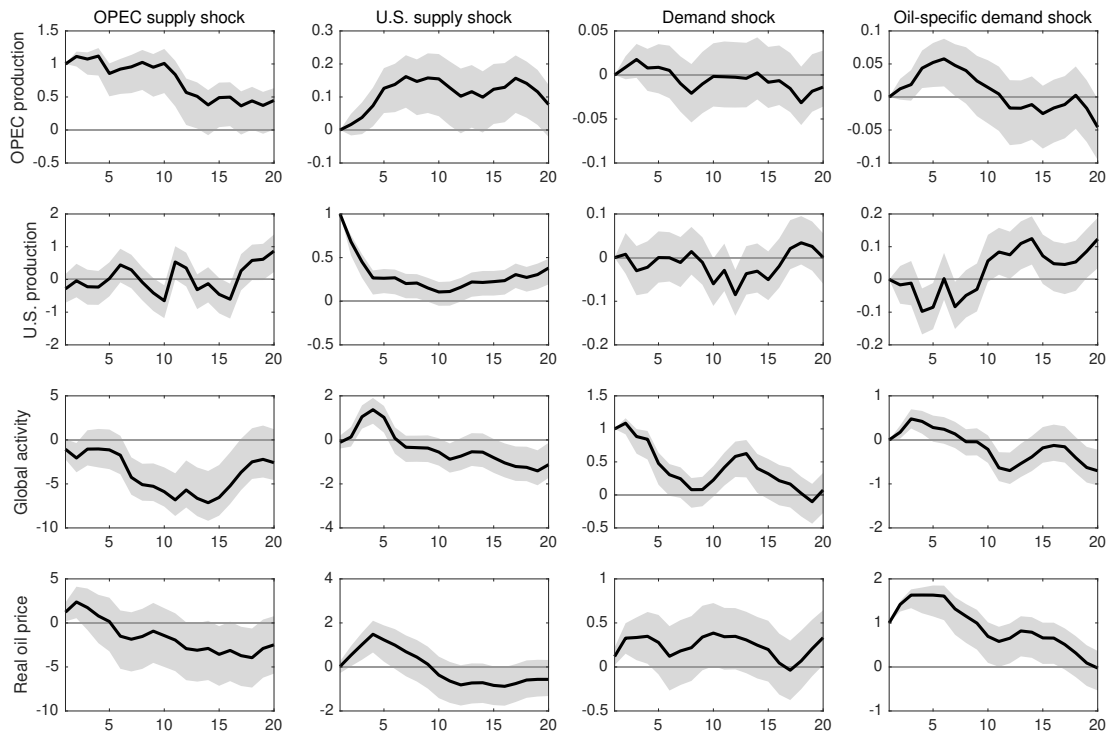


Figure 7: Impulse responses generated from the alternative SVAR model with U.S. production and the sample 2003:M01–2015:M12. They are all in levels of the variables. Shocks are normalised so that the response of the variables is 1 on impact, i.e. 1% for the OPEC supply shock, U.S. supply shock and aggregate demand shock while one log-unit for the oil price. The shaded areas represent 68% confidence bands calculated using a bootstrap with 10,000 draws.




Figure 8: Historical decomposition of the real price of oil derived from the alternative SVAR model with the 2003:M01–2015:M12 sample. The shaded areas correspond to the cumulative effects of the different shocks on the oil price.

Variance decomposition of the real price of oil

Shocks	Horizons			
	1 month	6 months	12 months	18 months
OPEC supply	2.53 (0.33; 8.71)	1.45 (1.57; 9.92)	2.63 (2.98; 17.42)	6.61 (4.87; 25.55)
U.S. supply	0.01 (0.07; 3.30)	9.26 (2.05; 22.33)	8.46 (4.71; 23.93)	10.01 (6.22; 26.37)
Aggregate demand	2.08 (0.24; 6.68)	4.31 (1.74; 14.81)	6.90 (3.32; 22.54)	6.80 (4.21; 22.91)
Oil-specific demand	95.37 (84.39; 96.76)	84.97 (61.16; 87.38)	82.00 (47.13; 78.63)	76.58 (38.82; 72.30)

Table 3: Variance decomposition (in percentages) of the real price of oil for different time horizons, generated from the alternative SVAR model with sample 2003:M01–2015:M12. The confidence intervals (in brackets) are at the 68% level and computed using a bootstrapping method with 10,000 draws.

D Section 5 — Robustness and sensitivity checks

D.1 Alternative restrictions

In my baseline model I assume that OPEC oil production cannot respond contemporaneously to U.S. supply shocks. To investigate whether my results are sensitive to this assumption, I report results for two alternative identification schemes.

D.1.1 Alternative ordering

A simple check of whether the results are sensitive to the restriction on s_{12} is to re-arrange the equations so that s_{12} now corresponds to the contemporaneous OPEC parameter in the United States imports equation. While often perceived as an infeasible exercise without any prior considerations and in large systems (see Kilian and Lütkepohl 2017), only two different models (orderings) are considered here as the other restrictions are taken as given following Kilian (2009). The results from this model are reported in figure 9. The first two columns show that the main results are largely insensitive to the ordering of the equations. OPEC production, no longer restricted to zero on impact, now responds significantly negative to a U.S. supply shock on impact and the oil price responds slightly less and slightly following a U.S. supply shock and an OPEC supply shock respectively. The qualitative interpretation of the results remains unchanged.

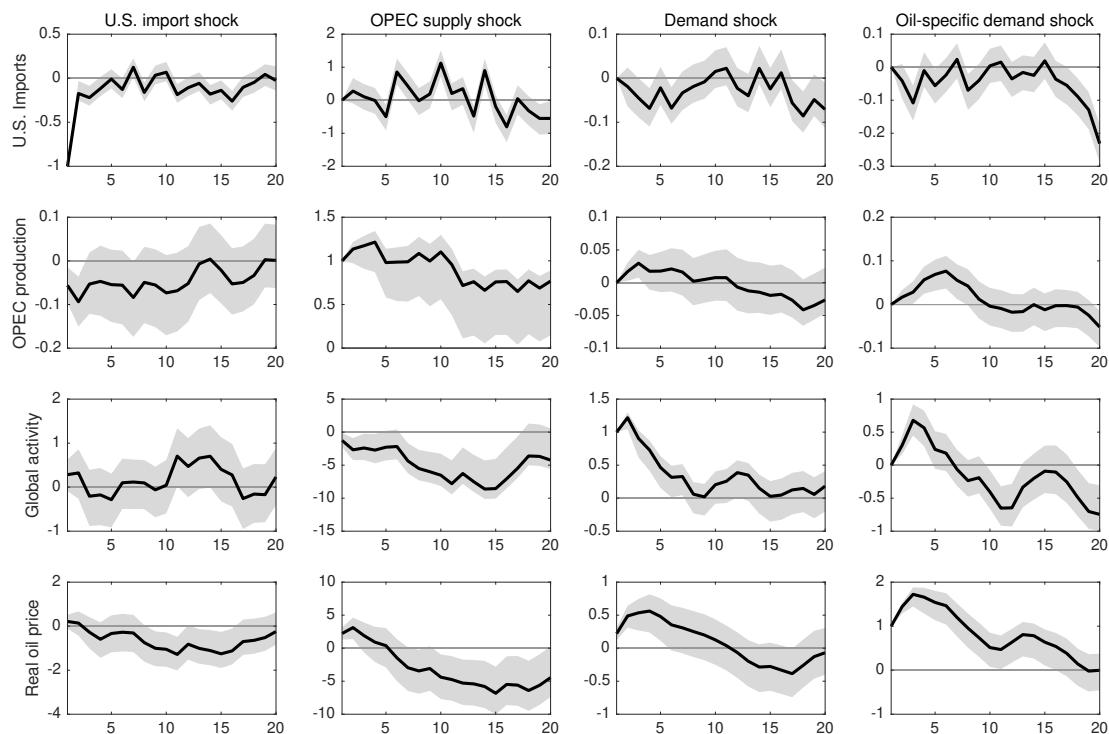


Figure 9: Impulse responses generated from the baseline model but where the ordering of the supply variables have been interchanged. The sample is 2003:M01–2015:M12. They are all in levels of the variables. Shocks are normalised so that the response of the variables is 1 on impact, i.e. 1% for the OPEC supply shock, U.S. supply shock and aggregate demand shock while one log-unit for the oil price. The shaded areas represent 68% confidence bands calculated using a bootstrap with 10,000 draws.

D.1.2 Mixed restrictions

One limitation of using a recursive identification scheme is that OPEC oil production and U.S. imports cannot both affect each other contemporaneously. To allow for this, I identify shocks using a combination of sign and zero-restrictions. Estimation of the reduced form model is done by applying Bayesian methods with diffuse priors.²⁰

Sign restrictions have become a popular way of identifying structural shocks and date back to Faust (1998), Canova and De Nicolò (2003) and Uhlig (2005). For simplicity, only the zero restriction on s_{12} will be relaxed. While not very common, imposing a mix of identifying restrictions has been done previously in Aastveit et al. (2015). Behar and Ritz (2017) show that a shift to a market-share strategy by OPEC can be optimal when facing competition from high-cost suppliers. OPEC thus will respond with the opposite sign to a change in U.S. imports. In other words, following a negative U.S. import shock reflecting a higher domestic supply of crude oil, OPEC will respond by increasing their own production.²¹ The restriction is imposed only on impact following Canova and Paustian (2011) to not to be more restrictive than necessary.

$$\begin{bmatrix} e^{\Delta opecprod} \\ e^{\Delta usimp^S} \\ e^{rea} \\ e^{lrpo} \end{bmatrix}_t = \begin{bmatrix} + & - & 0 & 0 \\ \times & + & 0 & 0 \\ \times & \times & + & 0 \\ \times & \times & \times & + \end{bmatrix} \begin{bmatrix} \varepsilon^{\Delta opecprod} \\ \varepsilon^{\Delta usimp^S} \\ \varepsilon^{rea} \\ \varepsilon^{lrpo} \end{bmatrix}_t \quad (6)$$

To produce impulse response functions that are consistent with the restrictions described in equation 6, a procedure based on the Rubio-Ramirez, Waggoner and Zha (2010) algorithm is implemented.

First, the Cholesky decomposition of the covariance matrix of the reduced form model is computed, $\Sigma_e = SS'$. With $\tilde{S} = SP$, $\Sigma_e = \tilde{S}\tilde{S}'$ also holds if P is an orthogonal matrix with the same dimensions as S . By drawing P randomly, we can have as many candidate draws of \tilde{S} we want. Each draw of the orthogonal matrix P is constructed in the following way: First draw a 2×2 matrix W with elements

²⁰The posterior distribution is then dominated by the likelihood function. Further, assuming normally distributed reduced form errors, the posterior will be Normal-Inverse-Wishart with mean and variance parameters corresponding to the OLS estimates of the parameters and covariance matrix of the reduced form model. See Kadiyala and Karlsson (1997) and Canova (2007) for details.

²¹In order to distinguish the OPEC supply shock from the U.S. import shock, it is necessary to restrict the response of OPEC to move in the opposite direction of U.S. imports. Depending on the signs of the unrestricted parameters, the OPEC supply shock and the U.S. import shock could be observationally equivalent if not.

$w \sim N(0,1)$. Then perform a QR-decomposition $W = QR$ with the property that $QQ' = Q'Q = I$. Finally, construct the orthogonal matrix $P = \begin{bmatrix} Q_{2 \times 2} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$.

Each draw of \tilde{S} is checked against the sign-restrictions posited in equation 6 and stored if it passes. If the restrictions are not satisfied, the draw is discarded and a new matrix P is drawn. Among all the accepted draws, the impulse response functions are computed.

The results are shown in figure 10. The solid lines correspond to the mean impulse responses at each horizon. Note that because of the identification strategy implemented, only the responses of the OPEC supply and the U.S. import shocks will differ from the baseline model. The responses to the OPEC supply shock are mostly unchanged. This is to be expected as the responses to this shock were not restricted in the baseline model. The main difference is that the response of the real oil price is now insignificant. OPEC responds to a negative U.S. import shock by increasing output. However, this increase is very small in magnitude. Global real activity responds negatively, but is only significant from zero in some periods. The main result from the baseline model is robust to the identifying restrictions as the oil price decreases following a negative U.S. import shock in a similar way as with a pure recursive identification scheme.

D.2 SVAR with U.S. inventories

In order to address the issue of inventories not being included in the main model, I present here an adaptation of the Kilian and Murphy (2014) model where an extra equation for adjusted U.S. imports is included. The model is estimated on a 2000:M01–2015:M12 sample compared to the baseline model as the extra equation introduces additional parameters to be estimated. Consider the following reduced form model

$$\mathbf{Y}_t = \mu + \sum_{p=1}^P A_p \mathbf{Y}_{t-p} + \mathbf{e}_t \quad (7)$$

where $\mathbf{Y}_t = [\Delta opecprod_t, \Delta usimp_t^S, rea_t, lrpo_t, \Delta usinv_t]'$. The first four variables remain the same as in the baseline model while $\Delta usinv_t$ is the percentage change of total U.S. crude oil inventories and the number of lags is 18. The estimation of the reduced form model in equation 7 is done by applying Bayesian methods with diffuse priors so that posterior distributions will be Normal-Inverse-Wishart.

A total of 5 structural shocks will be identified by impact sign-restrictions. The

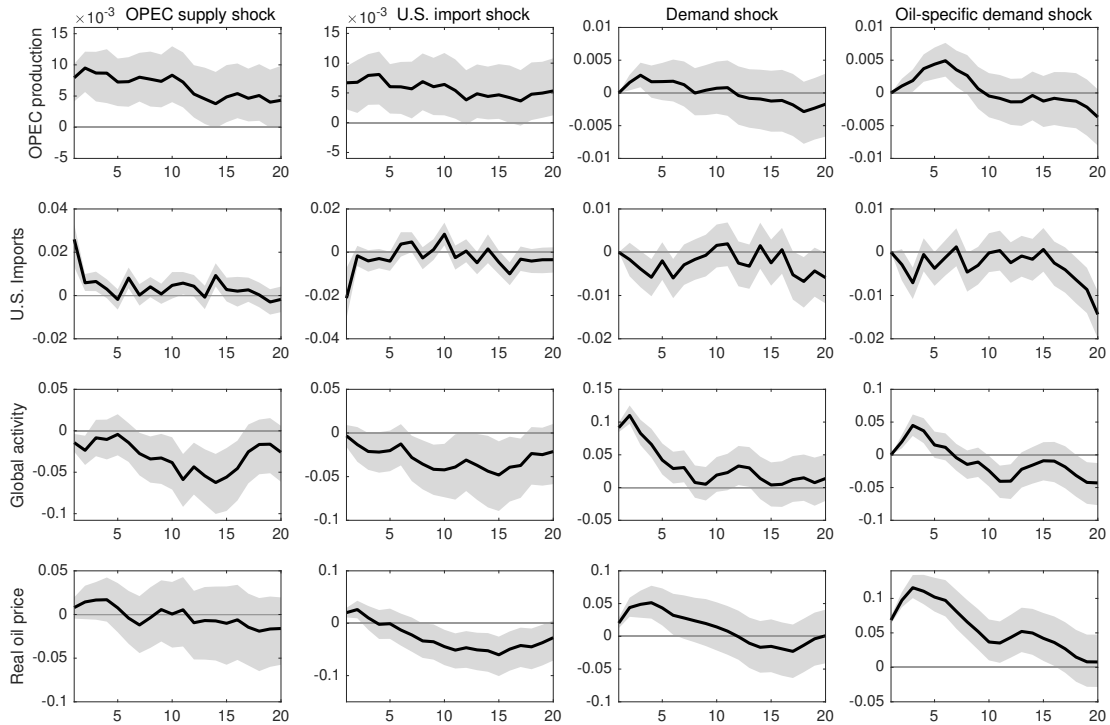


Figure 10: One standard deviation impulse responses generated from the imports model described in equation 6. The sample is 2003:M01–2015:M12. They are all in levels of the variables. The U.S. import shock is normalised to be negative. The solid line represents the median response at each horizon and the shaded areas represent 68th posterior probability regions of the estimated impulse responses.

first three, the OPEC supply shock, the U.S. import shock and the aggregate demand shock have similar interpretations as in the baseline model. The OPEC supply shock is an unexpected increase in OPEC production that lowers oil prices and increases real activity. As prices fall, U.S. producers will reduce output and U.S. refineries would have to import more crude oil to cover petroleum product demand. The effect on inventories is ambiguous, but the logic is the same as in Kilian and Murphy (2014). Storage demand might fall as the OPEC supply shock triggers a predictable decrease in prices, but it might also increase to smooth out consumption. The same logic applies to the U.S. import shock which represents a sudden change in the domestic availability of crude oil that moves U.S. demand for imports. However, as has been shown by Behar and Ritz (2017), it can be optimal for OPEC to switch to a market-share strategy when facing competition from high-cost suppliers and OPEC will respond by increasing output following a negative U.S. import shock. This restriction is necessary in order to separately identify the OPEC supply shock and the U.S. import shock with sign-restrictions. A positive aggregate demand shock raises oil price and stimulates oil production among OPEC producers and in the United States. The effect on inventories are

	OPEC supply shock	U.S. import shock	Aggregate demand shock	Speculative demand shock
OPEC production	+	+	+	+
U.S. imports	+	-	-	-
Real economic activity	+	+	+	-
Real price of oil	-	-	+	+
U.S. inventories				+

Table 4: Sign-restrictions imposed on impact responses

again ambiguous. Kilian and Murphy (2014) named the next shock *the speculative demand shock* and is the parallel to the oil-specific demand shock in the baseline model. It captures changes in the expectations regarding the future availability of crude oil. A speculative demand shock that raises the oil price in anticipation of future oil shortages will induce market agents to put oil in storage before the actual shortfall has materialised, hence *speculative*. Such a shock will make producers increase output while consumption will fall to accommodate the storage of oil. The final structural shock is a residual shock that will account for anything not captured by the other shocks. Following Kilian and Murphy, I do not impose any restrictions on the responses from this shock nor do I report any results stemming from this shock. Table 4 summarises the impact sign-restrictions implied by the above description of the properties of each structural shock. The U.S. import shock is normalised to decrease the real price of oil.

In addition to these sign-restrictions, I follow the example of Kilian and Murphy (2014) and impose a dynamic restriction on the shape of the response of the real price of oil to a OPEC supply shock. This is to ensure that models, where a positive supply shock do not result in lower prices and higher quantities, are ruled out.

The estimation procedure follows the standard Rubio-Ramirez, Waggoner and Zha (2010) algorithm. With the estimates of the reduced form model and the covariance matrix Σ_e , let the reduced form errors be written as a linear combination $e_t = S\varepsilon_t$ where S is the $N \times N$ contemporaneous multiplier matrix and ε_t the vector of structural shocks with $\mathbb{E}[\varepsilon_t\varepsilon_t'] = I$. To obtain an estimate of S , perform an eigendecomposition of $\Sigma_e = EDE'$ and let $S = ED^{0.5}$ so that $\Sigma_e = SS'$. With $\tilde{S} = SP$, $\Sigma_e = \tilde{S}\tilde{S}'$ will also hold as long as P is an orthogonal $N \times N$ matrix. This means that we can construct as many candidate draws of \tilde{S} as we want by drawing matrices P randomly. Each draw of \tilde{S} is checked against the sign-restrictions and stored if in compliance. P will be drawn in the following way: First draw an $N \times N$ matrix K of standard normal variables and perform a QR decomposition such that $K = QR$ where $QQ' = Q'Q = I$ and let $P = Q'$.

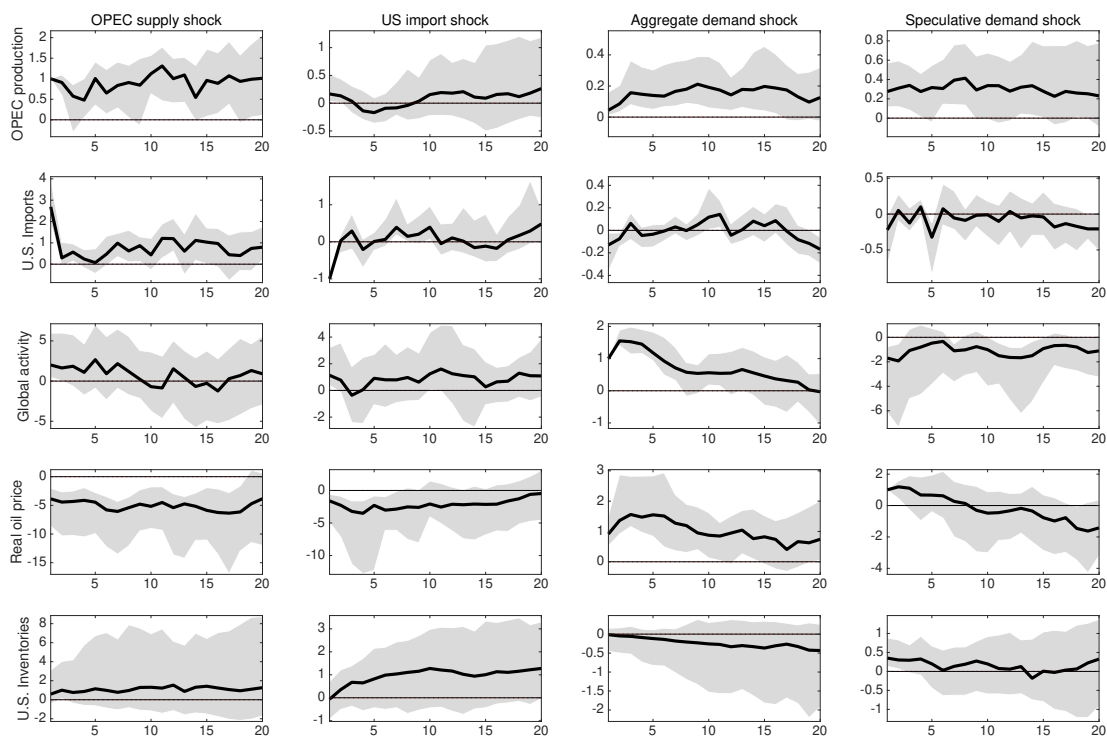


Figure 11: Impulse responses generated from the model with inventories described in Equation 7 and with impact sign-restrictions as described in Table 4. The sample is 2000:M01–2015:M12. They are all in levels of the variables. The shocks are normalised so that the response of the variables is 1 on impact. The U.S. import shock is normalised to be negative. The solid line represents the median response at each horizon and the shaded areas represent 68th posterior probability regions of the estimated impulse responses.

Finally, the impulse response functions are computed and checked against the dynamic sign-restriction.

The impulse responses are shown in Figure 11. Overall, the conclusions drawn based on the baseline model is not changed. However, there are some significant differences stemming from the new identification scheme as well as the inclusion of inventories in the model. I will discuss each in turn. Following a OPEC supply shock, the oil price falls immediately and is persistently negative for over 20 months. In the baseline results where there are no restrictions on this response, the oil price is positive initially. Global activity increases as oil prices fall for a couple of months contrary to the baseline model where activity is persistently negative. Restricting the oil price to respond negatively on impact may have remedied this contradictory result. Demand for U.S. imports is persistently positive and the level of inventories responds positively on impact. This suggests that the consumption smoothing effect dominates in this case.

The U.S. import shock that reduces U.S. demand for crude oil imports due to

higher domestic supply causes the oil price to fall immediately. The increase is only significantly different from zero for 6 months compared to 16 in the baseline model. This is possibly due to global activity responding positively for the first 10 months rather than negatively as in the baseline model. In fact, relaxing the impact restriction on the response of global activity brings these responses closer to the baseline model results.

The responses from the aggregate demand shock are similar to those in the baseline model. The speculative demand shock that raises the oil price and inventory demand has much less persistent effect on the oil price than the oil-specific demand shock in the baseline model. This might stem from global activity responding negatively on impact rather than positive as in the baseline model and in Kilian (2009).

Based on the above results, controlling for inventories does not seem to produce results that contradicts those of the baseline model. The identification scheme however, has put more structure on the variable responses, especially the responses of global real activity, and thus muted the oil price responses to many of the structural shocks.